

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

AIX

5.0 Maintenance Pack 3



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 5.0MP3

Documentation version 5.0MP3.0

Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

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

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

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
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

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

www.symantec.com/techsupp/

Contacting Technical Support

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

www.symantec.com/business/support/assistance_care.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:

<https://licensing.symantec.com>

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/assistance_care.jsp

Customer Service is available to assist with 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 maintenance 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

Maintenance agreement resources

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

Asia-Pacific and Japan	contractsadmin@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4
Section 1 SF Oracle RAC requirements and planning	23
Chapter 1 Introducing Veritas Storage Foundation for Oracle RAC	25
About Veritas Storage Foundation for Oracle RAC	25
SF Oracle RAC basics	26
About SF Oracle RAC components	27
About SF Oracle RAC optional components and features	28
Veritas Installation Assessment Service	28
Symantec Product Authentication Service	29
Veritas Cluster Server Management Console	29
Notification for VCS events	29
Global clusters	29
Veritas Volume Replicator	29
Veritas Storage Foundation for Databases (SFDB) repository	30
Chapter 2 Installation planning and requirements	33
Important preinstallation information	33
Installation requirements	34
Hardware requirements	34
Supported software	38
Supported operating systems	39
Recommended Oracle configuration	40
Recommended database storage configuration	41
SF Oracle RAC cluster setup	42
Basic SF Oracle RAC cluster setup	42
SF Oracle RAC cluster optional features setup	44
SF Oracle RAC global cluster setup	45
Overview of SF Oracle RAC installation and configuration tasks	46
Tasks to prepare for optional SF Oracle RAC features	47
Tasks to prepare for SF Oracle RAC installation	49

Tasks for SF Oracle RAC installation and configuration	49
Tasks for installing Oracle RAC and creating an Oracle database	50
Tasks for setting up VCS to manage Oracle RAC resources	51
Tasks for setting up database management features for SF Oracle RAC	51
Tasks for setting up SF Oracle RAC global clusters (optional)	52
Required information for installing and configuring SF Oracle RAC	54

Section 2 SF Oracle RAC installation and configuration 55

Chapter 3 Preparing to install and configure SF Oracle RAC	57
About preparing to install and configure SF Oracle RAC	57
Preparing to configure the clusters in secure mode	58
Installing the root broker for the security infrastructure	62
Creating authentication broker accounts on root broker system	63
Creating encrypted files for the security infrastructure	64
Preparing the installation system for the security infrastructure	66
Preparing to install SF Oracle RAC	67
Setting umask before installation	67
Synchronizing time settings on cluster nodes	68
Setting up inter-system communication	68
Setting up shared storage	73
Mounting the product disc	73
Setting the MANPATH variable	74
Setting the PATH variable	74
Obtaining SF Oracle RAC license keys	75
Removing pre-existing license keys	76
Obtaining and installing license keys for a global cluster	77
Verifying the systems before installation	78

Chapter 4 Installing SF Oracle RAC	81
About installing SF Oracle RAC	81
Invoking SF Oracle RAC installation program	82
Mounting the product disc	82
Starting software installation	82

	Specifying systems for installation	84
	Licensing SF Oracle RAC	84
	Installing SF Oracle RAC filesets	86
Chapter 5	Configuring SF Oracle RAC	87
	About configuring SF Oracle RAC	87
	Invoking SF Oracle RAC configuration program	89
	Starting software configuration	89
	Specifying systems for configuration	91
	Choosing the configuration task	91
	Performing basic system checks	92
	Preparing private interconnects for LLT	93
	Configuring SF Oracle RAC components	95
	Configuring the SF Oracle RAC cluster	96
	Setting permissions for database administration	104
	Configuring the Cluster Volume Manager	105
	Configuring VVR on each node	106
	Creating SF Oracle RAC configuration files	108
	Configuring standalone hosts for Storage Foundation	108
	Starting SF Oracle RAC processes	109
	Enabling LDAP authentication for clusters that run in secure mode	110
Chapter 6	Configuring SF Oracle RAC clusters for data integrity	113
	About configuring SF Oracle RAC clusters for data integrity	113
	About I/O fencing components	114
	About data disks	114
	About coordination points	115
	About setting up I/O fencing	116
	About setting up shared storage for I/O fencing	119
	Preparing to configure I/O fencing	119
	Initializing disks as VxVM disks	119
	Identifying disks to use as coordinator disks	121
	Checking shared disks for I/O fencing	122
	Setting up I/O fencing	125
	Setting up coordinator disk groups	125
	Configuring I/O fencing	126
	Modifying VCS configuration to use I/O fencing	127
	Starting SF Oracle RAC on all nodes	128
	Verifying I/O fencing configuration	130
	Removing permissions for communication	130

Chapter 7	Verifying cluster installation	131
	Verifying SF Oracle RAC installation using VCS configuration file	131
	About the LLT and GAB configuration files	131
	About VCS configuration file after SF Oracle RAC installation	133
	Sample main.cf file after SF Oracle RAC installation	134
	Verifying LLT, GAB, and cluster operation	136
	Verifying LLT	136
	Verifying GAB	139
	Verifying the cluster	140
	Verifying the cluster nodes	141
Chapter 8	Upgrading to SF Oracle RAC 5.0 MP3	145
	Supported upgrade options for SF Oracle RAC 5.0 MP3	145
	Overview of tasks for upgrading SF Oracle RAC	146
	Preparing to upgrade SF Oracle RAC to 5.0 MP3	147
	Upgrading to SF Oracle RAC 5.0 MP3	151
	Upgrading from SF Oracle RAC 5.0x using the installmp installer	152
	Upgrading from pre-5.0 versions of SF Oracle RAC using the installsrac installer	154
	Performing post-upgrade tasks for SF Oracle RAC 5.0 MP3	157
	Upgrading CVM disk group version	163
	Minimum downtime upgrade from SF Oracle RAC 5.0 to 5.0 MP3	164
	Upgrading clusters with Oracle RAC 10g or Oracle RAC 11g R1	164
	Upgrading clusters with Oracle RAC 9iR2	178
	Upgrading licenses to SF Oracle RAC level	187
	Upgrading to SF Oracle RAC from Storage Foundation for Oracle	187
	Upgrading to SF Oracle RAC from Storage Foundation HA for Oracle	187
	Upgrading to SF Oracle RAC from Storage Foundation Cluster File System	188
	Upgrading the repository database	189

Section 3	Oracle RAC 10g installation in an SF Oracle RAC environment	193
Chapter 9	Preparing to install Oracle RAC 10g	195
	About preparing to install Oracle RAC 10g	195
	Pre-installation tasks	197
	Setting the system parameters	197
	Identifying storage for Oracle RAC 10g components	198
	Launch the SF Oracle RAC configuration program	198
	Creating Oracle user and groups	201
	Setting up Oracle user equivalence for RSH and RCP	206
	Setting Oracle Clusterware oracle user capabilities	207
	Configuring shell limits for Oracle user	207
	Configuring kernel parameter for processes per user	208
	Tuning the Virtual Memory Manager (VMM)	209
	Verifying the user "nobody" exists	209
	Configuring private IP addresses for CRS	210
	Identifying the public virtual IP addresses for use by Oracle	215
	Creating Oracle disk groups, volumes, and mount points	216
	Verifying SF Oracle RAC resources on all nodes	228
	Verifying GAB port membership	228
Chapter 10	Installing Oracle RAC 10g	231
	About installing Oracle RAC 10g	231
	Installing Oracle Clusterware (CRS)	233
	Installing Oracle RAC 10g database binaries	238
	Verifying the Oracle RAC 10g CRS and database installation	243
	Completing Oracle RAC 10g post-installation tasks	244
	Adding Oracle RAC 10g patchsets	244
	Relinking the SF Oracle RAC libraries to Oracle RAC 10g	249
	Upgrading Oracle databases	251
	Creating the Oracle RAC 10g database	251
	Configuring the Oracle RAC 10g service group in a VCS configuration	251
Chapter 11	Configuring VCS service groups for Oracle RAC 10g	253
	About VCS service group for Oracle RAC 10g dependencies	253
	VCS service group dependencies with or without the VCS Oracle agent	254

Automatic database starting prevention	259
Configuring VCS resources for Oracle	260
Editing the main.cf file to configure the CVM and Oracle service groups	260
Configuring the CVM and Oracle service groups using the CLI	262
Location of VCS log files	263

Chapter 12 Upgrading to Oracle 10g and migrating the database

About upgrading and migrating to Oracle RAC 10g software	265
Overview of upgrade and migration tasks	265
Migrating from Oracle RAC 9i R2 to Oracle RAC 10g	266
Performing pre-upgrade tasks	266
Installing Oracle RAC 10g	267
Migrating an existing Oracle RAC 9i R2 database to Oracle RAC 10g	268
Performing post-upgrade tasks	269
Migrating from Oracle RAC 10gR1 to Oracle RAC 10gR2	269
Performing pre-upgrade tasks	269
Installing Oracle RAC 10gR2	270
Migrating the existing Oracle RAC 10g R1 database to Oracle RAC 10g R2	271
Performing post-upgrade tasks	271

Section 4 Oracle RAC 11g installation in an SF Oracle RAC environment

Chapter 13 Preparing to install Oracle RAC 11g

About preparing to install Oracle RAC 11g	275
Performing pre-installation operations	277
Setting the system parameters	277
Identify storage for the Oracle RAC 11g components	278
Launching the SF Oracle RAC configuration program and menu	278
Creating Oracle user and groups	281
Setting up Oracle user equivalence for RSH and RCP	286
Setting Oracle Clusterware oracle user capabilities	288
Configuring shell limits	288
Configuring the user process parameter	289
Configuring Networks Tuning Parameters	290

	Tuning the Virtual Memory Manager (VMM)	290
	Verifying the user "nobody" exists	291
	Configuring CRS and UDP private IP addresses for failover	291
	Manually configuring MultiPrivNIC for UDP cache fusion	299
	Identifying the public virtual IP addresses for use by Oracle	301
	Creating Oracle disk groups, volumes, and mount points	302
	Verifying SF Oracle RAC resources on all nodes	313
	Verifying GAB port membership	313
Chapter 14	Installing Oracle RAC 11g	315
	About installing Oracle RAC 11g	315
	Installing Oracle Clusterware (CRS) using the installsrac command	317
	Installing Oracle Clusterware (CRS)	321
	Installing Oracle RAC 11g database binaries	323
	Verifying the Oracle RAC 11g CRS and database installation	328
	Completing Oracle RAC 11g post-installation tasks	329
	Adding Oracle RAC 11g patches or patchsets	329
	Relinking the SF Oracle RAC libraries to Oracle RAC 11g	330
	Creating the Oracle RAC 11g database	334
	Adding UDP IP addresses to the Oracle init file	334
	Upgrading databases	336
	Configuring the Oracle RAC 11g service group in a VCS configuration	336
Chapter 15	Installing Oracle RAC 11g manually	337
	About installing Oracle RAC 11g manually	337
	Installing Oracle RAC 11g manually	338
	Configuring MultiPrivNIC for CRS/UDP cache fusion	339
	Setting up environment variables	339
	Linking Veritas Membership library	341
	Installing the Oracle Clusterware (CRS)	343
	Installing the Oracle RAC 11g database binaries	345
	Adding Oracle RAC 11g patches or patchsets	346
	Link the ODM library	346
	Add CRS under VCS control	347
	Creating the Oracle RAC 11g database	348
	Configure cluster interconnects	348
	Upgrading databases	349
	Configuring the Oracle RAC 11g service group in a VCS configuration	349

Chapter 16	Configuring VCS service groups for Oracle RAC 11g	351
	About VCS service group for Oracle RAC 11g dependencies	351
	VCS service group dependencies with or without the VCS Oracle agent	352
	Automatic database starting prevention	356
	Configuring VCS resources for Oracle	357
	Editing the main.cf file to configure the Oracle service groups	357
	Configuring the Oracle service groups using the CLI	359
	Location of VCS log files	359
Chapter 17	Upgrading to Oracle RAC 11g and migrating the database	361
	About upgrading and migrating to Oracle RAC 11gR1 software	361
	Overview of upgrade and migration tasks	362
	Upgrading from Oracle RAC 9iR2 to Oracle RAC 11gR1	362
	Performing pre-upgrade tasks for Oracle RAC 11g	362
	Installing Oracle RAC 11gR1	364
	Migrating an existing Oracle RAC 9iR2 database to Oracle RAC 11g	365
	Performing post-upgrade tasks Oracle RAC 11gR1	365
	Upgrading from Oracle RAC 10g to Oracle RAC 11gR1	365
	Preparing to upgrade to Oracle RAC 11gR1	366
	Installing Oracle RAC 11g R1	366
	Migrating the existing Oracle RAC 10g database to Oracle RAC 11gR1	367
	Performing post-upgrade tasks	367
Section 5	Adding and removing SF Oracle RAC nodes	369
Chapter 18	Adding a node to SF Oracle RAC clusters using Oracle RAC 10g	371
	About adding a node to an Oracle RAC 10g cluster	371
	SF Oracle RAC add node procedures	372
	Modifying /etc/pse.conf to enable the dlpi driver	372
	Checking system requirements for new node	373
	Physically adding a new system to the cluster	373

	Installing Veritas Storage Foundation for Oracle RAC on the new system	373
	Starting Volume Manager	375
	Verifying the existing security setup on the node	375
	Configuring LLT, GAB, VCSMM, and VXFEN drivers	377
	Adding a node in a VxSS group	380
	Preparing the new node for installing Oracle	381
	Setting the system parameters	382
	Creating Oracle user and groups	382
	Setting up Oracle user equivalence for RSH and RCP	382
	Setting Oracle Clusterware oracle user capabilities	382
	Configuring shell limits for Oracle user	383
	Configuring kernel parameter for processes per user	383
	Tuning the Virtual Memory Manager (VMM)	383
	Verifying the user "nobody" exists	383
	Configuring private IP addresses for CRS	383
	Identifying the public virtual IP addresses for use by Oracle	384
	Creating Oracle disk groups, volumes, and mount points	385
	Preparing \$CRS_HOME and \$ORACLE_HOME on new node	385
	Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node	388
	Configuring CVM	388
	Linking Veritas membership library	389
	Using the Oracle RAC 10g add node procedure	391
	Relinking the SF Oracle RAC libraries to Oracle RAC 10g	391
	Adding the new system to the SFDB repository	391
Chapter 19	Removing a node from SF Oracle RAC clusters using Oracle RAC 10g	393
	About removing a node from an Oracle RAC 10g cluster	393
	Using the Oracle remove node procedure	394
	Removing SF Oracle RAC using the uninstallsfrac utility	394
	Editing VCS configuration files on existing nodes	396
Chapter 20	Adding a node to SF Oracle RAC clusters using Oracle RAC 11g	399
	About adding a node to an Oracle RAC 11g cluster	399
	Using the SF Oracle RAC add node procedures	399
	Modifying /etc/pse.conf to enable the dlpi driver	400
	Checking system requirements for new node	401
	Physically adding a new system to the cluster	401

Installing Veritas Storage Foundation for Oracle RAC on the new system	401
Starting Volume Manager	403
Verifying the existing security setup on the node	403
Configuring LLT, GAB, VCSMM, and VXFEN drivers	405
Adding a node in a VxSS group	408
Preparing the new node for installing Oracle	409
Setting the system parameters	410
Creating Oracle user and groups	410
Setting up Oracle user equivalence for RSH and RCP	410
Setting up Oracle Clusterware oracle user capabilities	410
Configuring shell limits for Oracle user	411
Configuring kernel parameters for processes per user	411
Tuning the Virtual Memory Manager (VMM)	411
Verifying the user "nobody" exists	411
Configuring private IP addresses for CRS	411
Identifying the public virtual IP addresses for use by Oracle	413
Creating Oracle disk groups, volumes, and mount points	413
Preparing \$CRS_HOME and \$ORACLE_HOME on new node	413
Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node	416
Configuring CVM	416
Linking Veritas membership library	417
Using the Oracle RAC 11g add node procedure	419
Relinking the SF Oracle RAC libraries to Oracle	419
Adding the new system to the SFDB repository	419
 Chapter 21	
Removing a node from SF Oracle RAC clusters using Oracle RAC 11g	421
About removing a node from an Oracle RAC 11g cluster	421
Using the Oracle remove node procedure	422
Removing SF Oracle RAC using the uninstallsfrac utility	422
Editing VCS configuration files on existing nodes	424

Section 6	Configuring an SF Oracle RAC global cluster	429
Chapter 22	Preparing for a replicated SF Oracle RAC global cluster	431
	Replication in the SF Oracle RAC environment	431
	Requirements for SF Oracle RAC global clusters	432
	Supported software and hardware for SF Oracle RAC	432
	Supported replication technologies for SF Oracle RAC	432
	Global cluster configuration overview	434
	Configuring an SF Oracle RAC cluster at the primary site	435
	Configuring an SF Oracle RAC cluster at the secondary site	438
	Configuring replication on clusters at both sites	441
	Configuring VCS service groups for global clusters	442
	Testing a global cluster configuration	443
Chapter 23	Configuring a global cluster using VVR	445
	About configuring global clustering using VVR	445
	Preparing clusters for replication using VVR	446
	Adding the VVR resource types to the VCS configuration	447
	Configuring global clustering using VVR	448
	Defining the remote cluster and heartbeat Cluster Objects	451
	Setting up replication using VVR	454
	Creating the SRL volume on the primary site	454
	Setting up replication objects on the primary site	455
	Configuring replication for the secondary site	456
	Starting replication of Oracle RAC database volume	460
	Configuring VCS to replicate the database volume using VVR	463
	About modifying the VCS configuration for replication	464
	Configuration examples before and after modification	465
	Dependencies	466
	Modifying the VCS Configuration on the Primary Site	468
	Modifying the VCS Configuration on the Secondary Site	472
	Starting VCS on all nodes in both clusters	476

Section 7	Uninstalling SF Oracle RAC	479
Chapter 24	Uninstalling SF Oracle RAC from an Oracle RAC 10g cluster	481
	About uninstalling SF Oracle RAC from an Oracle RAC 10g cluster	481
	Preparing to uninstall SF Oracle RAC from an Oracle RAC 10g cluster	482
	Stopping Oracle instances	483
	Backing up the Oracle database	484
	Moving or removing the Oracle database (optional)	484
	Uninstalling Oracle RAC 10g (optional)	484
	Removing repository database	485
	Stopping applications that use CFS (outside of VCS control)	486
	Unmounting CFS file systems (outside of VCS control)	486
	Stopping applications that use VxFS (outside of VCS control)	487
	Unmounting VxFS file systems (outside of VCS control)	487
	Stopping VCS	487
	Removing SF Oracle RAC from an Oracle RAC 10g cluster	488
	Removing the SF Oracle RAC filesets	488
	Removing other configuration files (optional)	490
	Rebooting the nodes	491
Chapter 25	Uninstalling SF Oracle RAC from an Oracle RAC 11g cluster	493
	About uninstalling SF Oracle RAC from an Oracle RAC 11g cluster	493
	Preparing to uninstall SF Oracle RAC from an Oracle RAC 11g cluster	494
	Stopping Oracle instances	495
	Backing up the Oracle database	496
	Moving or removing the Oracle database (optional)	496
	Uninstalling Oracle RAC 11g (optional)	496
	Removing repository database	497
	Stopping the applications that uses CFS (outside of VCS control)	498
	Unmounting CFS file systems (outside of VCS control)	498
	Stopping the applications that use VxFS (outside of VCS control)	499
	Unmounting VxFS file systems (outside of VCS control)	499
	Stopping VCS	499

	Removing SF Oracle RAC from an Oracle RAC 11g cluster	500
	Removing the SF Oracle RAC filesets	500
	Removing other configuration files (optional)	502
	Rebooting the nodes	503
Section 8	Reference	505
Appendix A	Sample configuration files	507
	About sample main.cf files	507
	Sample main.cf files for Oracle configurations	507
	Sample main.cf for Oracle RAC 10g without Oracle agent	508
	Sample main.cf for Oracle RAC 10g with Oracle agent	511
	Sample main.cf for Oracle RAC 10g with Oracle agent and CMC	514
	Sample main.cf for Oracle RAC 10g without the Oracle agent and with VxSS	518
	Sample main.cf for Oracle RAC 11g without Oracle agent	522
	Sample main.cf for Oracle RAC 11g with Oracle agent	525
	Sample main.cf files for Oracle replication configurations	529
	Sample main.cf for Oracle RAC 10g and CVM/VVR (primary replication site)	529
	Sample main.cf for Oracle RAC 10g and CVM/VVR (secondary replication site)	534
	Sample main.cf for Oracle RAC 10g for hardware-based replication primary site	538
	Sample main.cf for Oracle RAC 10g for hardware-based replication secondary site	543
	Sample main.cf files for adding and removing nodes	549
	Sample main.cf for adding and removing an Oracle RAC 10g node	549
	Sample main.cf for adding and removing an Oracle RAC 11g node	552
Appendix B	SF Oracle RAC installation and configuration sample values	559
	SF Oracle RAC installation and configuration information	559
	SF Oracle RAC filesets information	561
	Veritas Cluster Server component information	561
	Cluster Volume Manager information	562
	SF Oracle RAC secure mode cluster information	563
	SF Oracle RAC add user information	563

SMTP email notification information	564
SNMP trap notification information	565
Global cluster information	565
Veritas Volume Replicator information	566
I/O fencing information	566
Installation and configuration worksheets	567
SF Oracle RAC worksheet	567
Oracle worksheet	568
Replicated cluster using VVR worksheet	571
Replicated cluster using SRDF worksheet	572

Appendix C Response Files 575

About Response files	575
Veritas Cluster Server response files	576
Response file installation procedures	576
Re-installation using the response files	576
Installation using the response file on other clusters	580
Response file syntax	583
Install-only response file example	583
Configure response file example	584
Response file variable definitions	585

Appendix D Automatic Storage Management 593

About ASM	593
Veritas CVM and ASM integration	594
Dynamic MultiPathing	594
ASM on CVM rules	594
Supported SF Oracle RAC and ASM configurations	595
Requirements for the supported SF Oracle RAC and ASM configurations	597
ASM monitoring utilities	597
CVM and ASM configuration procedure	597
Configure the CVM disk group and volume for the ASM disk group	598
Install Automatic Storage Management in an ASM home directory	599
Use the Oracle DBCA procedure to create ASM disk group and instances	600
Verify your ASM installation	601
Relinking the SF Oracle RAC libraries to Oracle	602
Unlinking the SF Oracle RAC ODM library	602
ASM instance startup and shutdown commands	604

	Adding additional storage to an ASM disk group with CVM	605
	Veritas ASM agents	605
	Veritas CVM and Oracle service group with ASM	606
	Veritas CVM and an Oracle service group with ASM	606
	Veritas CVM and ASM main.cf file	609
Appendix E	Creating a test database	617
	About creating a test database	617
	Creating a database for Oracle	617
	Creating the database tablespaces for Oracle on raw volumes	618
	Creating the database tablespaces for Oracle on CFS	619
Appendix F	High availability agent information	621
	About agents	621
	VCS agents included within SF Oracle RAC	622
	Other agents included within SF Oracle RAC	622
	CVMCluster agent	623
	Entry points for CVMCluster agent	623
	Attribute definition for CVMCluster agent	624
	CVMCluster agent type definition	624
	CVMCluster agent sample configuration	625
	CVMVxconfigd agent	625
	Entry points for CVMVxconfigd agent	625
	Attribute definition for CVMVxconfigd agent	626
	CVMVxconfigd agent type definition	626
	CVMVxconfigd agent sample configuration	626
	CVMVolDg agent	627
	Entry points for CVMVolDg agent	627
	Attribute definition for CVMVolDg agent	628
	CVMVolDg agent type definition	628
	CVMVolDg agent sample configuration	629
	CFSMount agent	629
	Entry points for CFSMount agent	629
	Attribute definition for CFSMount agent	629
	CFSMount agent type definition	630
	CFSMount agent sample configuration	631
	PrivNIC agent	631
	Entry point for PrivNIC agent	631
	Attribute definition for PrivNIC agent	632
	PrivNIC agent type definition	634
	PrivNIC agent sample configuration	635
	MultiPrivNIC agent	635

MultiPrivNIC agent entry point	636
MultiPrivNIC agent attribute definition	636
Resource type definition for MultiPrivNIC agent	637
Sample configuration for MultiPrivNIC agent	637
CSSD agent	638
VCS agents for Oracle	639
Oracle agent functions	640
Oracle resource type	646
Netlsnr agent functions	653
Netlsnr resource type	654
ASMInst agent functions	657
ASMInst resource type	658
ASMDG agent functions	660
ASMDG resource type	661
Glossary	665
Index	669

SF Oracle RAC requirements and planning

- [Introducing Veritas Storage Foundation for Oracle RAC](#)
- [Installation planning and requirements](#)

Introducing Veritas Storage Foundation for Oracle RAC

This chapter includes the following topics:

- [About Veritas Storage Foundation for Oracle RAC](#)
- [SF Oracle RAC basics](#)
- [About SF Oracle RAC components](#)
- [About SF Oracle RAC optional components and features](#)

About Veritas Storage Foundation for Oracle RAC

Veritas Storage Foundation™ for Oracle® Real Application Clusters (SF Oracle RAC) is a storage management and clustering solution that enables the user to perform the following tasks:

- Create a standard approach toward application and database management in datacenters. While other clusterware can only work with an Oracle database, SF Oracle RAC provides flexible support for many types of applications and databases. Since SF Oracle RAC incorporates existing Veritas storage management and clustering technologies, administrators can apply existing expertise of Veritas technologies toward this product.
- Set up an infrastructure for Oracle RAC that simplifies database management while fully integrating with Oracle Clusterware.
- Enhance scalability and availability with access to multiple Oracle RAC instances per database in a cluster.
- Transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites.

- Back up and recover databases 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.
- Prevent data corruption at the storage layer with robust split-brain protection.
- Increase scalability with high throughput and low latency technology for use by Oracle Cache Fusion.
- Share all types of files, in addition to Oracle database files, across nodes.
- Increase availability and performance with dynamic multi-pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the HBAs and SAN switches.
- Model and test cluster configurations without affecting production systems using the simulator and fire drill clustering technologies.
- Optimize I/O performance through storage mapping technologies and tunable attributes.

SF Oracle RAC basics

Real Application Clusters (RAC) is a parallel database environment that takes advantage of the processing power of multiple computers. The Oracle database is the physical data stored in tablespaces on disk. An Oracle instance is a set of processes and shared memory that provides access to the physical database. Specifically, the instance involves server processes and background processes. The server processes act on behalf of clients to read data into shared memory and make modifications to it. The background processes write changed data to disk.

In traditional environments, only one instance accesses a database at a specific time. SF Oracle RAC enables all nodes to concurrently run Oracle instances and execute transactions against the same database. This software coordinates access to the shared data for each node to provide consistency and integrity. Each node adds its processing power to the cluster as a whole and can increase overall throughput or performance.

At a conceptual level, SF Oracle RAC is a cluster that manages applications (instances), networking, and the storage components that use the resources contained in service groups. SF Oracle RAC clusters have many of the same properties as Veritas Cluster Server (VCS) clusters:

- Each node runs its own operating system.
- A cluster interconnect enables cluster communications.
- A public network connects each node to a LAN for client access.

- Shared storage is accessible by each node that needs to run the application.

SF Oracle RAC adds the following technologies, engineered specifically to improve performance, availability, and manageability of Oracle RAC environments, to a failover cluster environment:

- Cluster File System (CFS) and Cluster Volume Manager (CVM) technologies to manage multi-instance database access to shared storage.
- An Oracle Disk Manager (ODM) library to maximize Oracle disk I/O performance.
- Interfaces to Oracle Clusterware and Oracle RAC for managing cluster membership and communication.

SF Oracle RAC provides an environment that can tolerate failures with minimal downtime and interruption to users. If a node fails as clients access the same database on multiple nodes, clients that are attached to the failed node can reconnect to a surviving node and resume access. Recovery after failure in the SF Oracle RAC environment is far quicker than recovery for a failover database. Fast recovery is possible because another Oracle instance is already up and running. The recovery process involves applying outstanding redo log entries from the failed node.


About SF Oracle RAC components

The required components for SF Oracle RAC are as follows:

Cluster Volume Manager (CVM)	It enables simultaneous access to the shared volumes that are based on technology from Veritas Volume Manager (VxVM).
Cluster File System (CFS)	It enables simultaneous access to the shared file systems that are based on technology from Veritas File System (VxFS).
Database Accelerator	It provides the interface with the Oracle Disk Manager (ODM) API.
Veritas Cluster Server (VCS)	It uses the technology from Veritas Cluster Server to manage Oracle RAC databases and infrastructure components.
RAC Extensions	It manages the cluster membership and communications between cluster nodes.

Figure 1-1 displays an illustration of the SF Oracle RAC components.

Figure 1-1 SF Oracle RAC components



About SF Oracle RAC optional components and features

SF Oracle RAC supports the following optional components and features:

- [Veritas Installation Assessment Service](#)
- [Symantec Product Authentication Service](#)
- [Veritas Cluster Server Management Console](#)
- [Notification for VCS events](#)
- [Global clusters](#)
- [Veritas Volume Replicator](#)
- [Veritas Storage Foundation for Databases \(SFDB\) repository](#)

By default, the installer installs all filesets that are required for optional components and features at the time of installation. These components and features are optional for configuration with SF Oracle RAC. The options you choose to configure with SF Oracle RAC influence your preparation and configuration procedures.

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (VIAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation

or upgrade. The VIAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

<https://vias.symantec.com/>

Symantec Product Authentication Service

Symantec Product Authentication Service is a feature that validates identities based on existing network operating system domains (such as NIS and NT) or private domains. The authentication service protects communication channels among Symantec application clients and services through message integrity and confidentiality services.

Veritas Cluster Server Management Console

Veritas Cluster Server Management Console is a high availability management solution that enables monitoring and administering SF Oracle RAC clusters from a single web console.

Notification for VCS events

You have the option to configure SMTP email notification and SNMP trap notification of VCS events by the VCS Notifier component.

Global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. You are required to have a separate license to configure global clusters. You may add this license during the installation or at any time after the installation completes.

Veritas Volume Replicator

Veritas Volume Replicator (VVR) replicates data to remote locations over any standard IP network to provide continuous data availability. It is a fully integrated component of Veritas Volume Manager (VxVM). VVR is available as an optional, separately-licensed feature of SF Oracle RAC.

VVR replicates the application writes on the volumes at the source location to one or more remote locations across any distance. It provides a consistent copy of application data at the remote locations. If a disaster occurs at the source location, you can use the copy of the application data at the remote location and restart the application at the remote location. The host at the source location on which the application is running is known as the Primary host. The host at the

target location is known as the Secondary host. You can have up to 32 Secondary hosts in a VVR environment. VVR provides several methods to initialize the application data between the primary location and the remote location. Some of the methods include using the network, using tape backup, and moving disks physically.

Veritas Storage Foundation for Databases (SFDB) repository

The Storage Foundation for Databases (SFDB) repository or repository database stores metadata information that is required by SF Oracle RAC. This information includes data about user databases, snapshot databases, storage configuration, scheduled tasks, and storage statistics.

The SFDB repository database consists of a database file, `dbed_db.db`, and transaction log files, `yymmddxx.log`. The repository is stored in a relational database and is managed by a lightweight embedded relational DBMS, called VxDBMS. VxDBMS supports remote client access from any host in the network that has proper authorization and configuration.

It provides the following database management capabilities:

- [Veritas Storage Checkpoint and Storage Rollback](#)
- [Veritas Database FlashSnap](#)
- [Veritas Dynamic Storage Tiering](#)
- [Veritas Storage Mapping](#)

For detailed information about these features, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Veritas Storage Checkpoint and Storage Rollback

The Veritas Storage Checkpoint feature is used for the efficient backup and recovery of Oracle databases. It is available with SF Oracle RAC as part of the Veritas File System package.

A Storage Checkpoint provides a consistent image of the database from the point in time the Storage Checkpoint was created. A Storage Checkpoint can be mounted as read-only or read-write, and allows access to the files as if it were a regular file system. The Storage Checkpoint image is managed and available through the Veritas Storage Foundation command line interface (CLI).

Storage Rollback rolls back the changed blocks in a Storage Checkpoint into the primary file system for faster database restoration.

Veritas Database FlashSnap

Database FlashSnap offers you a flexible way to efficiently manage multiple point-in-time snapshots of your data, and reduce resource contention on your business-critical servers.

A database snapshot can be used for the following off-host processing applications:

- Data backup
- Data warehousing
- Decision-support queries

Database FlashSnap commands are executed from the command-line interface.

Veritas Dynamic Storage Tiering

Database Dynamic Storage Tiering (DST) matches data storage with data usage requirements. After matching, the data can be relocated based upon data usage and other requirements that are determined by the database administrator (DBA).

To use Database Dynamic Storage Tiering, your storage must be managed using the following features:

- VxFS multi-volume file system
- VxVM volume set
- Volume tags
- Dynamic Storage Tiering policies

Veritas Storage Mapping

Veritas Storage Mapping lets you map data files, tablespaces, and tables to physical devices. The mapping information provides a detailed understanding of the storage hierarchy in which files reside and help you evaluate and troubleshoot I/O performance.

Installation planning and requirements

This chapter includes the following topics:

- [Important preinstallation information](#)
- [Installation requirements](#)
- [Recommended Oracle configuration](#)
- [Recommended database storage configuration](#)
- [SF Oracle RAC cluster setup](#)
- [Overview of SF Oracle RAC installation and configuration tasks](#)
- [Required information for installing and configuring SF Oracle RAC](#)

Important preinstallation information

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

- Review the current compatibility list to confirm compatibility of your hardware:
<http://entsupport.symantec.com/docs/283161>
- TechNote for late-breaking and new information on updates, patches, and software issues regarding this release:
<http://entsupport.symantec.com/docs/282024>
- General information regarding the release, installation notes, known issues, and fixed issues:
See *Veritas Storage Foundation for Oracle RAC Release Notes*.

Installation requirements

Make sure that each node on which you want to install or upgrade SF Oracle RAC meets the following installation requirements:

- [Hardware requirements](#)
- [Supported software](#)
- [Supported operating systems](#)

Before you install SF Oracle RAC, you must read the Late Breaking News TechNote: <http://entsupport.symantec.com/docs/282024>

Hardware requirements

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

- Review the current compatibility list to confirm compatibility of your hardware: <http://entsupport.symantec.com/docs/283161>
- Review recommendations in Oracle's documentation for your version of Oracle.
- Review the hardware requirements for an SF Oracle RAC installation using SCSI-3 disks.

[Table 2-1](#) lists the hardware requirements.

Table 2-1 Hardware requirements	
Item	Description
SF Oracle RAC systems	Two to eight systems with two or more CPUs at 2GHz or higher.
DVD drive	One drive that is accessible to all nodes in the cluster.
Disks	Typical SF Oracle RAC configurations require that shared disks support applications that migrate between systems in the cluster. The SF Oracle RAC I/O fencing feature requires that all disks used as data disks or as coordinator disks must support SCSI-3 Persistent Reservations (PR). Note: The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space. See “Performing basic system checks” on page 92.
Disk space	See “Disk space (local)” on page 35.

Table 2-1 Hardware requirements (*continued*)

Item	Description
RAM	Each SF Oracle RAC system requires at least 2 GB. Symantec recommends additional amount of at least twice the Oracle SGA size.
Swap space	The minimum requirement is 4 GB of swap space. Symantec recommends two times the size of RAM for swap space.
Network links	Two or more private links and one public link. Links must be 100BaseT or Gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication. These links must be of the same type; you cannot mix 100BaseT and Gigabit. Symantec recommends Gigabit Ethernet using enterprise-class switches for the private links. You can also configure aggregated interfaces.
Fibre channel or SCSI host bus adapters	SF Oracle RAC requires at least one built-in SCSI adapter per system to access the operating system disks, and at least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

Disk space (local)

Confirm that your system has enough free disk space to install SF Oracle RAC. You can evaluate your systems for available disk space by running the 'installmp' or the 'installsfrac' script with the '-precheck' option. The output of the script compares the required disk space against the available disk space on your system.

Use the 'installsfrac' script with the '-precheck' option if you are installing SF Oracle RAC 5.0 MP3 or upgrading from version 4.x.

```
# ./installsfrac -precheck node_name
```

Use the 'installmp' script with the '-precheck' option if you are upgrading from version 5.0x to SF Oracle RAC 5.0 MP3.

```
# ./installmp -precheck node_name
```

For example, to evaluate the system 'galaxy' for sufficient disk space before upgrading from version 5.0 to SF Oracle RAC 5.0 MP3, run the following command:

```
# ./installmp -precheck galaxy
```

Veritas Maintenance Pack 5.0 Pre-Installation Check Program
Copyright (c) 2008 Symantec Corporation. All rights reserved.
Symantec, the Symantec Logo 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 Licensed Software and Documentation are deemed to be "commercial
computer software" and "commercial computer software documentation" as
:

The following MP3 installation issues have been found on galaxy:
389518 KB is required in the /opt volume and
only 122216 KB is available on galaxy

844874 KB is required in the /usr volume and
only 436296 KB is available on galaxy

485767 KB is required in the /var volume and
only 473108 KB is available on galaxy

installmp log files are saved at /opt/VRTS/install/logs/installmp-y8ITho

You may also use the following information as a guideline to ensure that you have
sufficient local disk space to install the SF Oracle RAC filesets.

/opt	1.8 G
/usr	390 MB
/tmp	512 MB
/var	636 MB
/var/tmp	700 MB
/	120 MB
total	4.2 G

Additional space is required for Oracle. For details on the space required for Oracle,
consult the Oracle documentation.

To check for current RAM, swap space, and /tmp

- 1 To check the current RAM installed on a system, enter the following command:

```
# /usr/sbin/lsattr -E -l sys0 -a realmem
```

If the current RAM is insufficient, you must install more memory before proceeding with the installation.

- 2 To check the amount of swap space configured for a system, enter:

```
# /usr/sbin/lsp -a
```

Symantec recommends two times the main memory for swap space. Additionally, refer to the Oracle documentation for recommendations on swap space.

Refer to the operating system documentation for information about configuring swap space.

See the *Oracle Database Administrator's Reference for UNIX* for more information.

- 3 To check for the current space available in the /tmp directory, enter:

```
# df -k /tmp
```

If the space remaining in /tmp is insufficient, complete one of the following operations:

- Delete unnecessary files in /tmp, and recheck the available space.
- Extend the file system that contains the /tmp directory.
- Set the TEMP and TMPDIR environment variables when setting the Oracle user's environment.

Supported replication technologies

SF Oracle RAC supports the following hardware-based replication technologies in addition to Veritas Volume Replicator (VVR):

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

For information on replication requirements for SF Oracle RAC clusters:
See “[Requirements for SF Oracle RAC global clusters](#)” on page 432.

Supported software

Caution: SF Oracle RAC software must be at the same version across all nodes in an SF Oracle RAC cluster, in this case 5.0MP3++. All components and point products of SF Oracle RAC must be from the same version of the release, in this case 5.0MP3++.

Symantec strongly recommends upgrading all SF Oracle RAC component products to the same version at the same time. In a CVM/CFS environment, dependencies between SF Oracle RAC component products will not be met if you do not upgrade all components to the same version.

For information on upgrade options:
See “[Supported upgrade options for SF Oracle RAC 5.0 MP3](#)” on page 145.
Software versions that SF Oracle RAC supports include:

Oracle RAC	■ Oracle RAC 9i Release 2 ■ Oracle RAC 10g Release 1 ■ Oracle RAC 10g Release 2 ■ Oracle RAC 11g Release 1
AIX operating system	See “ Supported operating systems ” on page 39.
VCS, VxVM, VxFS, VVR	Use only versions of VCS, VxVM, VxFS, and VVR provided on the software disc. Remove other versions before you install the software from the SF Oracle RAC product disc.

To verify the latest information on support for Oracle database versions, see the Technical Support TechNote:
<http://entsupport.symantec.com/docs/280186>

Oracle software patches

The following Oracle patches are mandatory for proper functioning of SF Oracle RAC.

Oracle RAC 10g Release 2 Patchset 1 (10.2.0.1)	Patch 4637591 Patch 4435949 Patch 5082958
Oracle RAC 10g Release 2 Patchset 2 (10.2.0.2)	Patch 4637591
Oracle RAC 10g Release 2 Patchset 3 (10.2.0.3)	Patch 4430244 Patch 5752399 Patch 5769259 Patch 6613550 (AIX 6.1 only)
Oracle RAC 10g Release 2 Patchset 4 (10.2.0.4)	None
Oracle RAC 11g Release 1 Patchset 6 (11.1.0.6)	Patch 6849184 Patch 6442900 Patch 6792086

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

Supported operating systems

SF Oracle RAC can be installed on the following operating systems:

- AIX 5.3 TL7 SP2
- AIX 6.1 TL0 SP4

For the most current information on downloads, service pack availability, fixes, and other important issues related to 5.0 MP releases on AIX:

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

At the time this document is published, AIX 6.1 supports only Oracle RAC 10g R2.

APAR and fileset information

You will need the following AIX filesets for Oracle installation:

- bos.adt.libm
- perfagent.tools
- bos.perf
- Java14
- rsct.core

- rsct.compat.basic
- rsct.compat.clients
- rsct.basic
- rsct.clients

The Veritas High Availability Agent for Oracle on AIX requires the following runtime fileset level for the xlC.rte filesets:

xlC.aix61.rte	9.0.0.1	COMMITTED	XL C/C++ Runtime for AIX 6.1
xlC.cpp	9.0.0.0	COMMITTED	C for AIX Preprocessor
xlC.msg.en_US.cpp	9.0.0.0	COMMITTED	C for AIX Preprocessor
xlC.msg.en_US.rte	9.0.0.1	COMMITTED	XL C/C++ Runtime
xlC.rte	9.0.0.1	COMMITTED	XL C/C++ Runtime

Recommended Oracle configuration

The following recommendations provide best practices for configuring Oracle in an SF Oracle RAC environment.

Placement of Oracle binaries	Install Oracle binaries on local disk to enable phased upgrade of the cluster.
Placement of Oracle data files and logs	Create separate clustered file systems for Oracle data files. Place redo and archived logs on CFS rather than on local file systems.
Placement of OCR and Voting disk	Place the OCR and VOTE devices on CVM volumes or CFS file systems to provide high availability to VOTE disks with dynamic multipathing. Mirror the CVM volumes that contain OCR and VOTE devices for redundancy.
Configuration of Oracle private IPs	Configure Oracle private IP addresses under VCS PrivNIC or MultiPrivNIC to protect against link failures. This requires you to configure CRS private IP addresses and UDP cache fusion IP addresses on LLT interconnects.

ASM considerations

- Use VxVM mirrored volumes with dynamic multi-pathing with external redundancy to ensure high availability.
- Do not use VxVM volumes, which are used by ASM, for any other purpose such as creation of file systems.
- Do not enable ODM when databases are installed on ASM.

Recommended database storage configuration

The following recommendations provide best practices for improving performance, protection, and high availability of data in your SF Oracle RAC environment.

Storage

- Use multiple storage arrays to ensure protection against array failures.
- 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 PGR compliant storage.
- Use dynamic multi-pathing with multiple access paths to disks.

Volume design

- Mirror the volumes across two or more storage arrays, especially with host-based mirrors.
- Separate Oracle recovery structures from 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.
- Implement zoning 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.

File design

- Use ODM with Cluster File System 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.
- Always place redo and archived logs on CFS file systems rather than local file systems.

SF Oracle RAC cluster setup

The following SF Oracle RAC setup examples illustrate the diversity of supported cluster configurations. Each setup example implies a set of requirements, preparations, installation, and configuration procedures that may differ from what is required for the other examples. The basic setup example illustrates the minimum level of requirements for a successful SF Oracle RAC installation. If you use a basic configuration, you can skip the procedures for the optional features that you do not require.

The setup that is required for many optional features requires additional procedures that must be completed before SF Oracle RAC installation and configuration. For instance, to set up a cluster in secure mode a root broker must be configured before the SF Oracle RAC installation.


For global clusters, two SF Oracle RAC clusters must be configured and linked. The example for global clusters covers the specific setup for using replication in the SF Oracle RAC environment.

- [Basic SF Oracle RAC cluster setup](#)
- [SF Oracle RAC cluster optional features setup](#)
- [SF Oracle RAC global cluster setup](#)

Basic SF Oracle RAC cluster setup

[Figure 2-1](#) depicts a high-level view of an SF Oracle RAC configuration for a four-node cluster with the minimum required features.

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



A basic installation comprising SF Oracle RAC with Oracle and a database typically has the following characteristics:

- Nodes that are connected by at least two VCS private network links using 100 Base T or Gigabit Ethernet controllers on each system.
For maximum performance, Symantec recommends the use of switches over hubs. In either case, use a minimum of two switches or hubs to provide the necessary redundancy.
If multiple links (three or more in number) are present on a single switch, a separate VLAN must be constructed for each link. Symantec does not support the use of multiple links on a single hub.

- Nodes that are connected to shared storage devices through Fibre Channel switch.
Symantec does not support the use of shared SCSI with the SF Oracle RAC product. For a complete list of supported Fibre Channel storage devices, see the current hardware compatibility list on the Symantec Support Web site.
<http://entsupport.symantec.com/docs/283161>
- Nodes must be connected with private network links using similar network devices and matching port numbers.
For example, if you use en1 on one end of a link, the other end must also use en1.
- Each system has a local VxVM, which can be an encapsulated root drive or an internal disk.
Symantec does not support using the root disk group for storage that is shared between systems.
- Oracle database is configured on the shared storage that is available to each node. The shared storage can be a cluster file system or raw volumes. All shared storage, including coordinator disks, must support SCSI-3 PR.
- VCS is configured to enable agents to direct and manage the resources that are required by Oracle RAC; these resources run in parallel on each node.

SF Oracle RAC cluster optional features setup


SF Oracle RAC provides additional security and management capabilities through the following options:

- Symantec Product Authentication Service (AT)
- Veritas Cluster Server Management Console (VCS MC)
- Notification for VCS

The options you choose influence the preparation, installation, and configuration procedures you follow for SF Oracle RAC. You can choose the appropriate options during the installation of SF Oracle RAC.

[Figure 2-2](#) illustrates a sample deployment with the Symantec Product Authentication Service (AT) and Veritas Cluster Server Management Console (VCS MC) components.

Figure 2-2 Sample deployment with AT and VCS MC components



SF Oracle RAC global cluster setup

SF Oracle RAC supports global clusters for Oracle RAC. VCS provides a high availability disaster recovery (HA/DR) option for wide-area failover and disaster recovery, which works with Veritas Volume Replicator and other supported replication options.

Note: You must have an SF Oracle RAC HA/DR license to configure global clusters.

- You may add the SF Oracle RAC HA/DR license during the initial installation or at any time after the installation to configure a global cluster.
- You may configure a basic cluster initially and add the HA/DR license to configure a global cluster at a later time.

For replication requirements for SF Oracle RAC clusters:

See [“Requirements for SF Oracle RAC global clusters”](#) on page 432.


Overview of SF Oracle RAC installation and configuration tasks

Installing and configuring SF Oracle RAC involves the following tasks:

- [Tasks to prepare for optional SF Oracle RAC features](#)
- [Tasks to prepare for SF Oracle RAC installation](#)
- [Tasks for SF Oracle RAC installation and configuration](#)
- [Tasks for installing Oracle RAC and creating an Oracle database](#)
- [Tasks for setting up VCS to manage Oracle RAC resources](#)
- [Tasks for setting up database management features for SF Oracle RAC](#)
- [Tasks for setting up SF Oracle RAC global clusters \(optional\)](#)

[Figure 2-3](#) displays a high-level flow of the installation and configuration processes for SF Oracle RAC.

Figure 2-3 SF Oracle RAC installation and configuration



Tasks to prepare for optional SF Oracle RAC features

Some of the optional features for SF Oracle RAC require you to perform preparatory tasks before you configure these features.

See [“Setting up Symantec Product Authentication Service”](#) on page 48.

See [“Setting up Veritas Cluster Server Management Console”](#) on page 48.

See [“Setting up notification for VCS events”](#) on page 49.

Setting up Symantec Product Authentication Service

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 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.

- **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 58.

Setting up Veritas Cluster Server Management Console

Veritas Cluster Server Management Console is a high availability management solution that enables monitoring and administering multiple SF Oracle RAC clusters from a single web console. You can configure Veritas Cluster Server Management Console to manage a single cluster, multiple clusters, or both. If you want to use it to manage multiple clusters, you must install and set up a management server.

See *Veritas Cluster Server Management Console Implementation Guide* for more information.

Setting up notification for VCS events

You have the option to configure SMTP email notification and SNMP trap notification of VCS events by the VCS Notifier component. You must have configured a notification server before running the SF Oracle RAC installation.

If you choose the SMTP or SNMP notification option, the installer enables you to choose whether or not to use the same NIC as configured for the ClusterService group, which is the default. If you choose not to use the same networking information, you must specify appropriate values when the installer prompts you.

See *Veritas Cluster Server User's Guide* for more information on SMTP and SNMP notification.

Tasks to prepare for SF Oracle RAC installation

A basic SF Oracle RAC installation involves the following preparatory tasks:

- [Setting umask before installation](#)
- [Synchronizing time settings on cluster nodes](#)
- [Setting up inter-system communication](#)
- [Setting the MANPATH variable](#)
- [Setting the PATH variable](#)
- [Obtaining SF Oracle RAC license keys](#)
- [Removing pre-existing license keys](#)
- [Verifying the systems before installation](#)

Tasks for SF Oracle RAC installation and configuration

[Table 2-2](#) lists the tasks that are required to complete SF Oracle RAC installation and configuration.

Table 2-2 SF Oracle RAC installation and configuration tasks

Task	Procedural reference
Invoking SF Oracle RAC installation program	■ Mounting the product disc ■ Starting software installation ■ Specifying systems for installation

Table 2-2 SF Oracle RAC installation and configuration tasks *(continued)*

Task	Procedural reference
Licensing SF Oracle RAC and adding SF Oracle RAC filesets to each cluster node	■ Licensing SF Oracle RAC ■ Installing SF Oracle RAC filesets
Invoking SF Oracle RAC configuration program	■ Mounting the product disc ■ Starting software configuration ■ Specifying systems for configuration ■ Choosing the configuration task
Performing basic system checks	■ Preparing private interconnects for LLT
Configuring SF Oracle RAC components	■ Veritas Cluster Server component information ■ Setting permissions for database administration ■ Configuring the Cluster Volume Manager ■ Configuring VVR on each node ■ Creating SF Oracle RAC configuration files ■ Configuring standalone hosts for Storage Foundation ■ Starting SF Oracle RAC processes ■ Enabling LDAP authentication for clusters that run in secure mode
Configuring I/O fencing	■ Preparing to configure I/O fencing ■ Setting up I/O fencing
Verifying the cluster	■ Verifying SF Oracle RAC installation using VCS configuration file

The example procedure installs and configures SF Oracle RAC on two nodes: galaxy and nebula. For the sample installation, the cluster’s name is rac_cluster101 and the cluster’s ID is 101. The example installation chooses all optional features.

Tasks for installing Oracle RAC and creating an Oracle database

After installing and configuring components of Storage Foundation for Oracle RAC, proceed to install Oracle RAC.

For supported Oracle versions:

See [“Supported software”](#) on page 38.

The following sections describe how to prepare for and install Oracle RAC:

- Prepare to install Oracle RAC
 See [“About preparing to install Oracle RAC 10g”](#) on page 195.
 See [“About preparing to install Oracle RAC 11g”](#) on page 275.
- Install Oracle RAC.
 See [“About installing Oracle RAC 10g”](#) on page 231.
 See [“About installing Oracle RAC 11g”](#) on page 315.
- Create a raw database on raw volumes within a VxVM disk group or on a Veritas cluster file system.
 You may create the database using the Oracle DBCA utility. For creating a database using Oracle DBCA, see the Oracle documentation.
 You may also create a test database for testing and troubleshooting purposes.
 For creating a test database:
 See [“About creating a test database”](#) on page 617.

Tasks for setting up VCS to manage Oracle RAC resources

SF Oracle RAC provides the capability to completely automate the Oracle RAC environment. This capability ranges from enabling automatic control of the entire database environment to having VCS mount cluster file systems or enable CVM and CFS daemons. The user or DBA is free to choose the level of control and automation.

VCS uses the `main.cf` configuration file to manage resources in the cluster. The SF Oracle RAC installation process creates a basic VCS configuration file. After installing Oracle and creating the database, you can modify the `main.cf` file on one of the cluster nodes to reflect the new resources and their configuration.

Refer to the following sections when configuring VCS.

- See [“About VCS service group for Oracle RAC 10g dependencies”](#) on page 253.
- See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.
- See [“About sample `main.cf` files”](#) on page 507.

For Oracle RAC 9i, see the *Veritas Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide*.

For Oracle RAC 10g and Oracle RAC 11g, see the respective chapters in this guide.

Tasks for setting up database management features for SF Oracle RAC

SF Oracle RAC supports the following database management features:

- Storage Checkpoint and Storage Rollback
- Database FlashSnap

- Dynamic Storage Tiering
- Storage Mapping

To use the repository database, you need to create and set up the repository database after installing SF Oracle RAC.

The basic tasks for setting up the SFDB repository are as follows:

- Run the `sfua_db_config` script as the root user to create and configure the repository database.
- Set appropriate permissions to allow database administration using SF Oracle RAC.

Storage Mapping does not require any separate setup tasks. The required Veritas Mapping Service (VxMS) packages are automatically installed during the SF Oracle RAC installation. By default, Storage Mapping is licensed for use with SF Oracle RAC.

You may verify that Storage Mapping is enabled for use with SF Oracle RAC by performing the following basic tasks:

- Verify that you have a license key for the storage mapping option.
- Verify that the `VRTSvxmsa` package is installed.

For detailed information on these database management features, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Tasks for setting up SF Oracle RAC global clusters (optional)

You can choose to set up a global cluster environment for disaster recovery after you have SF Oracle RAC and Oracle RAC running successfully on a local cluster, or you can configure both the local and remote clusters at the same time. The general process for setting up a global cluster involves the following tasks on both sites:

- Creating a cluster at each site with hardware set up
- Installing SF Oracle RAC
- Configuring SF Oracle RAC
- Installing Oracle RAC
- Configuring VCS service groups

You do not need to create a database for the secondary site. The database is replicated from the primary site. You can use Veritas Volume Replicator (VVR) or you can use a supported hardware-based replication product.

VVR is an optional, separately-licensed replication product that is supported by the High Availability Disaster Recovery (HA/DR) version of SF Oracle RAC. Veritas Volume Replicator is a fully integrated component of Veritas Volume Manager that replicates data to remote locations over any standard IP network to provide continuous data availability. You may install VVR during your initial SF Oracle RAC installation process or you may add it when configuring a global cluster at a later time.

If you are using SF Oracle RAC HA/DR with hardware-based replication, you need the VCS agent that supports your replication product.

SF Oracle RAC supports the following hardware-based replication options.

- For EMC SRDF:
 See *Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide*
- For Hitachi TrueCopy:
 See *Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide*
- For IBM Metro Mirror:
 See *Veritas Cluster Server Agent for IBM Metro Mirror Installation and Configuration Guide*
- For IBM SVC:
 See *Veritas Cluster Server Agent for IBM SVCCopyServices Installation and Configuration guide*
- For EMC MirrorView:
 See *Veritas Cluster Server Agent for EMC MirrowView Installation and Configuration Guide*
- For Oracle Data Guard:
 See *Veritas Cluster Server Agent for Oracle Data Guard Installation and Configuration Guide*

See the Technical Support TechNote for the latest updates or software issues for replication agents:

<http://entsupport.symantec.com/docs/282004.htm>

For detailed information about configuring SF Oracle RAC global clusters:

See “[Replication in the SF Oracle RAC environment](#)” on page 431.

Required information for installing and configuring SF Oracle RAC

The SF Oracle RAC installation and configuration program prompts you for information about some SF Oracle RAC components. This installation program provides default values for some information which you can choose to use.

To shorten the time required for your installation and configuration processes and improve the accuracy of your entries:

- Review the information requirements before you begin SF Oracle RAC installation and configuration.
See [“SF Oracle RAC installation and configuration information”](#) on page 559.
- Use the sample values in the worksheets as examples of the information that is required for installing and configuring SF Oracle RAC.
See [“Installation and configuration worksheets”](#) on page 567.
- Assemble the required information for your configuration.
- Have the required information for your configuration at hand during the installation process.

SF Oracle RAC installation and configuration

- [Preparing to install and configure SF Oracle RAC](#)
- [Installing SF Oracle RAC](#)
- [Configuring SF Oracle RAC](#)
- [Configuring SF Oracle RAC clusters for data integrity](#)
- [Verifying cluster installation](#)
- [Upgrading to SF Oracle RAC 5.0 MP3](#)

Preparing to install and configure SF Oracle RAC

This chapter includes the following topics:

- [About preparing to install and configure SF Oracle RAC](#)
- [Preparing to configure the clusters in secure mode](#)
- [Preparing to install SF Oracle RAC](#)

About preparing to install and configure SF Oracle RAC

Before you perform the pre-installation tasks, make sure that you reviewed the installation requirements, set up the basic hardware, and planned your SF Oracle RAC setup.

See [“Installation requirements”](#) on page 34.


See [“Overview of SF Oracle RAC installation and configuration tasks”](#) on page 46.

The SF Oracle RAC installation and configuration procedures consist of both optional and mandatory steps. The optional steps are performed for optional SF Oracle RAC features.


[Figure 3-1](#) displays both the optional and mandatory steps.

Figure 3-1 SF Oracle RAC installation and configuration procedures

Optional steps to
perform before SF
Oracle RAC installation:



Mandatory steps to
perform before SF
Oracle RAC installation:



Note: If you do not want to configure optional features for your SF Oracle RAC cluster setup, then proceed directly to [“Preparing to install SF Oracle RAC”](#) on page 67.

Preparing to configure the clusters in secure mode

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

Refer to the *Veritas Cluster Server User's Guide* for instructions to configure AT in a cluster that does not run in secure mode.

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

- A system in your enterprise is configured as root broker (RB).
If a root broker system does not exist, install and configure root broker on a system.
See [“Installing the root broker for the security infrastructure”](#) on page 62.
- 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 63.
- The system clocks of the root broker and authentication brokers must be in sync.

The `installsrac` program provides the following configuration modes:

Automatic mode	The root broker system must allow rsh or ssh passwordless login to use this mode.
Semi-automatic mode	This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 68.
Manual mode	This mode requires <code>root_hash</code> file and the root broker information from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 68.

[Figure 3-2](#) depicts the flow of configuring SF Oracle RAC cluster in secure mode.

Figure 3-2 Workflow to configure SF Oracle RAC cluster in secure mode




Table 3-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform.

Table 3-1 Preparatory tasks to configure a cluster in secure mode

Tasks	Who performs this task
Decide one of the following configuration modes to set up a cluster in secure mode: ■ Automatic mode ■ Semi-automatic mode ■ Manual mode	VCS administrator
Install the root broker on a stable system in the enterprise. See “Installing the root broker for the security infrastructure” on page 62.	AT administrator
On the root broker system, create authentication broker accounts for each node in the cluster. See “Creating authentication broker accounts on root broker system” on page 63. AT administrator requires the following information from the VCS administrator: ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker	AT administrator
To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator. See “Creating encrypted files for the security infrastructure” on page 64. AT administrator requires the following additional information from the VCS administrator: ■ Administrator password for each authentication broker Typically, the password is the same for all nodes.	AT administrator
To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.	AT administrator
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. See “Preparing the installation system for the security infrastructure” on page 66.	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. 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 install the root broker on an AIX, HP-UX, Linux, or Solaris system.

See Symantec Product Authentication Service documentation for more information.

See [“Setting up Symantec Product Authentication Service”](#) on page 48.

To install the root broker

- 1 Change to the directory where you can start the installsfrac program:
- 2 Start the Root Broker installation program:

```
# ./installsfrac -configure -security
```

- 3 Select to install the Root Broker from the three choices that the installer presents:

```
3 Install Symantec Security Services Root Broker
```

- 4 Enter the name of the system where you want to install the Root Broker.

```
Enter the system name on which to install VxSS: venus
```

- 5 Review the output as the installer does the following:
 - Checks to make sure that SF Oracle RAC supports the operating system
 - Verifies that you install from the global zone (only on Solaris)
 - Checks if the system is already configured for security
- 6 Review the output as the installsfrac program checks for the installed filesets on the system.

The installsfrac program lists the filesets that the program is about to install on the system. Press Enter to continue.

- 7 Review the output as the installer installs the root broker on the system.
- 8 Enter **y** when the installer prompts you to configure the Symantec Product Authentication Service.

- 9 Enter a password for the root broker. Make sure the password contains a minimum of five characters.
- 10 Enter a password for the authentication broker. Make sure the password contains a minimum of five characters.
- 11 Press the Enter key to start the Authentication Server processes.

```
Do you want to start Symantec Product Authentication Service
processes now? [y,n,q] y
```

- 12 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 <code>addprpl</code> command. See “Creating authentication broker accounts on root broker system” on page 63.
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 <code>addprpl</code> command. See “Creating authentication broker accounts on root broker system” on page 63.
broker_admin_password	The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.

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=vx:root@venus.symantecexample.com
root_broker=venus.symantecexample.com:2821
broker_admin_password=ab_admin_password
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 files to the installer node.

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 Do the following:

- Copy the encrypted files (BLOB files) to the system from where you plan to install VCS.
Note the path of these files that you copied to the installation system.
- During SF Oracle RAC configuration, choose the configuration option 2 when the installsfrac program prompts.

Manual mode

Do the following:

- Copy the root_hash file that you fetched to the system from where you plan to install VCS.
Note the path of the root hash file that you copied to the installation system.
- Gather the root broker information such as name, fully qualified domain name, domain, and port from the AT administrator.
- Note the principal name and password information for each authentication broker that you provided to the AT administrator to create the authentication broker accounts.
- During SF Oracle RAC configuration, choose the configuration option 3 when the installsfrac program prompts.

Preparing to install SF Oracle RAC

Perform the following tasks before proceeding to install SF Oracle RAC:

- [Setting umask before installation](#)
- [Synchronizing time settings on cluster nodes](#)
- [Setting up inter-system communication](#)
- [Setting up shared storage](#)
- [Setting the MANPATH variable](#)
- [Setting the PATH variable](#)
- [Mounting the product disc](#)
- [Obtaining SF Oracle RAC license keys](#)
- [Removing pre-existing license keys](#)
- [Obtaining and installing license keys for a global cluster](#)
- [Verifying the systems before installation](#)

Setting umask before installation

Use the following command to set the default file creation mode on each system where you plan to install Oracle RAC.

```
# umask 0022
```

Synchronizing time settings on cluster nodes

If you plan to configure the SF Oracle RAC in secure mode, Symantec recommends all cluster nodes have the same time. If you do not run the Network Time Protocol (NTP) daemon, make sure to synchronize the time settings on each node.

Setting up inter-system communication

When you install SF Oracle RAC using the `installsrac` program, to install and configure the entire cluster at one time, make sure that communication between systems exists. By default, the installer uses `ssh` for inter-system communication. You must grant permissions for the system where you run `installsrac` program to issue `ssh` or `rsh` commands as root on all systems in the cluster.

- If you configured `ssh` for communication, you must configure it in a way such that it operates without requests for passwords or passphrases.
- If you configured `rsh` for communication, on each node, placing a "+" character in the first line of the `/.rhosts` file gives remote access to the system running the install program.

You can limit the remote access to specific nodes. Remove the remote `rsh` access permissions after the installation and disk verification process. Refer to the manual page for the `/.rhosts` file for more information.

Configuring and enabling SSH

The SSH program enables you to log into and execute commands on a remote system. SSH enables encrypted communications and an authentication process between two untrusted hosts over an insecure network. SSH is the preferred method of remote communication because it provides a greater level of security than the RSH suite of protocols.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the `authorized_keys` file on the target systems.

If you are installing Oracle, you must configure a DSA key and an RSA key for the Oracle user in addition to the DSA key required for the root user to install SF Oracle RAC.

Read the SSH documentation and online manual pages before enabling SSH. Contact your operating system support provider for issues regarding SSH configuration.

Visit the OpenSSH website that is located at: <http://openssh.org> to access online manuals and other resources.

To create the DSA key pair

- 1 On the source system (*galaxy*), log in as root, and navigate to the root directory.

```
galaxy # cd /
```

- 2 To generate a DSA key pair on the source system, type the following command:

```
galaxy # 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`.
- 4 When the program asks you to enter the passphrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

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 (*nebula* in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

```
nebula # cd /  
nebula # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
nebula # chmod go-w /.ssh
```

To append the public key from the source system to the `authorized_keys` file on the target system, using secure file transfer

- 1 From the source system (*galaxy*), move the public key to a temporary file on the target system (*nebula*).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
galaxy # sftp nebula
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to nebula ...  
The authenticity of host 'nebula (10.182.00.00)'  
can't be established. DSA key fingerprint is  
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.  
Are you sure you want to continue connecting (yes/no)?
```

- 2 Enter "yes".

Output similar to the following is displayed:

```
Warning: Permanently added 'nebula,10.182.00.00'  
(DSA) to the list of known hosts.  
root@nebula password:
```

- 3 Enter the root password of *nebula*.
- 4 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
```

- 5 To quit the SFTP session, type the following command:

```
sftp> quit
```

- 6 To begin the `ssh` session on the target system (nebula in this example), type the following command on *galaxy*:

```
galaxy # ssh nebula
```

Enter the root password of nebula at the prompt:

```
password:
```

- 7 After you log in to nebula, enter the following command to append the `id_dsa.pub` file to the authorization key file:

```
nebula # cat /id_dsa.pub >> /.ssh/authorized_keys
```

- 8 After the `id_dsa.pub` public key file is copied to the target system (nebula), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, type the following command on nebula:

```
nebula # rm /id_dsa.pub
```

- 9 To log out of the `ssh` session, type the following command:

```
nebula # exit
```

- 10 When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
galaxy # cat /.ssh/id_dsa.pub >> /.ssh/authorized_keys
```

- 11 Run the following commands on the source installation system. If your `ssh` session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and makes the key globally available for the user root:

```
galaxy # exec /usr/bin/ssh-agent $SHELL
```

```
galaxy # ssh-add
```

```
Identity added: //./ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

- 1 On the source system (*galaxy*), type the following command:

```
galaxy # ssh -l root nebula uname -a
```

where nebula is the name of the target system.

- 2 The command should execute from the source system (*galaxy*) to the target system (*nebula*) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

Restarting ssh

After you complete this procedure, ssh can be restarted in any of the following scenarios:

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted
- After too much time has elapsed, to refresh ssh

To restart ssh

- 1 On the source installation system (*galaxy*), bring the private key into the shell environment.

```
galaxy # exec /usr/bin/ssh-agent $SHELL
```

- 2 Make the key globally available for the user `root`

```
galaxy # ssh-add
```

Enabling RSH

To enable `rsh`, create a `.rhosts` file on each target system. Then add a line to the file specifying the full domain name of the source system. For example, add the line:

```
.sysname.domainname.com root
```

Change permissions on the `.rhosts` file to 600 by typing the following command:

```
# chmod 600 /.rhosts
```


After you complete an installation procedure, delete the `.rhosts` file from each target system to ensure security:

```
# rm -f /.rhosts
```

Setting up shared storage

You need to set up shared storage so that it is visible to the SCSI layer from all the nodes in the cluster. The shared storage that you add for use with SF Oracle RAC software must support SCSI-3 persistent reservations, a functionality that enables the use of I/O fencing.

See [“About setting up I/O fencing”](#) on page 116.

See also the *Veritas Cluster Server User's Guide* and the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more information on I/O fencing.

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 your SF Oracle RAC installation and configuration.

To mount the product disc

- 1 Log in as 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 connected to the system.
- 3 Mount the disc.

Determine the device access name of the DVD drive. The format for the device access name is `cdx` where `x` is the device number. After inserting the disc, type the following commands:

```
# mkdir -p /cdrom
```

```
# mount -V cdrfs -o ro /dev/cdX /cdrom
```

- 4 Navigate to the location of the filesets.

```
# cd /cdrom/dvd1/sfrac
```

Setting the MANPATH variable

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

To set the MANPATH variable

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

```
Bourne Shell (sh or ksh)      # MANPATH=/usr/share/man:/opt/VRTS/man;\n                                export MANPATH
```

```
C Shell (csh)          # setenv MANPATH /usr/share/man:\n                        /opt/VRTS/man
```

Setting the PATH variable

You can find the installation and other commands located in various directories. Add these directories to your PATH environment variable in the .profile file (or other appropriate shell setup file for your system) on each system.

To set the PATH variable

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

```
Bourne Shell (sh or ksh)      # PATH=/usr/sbin:/sbin:/usr/bin:\n                                /usr/lib/vxvm/bin:/opt/VRTSvxfs/sbin:\n                                /opt/VRTSvcs/bin:/opt/VRTS/bin:\n                                /opt/VRTSvcs/rac/bin:/opt/VRTSob/bin:\n                                $PATH; export PATH
```

```
C Shell (csh)          # setenv PATH /usr/sbin:/sbin:/usr/bin:\n                        /usr/lib/vxvm/bin:/opt/VRTSvxfs/sbin:\n                        /opt/VRTSvcs/bin:/opt/VRTS/bin:\n                        /opt/VRTSvcs/rac/bin:/opt/VRTSob/bin:$PATH
```

Remember the following while setting the PATH variable:

- You must set the PATH variable for java/bin before executing the `installsf -configure` command:

```
# export PATH=$PATH:/usr/javaXX/bin
```

- Do not define \$ORACLE_HOME/lib in LIBPATH for root user. You should define \$ORACLE_HOME/lib in LIBPATH for the oracle user.

Obtaining SF Oracle RAC license keys

SF Oracle RAC includes a license key certificate. The certificate specifies the product keys and the number of product licenses purchased. A single key lets you install the product on the number and type of systems for which you purchased the license. A key may enable the operation of more products than are specified on the certificate; however, you are legally limited to the number of product licenses purchased.

Note: If you have pre-existing license keys, you should remove them before proceeding. For directions on how to remove a pre-existing license key, see [Removing pre-existing license keys](#).

[Table 3-2](#) lists the features enabled for SF Oracle RAC components.

Table 3-2 Features enabled for SF Oracle RAC components

Feature	Component
VXCFS	Veritas File System
File Change Log	Veritas File System
Cross-platform Data Sharing	Veritas File System
Extra-Big File Systems	Veritas File System
Multi-Volume Support	Veritas File System
Quality of Storage Service	Veritas File System
VXCKPT	Veritas File System
VxVM	Veritas Volume Manager
CVM_FULL	Veritas Volume Manager
PGR	Veritas Volume Manager
FMR_DGSJ	Veritas Volume Manager
Dynamic Lun Expansion	Veritas Volume Manager
Hardware assisted copy	Veritas Volume Manager

Table 3-2 Features enabled for SF Oracle RAC components *(continued)*

Feature	Component
Cross-platform Data Sharing	Veritas Volume Manager
VVS_CONFIG	Veritas Volume Manager
PGR_TRAINING	Veritas Volume Manager
Site Awareness	Veritas Volume Manager
VVR	Veritas Volume Replicator
GCO	Veritas Storage Foundation for Oracle RAC HA/DR Option
Found_Edi_map	Veritas Mapping Services
Mode=VCS_RAC	Veritas Cluster Server
Mode=VCS_0	Veritas Cluster Server

The product installation procedure includes instructions on how to activate the key. If you encounter problems while licensing this product, visit the Symantec licensing support website at:

<https://licensing.symantec.com>

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

vxlicinst	Installs a license key for a Symantec product
vxlicrep	Displays currently installed licenses
vxlictest	Retrieves features and their descriptions encoded in a license key

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

If you will use VVR for replication, you can enter the VVR license during the SF Oracle RAC install and it will be installed concurrently with SF Oracle RAC. If you do not currently have a VVR license, you can install it at a later time.

Removing pre-existing license keys

This procedure is used in cases where a pre-existing license is replaced or removed:

- Upgrading SF Oracle RAC
- Adding a node to a cluster
- Removing a node from a cluster
- Adding HA/DR capability to a cluster

If you are installing on a new system, you can skip this procedure.

See [“Obtaining SF Oracle RAC license keys”](#) on page 75.

To remove pre-existing license keys

- 1 View license key files currently installed on a node:

```
# cd /opt/VRTS/bin
# ./vxlicrep
```

The output lists the license keys and information about their respective products.

- 2 Navigate to the directory containing the license key files and list the files.

```
# cd /etc/vx/licenses/lic
# ls -a
```

- 3 Move the license key files to another location for future reference.

For example, in the directory containing the license key files, create a sub-directory and move the files to that new location:

```
# mkdir OLD
# mv * OLD
```

Obtaining and installing license keys for a global cluster

If you are setting up a global cluster, make sure you have licenses for the following products:

- Veritas Storage Foundation for Oracle RAC HA/DR
- Veritas Volume Replicator (VVR) if you are using it for replication

If you are using software for hardware-based replication, make sure the required licenses from your vendor are enabled.

For licensing details:

See [“Obtaining SF Oracle RAC license keys”](#) on page 75.

See [“Removing pre-existing license keys”](#) on page 76.

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

The product installation procedure includes instructions on how to activate the key. If you encounter problems while licensing this product, visit the Symantec licensing support website at:

<https://licensing.symantec.com>

Verify if HA/DR is enabled:

```
# vxlicrep | grep -i global
```

If HA/DR is not enabled, install it.

To install Veritas Storage Foundation for Oracle RAC HADR license keys

- 1 Install the HA/DR-enabled license:

```
# vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xx (CVM/VVR)
```

- 2 Make sure that HA/DR is enabled:

```
# vxlicrep | grep -i global
```

If you are using VVR, verify if it is enabled:

```
# vxlicrep | grep -i vvr
```

If VVR is not enabled, install it.

To install VVR license keys (if using VVR)

- 1 Install the VVR license:

```
# vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xx (CVM/VVR)
```

- 2 Make sure that VVR is enabled:

```
# vxlicrep | grep -i vvr
```

Verifying the systems before installation

Before beginning the installation of SF Oracle RAC software, you can check the readiness of the systems where you plan to install SF Oracle RAC.

You can also use the Veritas Installation Assessment Service utility for a detailed assessment of your setup.

See [“Veritas Installation Assessment Service”](#) on page 28.

To verify the systems

- 1 Navigate to the folder containing the installsfrac program.

```
# cd /cdrom/dvd1/sfrac
```

- 2 Start the pre-installation check:

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

The program proceeds in a non-interactive mode, examining the systems for licenses, filesets, disk space, and system-to-system communications. The program displays the results of the check and saves the results of the check in a log file.

Installing SF Oracle RAC

This chapter includes the following topics:

- [About installing SF Oracle RAC](#)
- [Invoking SF Oracle RAC installation program](#)

About installing SF Oracle RAC

You can install SF Oracle RAC on clusters of up to eight nodes. Use one of the following tools to install and configure SF Oracle RAC:

Veritas product installer	Offers a high-level approach to installing or configuring multiple Veritas products. You can use the Veritas product installer to install Symantec Product Authentication Service, and Veritas Cluster Management Console For more information on these product and features, refer to the respective product documentation.
installsfrc program	Offers a direct approach to specifically installing and configuring SF Oracle RAC. You can choose one of the following ways: ■ Install the filesets and proceed to configure SF Oracle RAC. ■ Install the filesets (<code>installsfrc -installonly</code>) and leave the cluster configuration steps for later (<code>installsfrc -configure</code>).

Note: If you obtained SF Oracle RAC from an electronic download site, you must use the `installsfrc` program instead of the Veritas product installer.

Invoking SF Oracle RAC installation program

Mount the product disc and launch the installation program. At the end of each product installation, the installer creates a new directory with the following files:

- A log file containing any system commands executed, and their output.
- A response file that can be used in conjunction with the `-responsefile` option of the installer.

See [“About Response files”](#) on page 575.

- A summary file containing the output of the install scripts.

The product installation consists of the following steps:

- [Mounting the product disc](#)
- [Starting software installation](#)
- [Specifying systems for installation](#)
- [Licensing SF Oracle RAC](#)

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 your SF Oracle RAC installation and configuration.

See [“Mounting the product disc”](#) on page 73.

Starting software installation

You can install SF Oracle RAC using one of the following two programs:

- Veritas product installer
- `installsfrac` program

For a simple install, Symantec recommends using the Veritas product installer.

Note: SSH or RSH must be set up prior to starting the software installation. If SSH or RSH is not set up and enabled, then the installation process will fail.

To install SF Oracle RAC using the product installer

- 1 Confirm that you are logged in as the superuser.
- 2 Navigate to the folder containing the installer program.

```
# cd /cdrom
```

- 3 Start the installer.

```
# ./installer
```

Use the above command if using SSH. If using RSH, then enter the following command:

```
# ./installer -rsh
```

The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.

- 4 From the opening Selection Menu, choose: "I" for "Install/Upgrade a Product."
- 5 From the displayed list of products to install, choose: Veritas Storage Foundation for Oracle RAC.

To install SF Oracle RAC using the installsfrac program

- 1 Confirm that you are logged in as the superuser.
- 2 Navigate to the folder containing the installsfrac program.

```
# cd /cdrom/dvd1/sfrac
```

- 3 Prior to starting the `installsfrac` program, review the SF Oracle RAC installation and configuration worksheets:

See [“SF Oracle RAC installation and configuration information”](#) on page 559.

Make sure that you have values ready for your installation. You will be prompted for these values during the installation process.

- 4 Start the `installsfrac` program.

```
# ./installsfrac -installonly
```

The installer begins with a copyright message and specifies the directory where the logs are created.

By default, the `installsfrac` program uses `ssh` for remote communication. However, to use `rsh`, specify the `-rsh` option with the `installsfrac` program.

```
# ./installsfrac -rsh -installonly
```

Specifying systems for installation

The installer prompts for the system names on which you want to install and then performs an initial system check.

To specify system names for installation

- 1 Enter the names of the systems where you want to install SF Oracle RAC.

Enter the system names separated by spaces on which to install SF Oracle RAC: **galaxy nebula**

- 2 Review the output as the installer checks that the local node running the script can communicate with remote nodes and checks whether a previous version of SF Oracle RAC is installed.

If a previous version of SF Oracle RAC is installed, the installer provides an option to upgrade.

See [“Supported upgrade options for SF Oracle RAC 5.0 MP3”](#) on page 145.

Licensing SF Oracle RAC

The installer checks whether SF Oracle RAC license keys are currently in place on each system. If license keys are not installed, the installer prompts you for the license keys.

To license SF Oracle RAC

- 1 Review the output as the utility checks system licensing and installs the licensing fileset.
 - If the VRTSvlic fileset is not present, the utility installs it on each node after checking for sufficient disk space.
 - If a previous version of the VRTSvlic fileset is present, the utility replaces the fileset with the current version.
- 2 Enter the license key for Veritas Storage Foundation for Oracle RAC as the installer prompts for each node.

```
Enter a SFRAC license key for galaxy: [?]  
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX  
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX successfully registered on  
galaxy  
SFRAC license registered on galaxy
```

- 3 Enter keys for additional product features such as SF Oracle RAC HADR and VVR, if you want to set up a global cluster and disaster recovery environment. You can add these keys at a later time if you are not prepared to configure these features now.

```
Do you want to enter another license key for galaxy? [y,n,q,?]  
(n) y
```

```
Enter a SFRAC license key for galaxy: [?]  
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX  
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX successfully registered on  
galaxy
```

```
Do you want to enter another license key for galaxy? [y,n,q,?]  
(n)
```

- 4 Enter the license keys for other nodes.

```
SFRAC is not licensed on nebula  
.  
.  
Do you want to enter another license key for nebula? [y,n,q,?]  
(n)  
SFRAC licensing completed successfully.
```

Installing SF Oracle RAC filesets

The installer verifies for any previously installed filesets and then based on your choice installs all the SF Oracle RAC filesets or only the required filesets.

To install SF Oracle RAC filesets

- 1 Review the output as the installer checks the filesets that are already installed.
- 2 Review the output as the installer makes sure that the required OS patches are available on all nodes.

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

- 3 Review the output as the installer checks for the required file system space. The installer also checks for the presence of processes that could conflict with the installation.
- 4 View the list of filesets that the installer would install on each node.
- 5 Verify that the installation process completes successfully. Review the output at the end of the installation and note the location of the summary and log files for future reference.
- 6 Shut down and reboot all the nodes by entering the following command on each node:

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

Configuring SF Oracle RAC

This chapter includes the following topics:

- [About configuring SF Oracle RAC](#)
- [Invoking SF Oracle RAC configuration program](#)
- [Performing basic system checks](#)
- [Configuring SF Oracle RAC components](#)

About configuring SF Oracle RAC

After you install the SF Oracle RAC filesets, you must configure SF Oracle RAC using the Veritas product installer or the installsfrac program. The Veritas product installer offers a high-level approach to install or configure multiple Veritas products. The installsfrac program offers a direct approach to specifically install or configure SF Oracle RAC.

The configuration program provides the following options:

Check system for SFRAC

Provides guidelines to verify private interconnects for LLT and to verify the disks intended for shared storage support.

See [“Performing basic system checks”](#) on page 92.

Configure SFRAC

Asks you a set of questions to configure SF Oracle RAC component products VCS, CVM, and CFS.

See [“Configuring SF Oracle RAC components”](#) on page 95.

Provides guidelines to configure I/O fencing.

See [“About configuring SF Oracle RAC clusters for data integrity”](#) on page 113.

Prepare to install Oracle	Provides guidelines and prompts you to complete pre-installation tasks for Oracle in an SF Oracle RAC environment.
Install or relink Oracle	Launches Oracle Installer to install Oracle RAC, and links Oracle with the Veritas IPC library.

Figure 5-1 illustrates the tasks involved to configure SF Oracle RAC.

Figure 5-1 Workflow to configure SF Oracle RAC




Table 5-1 lists the high-level tasks and references to the procedures to configure SF Oracle RAC.

Table 5-1 Tasks to configure SF Oracle RAC

Task	Reference
Invoke the configuration program	See “Invoking SF Oracle RAC configuration program” on page 89.
Check private interconnects for LLT	See “Preparing private interconnects for LLT” on page 93.
Configure SF Oracle RAC components	See “Configuring SF Oracle RAC components” on page 95.
Check shared disks for I/O fencing	See “About configuring SF Oracle RAC clusters for data integrity” on page 113. See “Preparing to configure I/O fencing” on page 119.
Configure I/O fencing	See “Setting up I/O fencing” on page 125.

Invoking SF Oracle RAC configuration program

If you installed SF Oracle RAC and did not choose to configure SF Oracle RAC immediately, use the `installsrac -configure` option to configure SF Oracle RAC when you are ready for cluster configuration. The `-configure` option can be used to reconfigure a SF Oracle RAC cluster. SF Oracle RAC must not be running on systems when this reconfiguration is performed.

Mount the product disc and launch the configuration program. At the end of each product configuration, the installer creates a new directory with three files:

- A log file containing any system commands executed, and their output.
- A response file that can be used in conjunction with the `-responsefile` option of the installer. See [“About Response files”](#) on page 575.
- A summary file containing the output of the install scripts.

Make sure you have superuser (root) privileges to load the SF Oracle RAC software. See [“Mounting the product disc”](#) on page 73.

Starting software configuration

After mounting the product disc, you can configure SF Oracle RAC using the Veritas product installer or the `installsrac` program.

To configure SF Oracle RAC using the product installer

- 1 Confirm that you are logged in as the superuser.
- 2 Make sure you have set the PATH variable for java/bin before executing `installsfrac -configure`:

```
# export PATH=$PATH:/usr/javaXX/bin
```

- 3 Navigate to the folder containing the installer program.

```
# cd /cdrom
```

- 4 Start the installer.

```
# ./installer
```

The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.

- 5 From the opening Selection Menu, choose: **C** for "Configure an Installed Product."
- 6 From the displayed list of products to install, choose: Veritas Storage Foundation for Oracle RAC.

To configure SF Oracle RAC using the installsfrac program

- 1 Confirm that you are logged in as the superuser.
- 2 Make sure you have set the PATH variable for java/bin before executing `installsfrac -configure`:

```
# export PATH=$PATH:/usr/javaXX/bin
```

- 3 Navigate to the folder containing the installsfrac program .

```
# cd /cdrom/dvd1/sfrac
```

- 4 Start the installsfrac program with the `-configure` option.

```
# ./installsfrac [-rsh] -configure
```

By default, the installsfrac program uses ssh for remote communication. However, to use `rsh`, specify the `-rsh` option with the installsfrac program.

The installer begins with a copyright message and specifies the directory where the logs are created.

Specifying systems for configuration

The configuration program prompts for the system names on which you want to configure and then performs an initial system check for communication.

To specify system names for configuration

- 1 Confirm that you want to proceed with the configuration at the prompt.
- 2 Enter the names of the nodes where you want to configure the software.

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

- 3 Review the output as the program checks that the local node running the script can communicate with remote nodes and checks whether SF 5.0 MP3 for Oracle RAC is installed successfully.

If SF 5.0 MP3 for Oracle RAC is not installed, the program exits.

Choosing the configuration task

The configuration program provides you menu options from which you can choose the configuration task. The menu options that are marked ****INSTRUCTIONS ONLY**** require you to:

- Follow the instructions in the order mentioned
- Keep the shells on all systems open with superuser privileges

To choose the configuration task

- 1 After the program verifies the license keys for SF Oracle RAC, review the configuration options that the installer presents.

See [“About configuring SF Oracle RAC”](#) on page 87.

- 2 Perform the following tasks to configure SF Oracle RAC cluster:

Check private interconnects for LLT.

From the main menu, select **Check systems for SFRAC > Check LLT links** for instructions to perform LLT checks.

Note: If you use raw devices for shared storage, you can also check shared disks for I/O fencing at this time.

See [“Performing basic system checks”](#) on page 92.

Configure SF Oracle RAC components.

From the main menu, select **Configure SFRAC > Configure VCS, CVM, and CFS**.

Review the output as the program stops various response files and processes. Answer the prompts to configure the different components and their optional features.

See [“Configuring SF Oracle RAC components”](#) on page 95.

Check shared disks for I/O fencing.

Launch the configuration program and select **Check systems for SFRAC > Check I/O fencing disks** for instructions to check I/O fencing disks.

See [“Preparing to configure I/O fencing”](#) on page 119.

Configure I/O fencing.

From the main menu, select **Configure SFRAC > Configure I/O fencing** for instructions to configure I/O fencing.

See [“Setting up I/O fencing”](#) on page 125.

Performing basic system checks

Choose Check systems for SFRAC from the configuration program’s menu to perform the basic system checks for SF Oracle RAC.

You can follow the SF Oracle RAC configuration program instructions to check the private interconnects for LLT and to check the shared disks for I/O fencing.

Note: If you use DMP disk devices, then you must verify the disks for I/O fencing compliance only after you configure SF Oracle RAC components.

To perform system checks

- 1 Choose one of the following from the SF Oracle RAC configuration menu:

- Perform all the following tasks
- Check LLT links ****INSTRUCTIONS ONLY****
- Check I/O fencing disks ****INSTRUCTIONS ONLY****

If you use DMP disk devices, then you must choose **Check LLT links** at this time. You must verify the disks for I/O fencing compliance only after you configure SF Oracle RAC components.

- 2 If you choose **Perform all the following tasks**, follow the guidelines as the program takes you through the other menu options individually.
- 3 If you choose **Check LLT links**, the program lists the prerequisites for the LLT links and prompts you to check each LLT link on all cluster nodes.

See [“Preparing private interconnects for LLT”](#) on page 93.

- 4 If you choose **Check I/O fencing disks**, follow the program guidelines to check I/O fencing disks on all nodes.

You must check I/O fencing disks that use Dynamic Multipathing (DMP) after you configure SF Oracle RAC components.

See [“Preparing to configure I/O fencing”](#) on page 119.

Preparing private interconnects for LLT

You must attach each network interface that is configured for LLT to a separate and distinct physical network.

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 switches or hubs used for the interconnects must match that of the interface cards. Incorrect settings can cause poor network performance or even network failure.

Guidelines for setting the media speed of the LLT interconnects

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

- If you have hubs or switches for LLT interconnects, Symantec recommends using the `Auto_Negotiation` media speed setting on each Ethernet card on each node.
- If you have hubs or switches for LLT interconnects and you do not use the `Auto_Negotiation` media speed setting, set the hub or switch port to the same setting as that used for the cards on each node.
- If you use directly connected Ethernet links, set the media speed to the highest value common to both cards, typically `100_Full_Duplex`.
- Settings on Ethernet cards for jumbo frames must match that of the switches or hubs. Symantec does not recommend use of jumbo frames in an SF Oracle RAC environment.
- Symantec does not recommend using dissimilar network cards for private links.

Displaying and setting the Ethernet media speed

Review the information about the current settings for interface cards and how to modify them if necessary.

To display and set the Ethernet auto-negotiation setting and the media speed

- 1 Use the command `lsattr E1 device_name` to display current capabilities of the interface card. For example:

```
# lsattr -R -l <device_name> -a media_speed
```

If you receive an error such as:

```
lsattr: 0514-528 The "media_speed" attribute does not exist
in the predefined device configuration database.
```

The card is a 10 Mbps card, capable of only 10 Mbps / half duplex.

Example output of a 100 Mbps card resembles:

```
10_Half_Duplex
10_Full_Duplex
100_Half_Duplex
100_Full_Duplex
Autonegotiation
```

- 2 To display the current settings of the card, enter:

```
# lsattr -EH -l <device_name> -a media_speed
```

- 3 To change the setting of the card, use the `chdev` command. For example, where the card is `ent0` and the desired speed is 100 half duplex, enter:

```
# chdev -P -l ent0 -a media_speed=100_Half_Duplex
```

The change in value does not take effect until the system is restarted or the device is rediscovered.

Configuring SF Oracle RAC components

Make sure that you gathered the required information to configure the SF Oracle RAC components. Also, make sure that you performed the pre-configuration tasks if you want to configure the cluster in secure mode, centralized cluster management, or centralized Storage Foundation management.

See [“About preparing to install and configure SF Oracle RAC”](#) on page 57.

To configure SF Oracle RAC components

- ◆ Choose **Configure SFRAC > Configure VCS, CVM and CFS** from the configuration program's menu.

Answer the configuration program prompts to configure the SF Oracle RAC components such as VCS, CVM, and CFS.

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.

Enter a cluster name and ID to perform the basic cluster configuration.

Configuring the basic cluster

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: [?] vcs_cluster2
```

```
Enter the unique Cluster ID number between 0-65535: [b,?] 7
```

- 3 Review the NICs available on the first system as the installer discovers and reports them.

The private heartbeats can either use NIC or aggregated interfaces. To use aggregated interfaces for private heartbeat, enter the name of the aggregated interface. To use a NIC for private heartbeat, enter a NIC which is not part of an aggregated interface.

4 Enter the network interface card details for the private heartbeat links.

You must choose the network interface cards or the aggregated interfaces that the installer discovers and reports. If you want to use aggregated interfaces that the installer has not discovered, then you must manually edit the `/etc/llttab` file before you start SF Oracle RAC when the installer prompts after product configuration.

See [“Starting SF Oracle RAC processes”](#) on page 109.

You must not enter the network interface card that is used for the public network (typically `en0` and `en1`.)

```
Enter the NIC for the first private heartbeat NIC on galaxy:
[b,?] en2
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,?] en3
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)
```

5 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.

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.

6 Verify and confirm the information that the installer summarizes.

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 `installsrc` program 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 58.

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 99.

- 2 Select one of the options to enable security.

Select the Security option you would like to perform [1-3,q,?]

Review the following configuration modes. Based on the configuration that you want to use, enter one of the following values:

Option 1. Automatic configuration	Enter the name of the Root Broker system when prompted. Requires a remote access to the Root Broker. Review the output as the installer verifies communication with the Root Broker system, checks vxatd process and version, and checks security domain.
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:
east.symantecexample.com
Enter root broker FQDN: [b]
(symantecexample.com)
symantecexample.com
Enter root broker domain: [b]
(root@east.symantecexample.com)
root@east.symantecexample.com
Enter root broker port: [b] (2821) 2821
Enter path to the locally accessible
root hash [b] (/var/tmp/
installvcs-1Lcljr/root_hash)
/root/root_hash
```

Enter the following Authentication Broker information as the installer prompts you for each node:

```
Enter authentication broker principal name on
galaxy [b]
(galaxy.symantecexample.com)
galaxy.symantecexample.com
Enter authentication broker password on galaxy:
Enter authentication broker principal name on
nebula [b]
(nebula.symantecexample.com)
nebula.symantecexample.com
Enter authentication broker password 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 *Veritas Cluster Server User'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 password for the Admin user
(default password='password')? [y,n,q] (n) y
```

```
Enter New 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: [?] smith
```

```
Enter New Password:*****
```

```
Enter Again:*****
```

```
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [?] 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 SF Oracle RAC 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 User'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] (y) **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 102.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server’s host name.

Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] **smtp.example.com**

- Enter the email address of each recipient.

Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] **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,?] **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,?] **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,?] **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.

```
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 SF Oracle RAC 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 User's Guide* for more information.

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of SF Oracle RAC.
- 2 Specify whether you want to configure the SNMP notification.

```
Do you want to configure SNMP notification? [y,n,q] (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 with other set of questions for CVM and CFS.

See [“Configuring global clusters”](#) on page 103.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the SNMP trap daemon port.

```
Enter the SNMP trap daemon port: [b,?] (162)
```

- Enter the SNMP console system name.

```
Enter the SNMP console system name: [b,?] 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,?] 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,?] 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,?] 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.

```
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. Note that you can also run the `goconfig` utility in each cluster later to update the VCS configuration file for global cluster.

See *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for instructions to set up SF Oracle RAC global clusters.

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] (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.

- 4 Verify and confirm the configuration of the global cluster.

```
Global Cluster Option configuration verification:
```

```
NIC: en0 and en1
```

```
IP: 10.10.12.1
```

```
Netmask: 255.255.240.0
```

```
Is this information correct? [y,n,q] (y)
```

Setting permissions for database administration

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

See [“Configuring the Cluster Volume Manager”](#) on page 105.

To set permissions for database administration

- 1 Review the required information to set up the permissions for database administration.
- 2 Specify whether you want to add single user access, group access, or both on each of the nodes as the installer prompts.
 - Provide information if you want to add single user access.


```
Do you want to add single user access on galaxy [y,n,q,?] (y)
Enter login account name for DBA user: dba
```

- Provide information if you want to add group access.

```
Do you want to add group access on galaxy [y,n,q,?] (y)
Enter group name for DBA users: oper
```

Configuring the Cluster Volume Manager

Cluster volume manager configuration tasks include:

- [Setting up naming scheme](#)
- [Setting up default disk group](#)

Setting up naming scheme

Disks on AIX systems typically use device names such as `/dev/rhdisk75` to identify disks on the system. It is possible to use the VxVM enclosure-based naming scheme, which allows disk arrays to be more readily recognizable. Dynamic Multipathing (DMP) is a prerequisite for enclosure-based naming schemes.

Note: If the disk path is too long, then the installer sets up the enclosure-based naming scheme even though you enter `n`.

Refer to the Veritas Volume Manager documentation for details on this scheme.

To set up the naming scheme

- 1 If you want to set up the enclosure-based naming scheme, enter `y`.

```
Do you want to set up the enclosure-based naming scheme?
[y,n,q,?] (n)
```

- 2 Specify if you want to use the new naming scheme for all eligible systems.

```
Do you want to use the enclosure-based naming scheme for all of
the eligible systems? [y,n,q,?] (y)
```

Setting up default disk group

If applicable, set up the default disk group. Because some VxVM commands require that a disk group be specified, the installer enables you to register the name of a

default VxVM disk group on each eligible node. Note that you can create the default disk group later.

To set up the default disk group procedure

- 1 If you want to set up a default disk group, enter `y`.

```
Do you want to set up a system wide default disk group?
[y,n,q,?] (y) y
```

```
Which disk group? [<group>,list,q,?] default_dg
```

- 2 If you specified setting up a default disk group, review the setup output.

```
Volume Manager default disk group setup and daemon startup
```

```
Setting default diskgroup to default_dg on galaxy ..... Done
Starting vxrelocd on galaxy ..... Started
Starting vxcached on galaxy ..... Started
Starting vxconfigbackupd on galaxy ..... Started
.
.
```

Configuring VVR on each node

If you added license for Veritas Volume Replicator during installation, you can now accept the default settings or modify the settings for VVR. The installer prompts you for the information on each node.

Setting up VVR ports

The installer identifies the default ports that would be used VVR. You can also assign different ports at this point. Note that the port settings must be identical for systems that will be part of the same Replicated Data Set. They must also be identical for all the systems in a cluster.

To set up VVR ports

- 1 Review and accept the default port values that the configuration program displays.

Following are the default ports that will be used by VVR on galaxy:

.
.

Do you want to change any of the VVR ports on galaxy? [y,n,q]
 (n)

- 2 If you want to change any of the VVR ports on the system, enter y.

Do you want to change any of the VVR ports on galaxy? [y,n,q]
 (n) **y**

- 3 Follow the instructions to change the port values. Note the following points:

- The port settings must be identical for systems that will be part of the same Replicated Data Set.
- The port settings must also be identical for all the systems in a cluster.

Configuring VVR statistics collector

The VVR administrative daemon vradmind collects and maintains various statistics, which are helpful in solving VVR performance issues. You can tune the collection using a few tunables:

frequency	for gathering the statistics default = 10 seconds
number of days	for which the collected statistics should be preserved, after which the earlier statistics are automatically deleted default = 3 days

The installation program provides an option to change the default settings.

To configure VVR statistics collector

- 1 Enter y at the prompt to change the default setting.
- 2 Enter the values when the installer prompts.

Configuring VVR tunables

As an advanced user, you can modify the VVR tunable parameters.

Refer to Veritas Volume Replicator documentation for more details.

To configure VVR tunables

- 1 Enter `y` to view or modify the VVR tunables.
- 2 Review the output to find whether the configuration is successful.

Creating SF Oracle RAC configuration files

After gathering the configuration information for SF Oracle RAC components, the configuration program creates the 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.

To create SF Oracle RAC configuration files

- 1 Review the output as the configuration program creates security principal, starts VCS, creates VCS configuration files, and copies the files to each node.
- 2 When the configuration prompts you, confirm the fully qualified host names of the cluster nodes.

```
Is the fully qualified hostname of system "galaxy" =  
"galaxy.example.com"? [y,n,q] (y)  
Is the fully qualified hostname of system "nebula" =  
"nebula.example.com"? [y,n,q] (y)
```

- 3 Review the output as the program verifies communication with the remote nodes.

Configuring standalone hosts for Storage Foundation

When the configuration program prompts you to enable Storage Foundation Manager Management, enter `n` to configure the standalone hosts for Storage Foundation.

Warning: If you press Enter to accept the default option, you must restart the configuration program. Storage Foundation Manager is not available on the Storage Foundation and High Availability Solutions release. Make sure to enter **n** at the prompt and configure standalone hosts.

To configure standalone hosts for Storage Foundation

- 1 Review the output on the ways to manage the Storage Foundation hosts.
- 2 Enter **n** at the prompt to configure standalone hosts.

```
Enable Storage Foundation Manager Management? [y,n,q]
(y) n
```

- 3 Review the output as the configuration program configures to manage the Storage Foundation hosts in standalone mode.

Starting SF Oracle RAC processes

The installer starts SF Oracle RAC, SF Oracle RAC processes, and configures the agents that you selected during the configuration phase.

Note: To use aggregated interfaces that the installer has not discovered for private heartbeats, do not opt to start SF Oracle RAC.

See [“Configuring the basic cluster”](#) on page 96.

To start SF Oracle RAC processes

- 1 Confirm that you want to start the SF Oracle RAC processes.

Enter **y** to start SF Oracle RAC processes, proceed to step 3.

If you want to use aggregated interfaces that the installer has not discovered for private heartbeats, enter **n**. Skip to step 2

```
Do you want to start Veritas Cluster Server processes now?
[y,n,q] (y) n
```

- 2 Do the following to use aggregated interfaces for private heartbeats:
 - Edit the `/etc/llttab` file to replace the names of NICs with the names of the aggregated interfaces.
 - Reboot the system for the configuration changes to take effect.

- 3 Review the output as the installer starts SF Oracle RAC and its 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.
See [“About setting up I/O fencing”](#) on page 116.
- 4 Review the output as the program starts the VxVM daemons and configures the SF Oracle RAC agents.
- 5 Review the output as the configuration program sets the default disk group on each node. You must have a specified a default disk group for VxVM during the configuration,
- 6 Review the output at the end of the configuration and note the location of the summary and log files for future reference.

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.

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)

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. To enable LDAP authentication plug-in, you must verify the LDAP environment, add the LDAP domain in AT, and then verify LDAP authentication. The AT component packaged with SF Oracle RAC requires you to manually edit the VRTSatlocal.conf file to enable LDAP authentication.

Refer to the *Symantec Product Authentication Service Administrator's Guide* for instructions.

If you have not already added VCS users during installation, you can add the users later.

See *Veritas Cluster Server User's Guide* for instructions to add VCS users.

Configuring SF Oracle RAC clusters for data integrity

This chapter includes the following topics:

- [About configuring SF Oracle RAC clusters for data integrity](#)
- [About I/O fencing components](#)
- [About setting up I/O fencing](#)
- [Preparing to configure I/O fencing](#)
- [Setting up I/O fencing](#)

About configuring SF Oracle RAC clusters for data integrity

When a node fails, SF Oracle RAC takes corrective action and configures its components to reflect the altered membership. If an actual node failure did not occur and if the symptoms were identical to those of a failed node, then such corrective action would cause a split-brain situation.

Some example scenarios that can cause such split-brain situations are as follows:

- **Broken set of private networks**
If a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner

- **System that appears to have a system-hang**

If a system is so busy that it appears to stop responding, the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

I/O fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster. SF Oracle RAC uses I/O fencing to remove the risk that is associated with split brain. I/O fencing allows write access for members of the active cluster. It blocks access to storage from non-members so that even a node that is alive is unable to cause damage.

After you install and configure SF Oracle RAC, you must configure I/O fencing in SF Oracle RAC to ensure data integrity.

About I/O fencing components

The shared storage for SF Oracle RAC must support SCSI-3 persistent reservations to enable I/O fencing. SF Oracle RAC involves two types of shared storage:

Data disks	Store shared data
Coordination points	Act as a global lock during membership changes

I/O fencing in SF Oracle RAC involves coordination points and data disks. Each component has a unique purpose and uses different physical disk devices. The fencing driver, known as vxfen, directs CVM as necessary to carry out actual fencing operations at the disk group level.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs). These disks must support SCSI-3 PR and are part of standard VxVM or CVM disk groups.

CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. Racing for control of the coordination points to fence data disks is the key to understand how fencing prevents split brain.

Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the SF Oracle RAC configuration. Users cannot store data on these disks or include the disks in a disk group for user data. The coordinator disks can be any three disks that support SCSI-3 PR. Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

Symantec recommends using the smallest possible LUNs for coordinator disks. Because coordinator disks do not store any data, cluster nodes need to only register with them and do not need to reserve them.

You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature. Dynamic Multipathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use either DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is raw by default.

See the *Veritas Volume Manager Administrator's Guide*.

You can use iSCSI devices as coordinator disks for I/O fencing. However, I/O fencing supports iSCSI devices only when you use DMP disk policy. If you use iSCSI devices as coordinator disks, make sure that the `/etc/vxfsenmode` file has the disk policy set to DMP.

For the latest information on supported hardware visit the following URL:

<http://entsupport.symantec.com/docs/283161>

When you configure I/O fencing in raw mode, I/O fencing disables secondary paths for DMP. I/O fencing disables secondary paths for DMP because for certain A/P arrays, DMP opens secondary paths with an exclusive flag, and keeps all paths permanently open. In raw mode, however, I/O fencing needs to open all paths and register keys down those paths.

Due to DMP's exclusive hold on the paths, the `vxfsenconfig` command cannot open those paths. To avoid this issue on the secondary paths, I/O fencing disables the secondary paths from DMP's perspective. This act causes DMP to close the

paths. This disabling of secondary paths is done using the `/etc/init.d/vxfen.rc` script.

During a cluster reconfiguration, if a primary path is not available for any reason, I/O fencing tries to use secondary paths for the operation. The path was disabled from DMP's perspective, but all SCSI-3 operations are performed to the path's raw interface.


When you try to run the `vxfcntlshdw` command, remember that DMP has an exclusive hold on the path, so run it only on primary paths. Use the `vxdisk list` command to determine the primary path to a disk.

About setting up I/O fencing


[Figure 6-1](#) illustrates the tasks involved to configure I/O fencing.

Figure 6-1 Workflow to configure I/O fencing

Preparing to set up I/O fencing



Setting up I/O fencing



See [“Preparing to configure I/O fencing”](#) on page 119.

See [“Setting up I/O fencing”](#) on page 125.

I/O fencing requires the coordinator disks be configured in a disk group. The coordinator disks must be accessible to each node in the cluster. These disks enable the vxfs driver to resolve potential split-brain conditions and prevent data corruption.

Review the following requirements for coordinator disks:

- You must have three coordinator disks.
- Each of the coordinator disks must use a physically separate disk or LUN.

- Each of the coordinator disks should exist on a different disk array, if possible.
- You must initialize each disk as a VxVM disk.
- The coordinator disks must support SCSI-3 persistent reservations.
- The coordinator disks must exist in a disk group (for example, vxfencoorddg).
- Symantec recommends using hardware-based mirroring for coordinator disks.

The I/O fencing configuration files include:

<code>/etc/vxfendg</code>	You must create this file to include the coordinator disk group information.
<code>/etc/vxfenmode</code>	You must set the I/O fencing mode to SCSI-3. You can configure the vxfen module to use either DMP devices or the underlying raw character devices. Note that you must use the same SCSI-3 disk policy on all the nodes. The SCSI-3 disk policy can either be raw or dmp. The policy is raw by default.
<code>/etc/vxfentab</code>	When you run the vxfen startup file to start I/O fencing, the script creates this <code>/etc/vxfentab</code> file on each node with a list of all paths to each coordinator disk. The startup script uses the contents of the <code>/etc/vxfendg</code> and <code>/etc/vxfenmode</code> files. Thus any time a system is rebooted, the fencing driver reinitializes the <code>vxfentab</code> file with the current list of all paths to the coordinator disks. Note: The <code>/etc/vxfentab</code> file is a generated file; do not modify this file.

An example of the `/etc/vxfentab` file on one node resembles as follows:

■ Raw disk:

```
/dev/rhdisk75  
/dev/rhdisk76  
/dev/rhdisk77
```

■ DMP disk:

```
/dev/vx/rdmp/rhdisk75  
/dev/vx/rdmp/rhdisk76  
/dev/vx/rdmp/rhdisk77
```

In some cases you must remove disks from or add disks to an existing coordinator disk group.

Warning: If you remove disks from an existing coordinator disk group, then be sure to remove the registration and reservation keys from these disks before you add the disks to another disk group.

About setting up shared storage for I/O fencing

You need to set up shared storage so that it is visible to the SCSI layer from all the nodes in the cluster. The shared storage that you add for use with SF Oracle RAC software must support SCSI-3 persistent reservations, a functionality that enables the use of I/O fencing.

Make sure that your system has at least SCSI-3 enabled disks set up and ready.

See [“Checking shared disks for I/O fencing”](#) on page 122.

Preparing to configure I/O fencing

The shared storage that you add for use with SF Oracle RAC software must support SCSI-3 persistent reservations, a functionality that enables the use of I/O fencing.

Perform the following preparatory tasks to configure I/O fencing:

Initialize disks as VxVM disks	See “Initializing disks as VxVM disks” on page 119.
Identify disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 121.
Check shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 122.
The tasks involved in checking the shared disks for I/O fencing are as follows:	
■ Verify that the nodes have access to the same disk ■ Test the disks using the vxfsentsthdw utility	

Initializing disks as VxVM disks

Install the driver and HBA card. Refer to the documentation from the vendor for instructions.

After you physically add shared disks to the nodes, you must do the following:

- Initialize them as VxVM disks
- Verify that all the nodes see the same disk

See the *Veritas Volume Manager Administrator's Guide* for more information on how to add and configure disks.

To initialize disks as VxVM disks

- 1 Scan for the new hdisk devices.

```
# /usr/sbin/cfgmgr
```

- 2 Make the new disks recognizable. On each node, enter:

```
# lsdev -Cc disk
```

- 3 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.

- 4 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
=====
libvxCLARiON.so                       DGC
libvxcscovrts.so                      CSCOVRTS
libvxemc.so                           EMC
```

- 5 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.

6 Determine the VxVM name by which a disk drive (or LUN) is known.

In the following example, VxVM identifies a disk with the AIX device name `/dev/rhdisk75` as `EMC0_17`:

```
# vxddmpadm getdmpnode nodename=rhdisk75
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
EMC0_17	ENABLED	EMC	1	1	0	EMC0

Notice that in the example command, the AIX device name for the block device was used.

As an option, you can run the command `vxdisk list vxvm_device_name` to see additional information about the disk like the AIX device name. For example:

```
# vxdisk list EMC0_17
```

7 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 Managers 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 EMC0_17
```

Repeat this command for each disk you intend to use as a coordinator disk.

Identifying disks to use as coordinator disks

After you add and initialize disks, 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 list
```

2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

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 `vxfcntlsthaw` 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 `vxfcntlsthaw` 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 *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Checking that disks support SCSI-3 involves the following tasks:

- Verifying that nodes have access to the same disk
 See [“Verifying that the nodes have access to the same disk”](#) on page 122.
- Testing the shared disks for SCSI-3
 See [“Testing the disks using vxfcntlsthaw utility”](#) on page 123.

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 `vxfcntlsthaw` 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 `vxfenadm` command to verify the disk serial number.

```
/sbin/vxfenadm -i diskpath
```

For A/P arrays, run the `vxfcntlsthaw` command only on secondary paths.

Refer to the `vxfenadm` (1M) manual page.

For example, an EMC disk is accessible by the `/dev/rhdisk75` path on node A and the `/dev/rhdisk76` path on node B.

From node A, enter:

```
# /sbin/vxfsadm -i /dev/rhdisk75

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/rhdisk76 path.

On a disk from another manufacturer, Hitachi Data Systems, the output is different and may resemble:

```
# /sbin/vxfsadm -i /dev/rhdisk77

Vendor id      : HITACHI
Product id     : OPEN-3
Revision       : 0117
Serial Number  : 0401EB6F0002
```

Testing the disks using vxfssthdw utility

This procedure uses the /dev/rhdisk75 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 vxfssthdw utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rhdisk75 is ready to be configured for I/O Fencing on
node galaxy
```

To test the disks using vxfssthdw utility

- 1 Make sure system-to-system communication functions properly.

See [“Setting up inter-system communication”](#) on page 68.

After you complete the testing process, remove permissions for communication and restore public network connections.

See [“Removing permissions for communication”](#) on page 130.

- 2 From one node, start the utility.

Do one of the following:

- If you use `ssh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw
```

- If you use `rsh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -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 galaxy in
the format: /dev/rhdiskx
/dev/rhdisk75
Enter the disk name to be checked for SCSI-3 PGR on node nebula in
the format: /dev/rhdiskx
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/rhdisk75
```

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:

```
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

```
ALL tests on the disk /dev/rhdisk75 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

7 Run the vxsfentsthdw utility for each disk you intend to verify.

Setting up I/O fencing

Make sure you completed the preparatory tasks before you set up I/O fencing.

Tasks that are involved in setting up I/O fencing include:

Table 6-1 Tasks to set up I/O fencing

Action	Description
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 125.
Configuring I/O fencing	See “Configuring I/O fencing” on page 126.
Modifying SF Oracle RAC configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 127.
Start SF Oracle RAC on all nodes	See “Starting SF Oracle RAC on all nodes” on page 128.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 130.

Setting up coordinator disk groups

From one node, create a disk group named vxfencoorddg. This group must contain three disks or LUNs. If you use VxVM 5.0 or later, 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 EMC0_12, EMC0_16, and EMC0_17.

To create the vxfencoorddg disk group

- 1 On any node, create the disk group by specifying the device names:

```
# vxdg init vxfencoorddg EMC0_12 EMC0_16 EMC0_17
```

- 2 If you use VxVM 5.0 or later, set the coordinator attribute value as "on" for the coordinator disk group.

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

Configuring I/O fencing

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 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

- 2 Import the disk group with the `-t` option to avoid automatically importing it when the nodes restart:

```
# vxdg -t import vxfencoorddg
```

- 3 Deport the disk group. Deporting the disk group prevents the coordinator disks from serving other purposes:

```
# vxdg deport vxfencoorddg
```

- 4 On each nodes, type:

```
# echo "vxencoorddg" > /etc/vxfendg
```

Do not use spaces between the quotes in the "vxencoorddg" text.

This command creates the /etc/vxfendg file, which includes the name of the coordinator disk group.

- 5 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
```

- 6 To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes. For example:

```
# more /etc/vxfenmode
```

Modifying VCS configuration to use I/O fencing

After you add coordinator disks 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:

```
# hstop -all
```

- 3 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

- 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
)
```

- 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:

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

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:

```
# hstart
```

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 139.

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 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

- ◆ On all nodes, type:

```
# hagrps -state cvm
```

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

- ◆ On one of the nodes, type:

```
# vxfenadm -d
```

```
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)
```

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.

Verifying cluster installation

This chapter includes the following topics:

- [Verifying SF Oracle RAC installation using VCS configuration file](#)
- [About the LLT and GAB configuration files](#)
- [About VCS configuration file after SF Oracle RAC installation](#)
- [Verifying LLT, GAB, and cluster operation](#)

Verifying SF Oracle RAC installation using VCS configuration file

You can verify the SF Oracle RAC installation and configuration by examining the VCS configuration file, `main.cf`, in the directory `/etc/VRTSvcs/conf/config`.

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.

The information that these LLT and GAB configuration files contain is as follows:

- **The `/etc/llthosts` file**
The file `llthosts` 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 is identical on each node in the cluster.
For example, the file `/etc/llthosts` contains the entries that resemble:

```
0      galaxy
1      nebula
```

■ The `/etc/llttab` file

The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the private network links that correspond to the specific system.

For example, the file `/etc/llttab` contains the entries that resemble:

```
set-node galaxy
set-cluster 2
link en1 /dev/en:1 - ether - -
link en2 /dev/en:2 - ether - -
```

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 `link` command. These lines identify the two network cards that the LLT protocol uses.

Refer to the `llttab(4)` manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the `llttab` file.

■ The `/etc/gabtab` file

After you install SF Oracle RAC, the file `/etc/gabtab` contains a `gabconfig(1)` command that configures the GAB driver for use.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

The `-c` option configures the driver for use. The `-nN` specifies that the cluster is not formed until at least *N* nodes are ready to form the cluster. By default, *N* is the number of nodes in the cluster.

Note: The use of the `-c -x` option for `/sbin/gabconfig` is not recommended. The Gigabit Ethernet controller does not support the use of `-c -x`.

About VCS configuration file after SF Oracle RAC installation

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 types files for VCS (types.cf), CFS (CFSTypes.cf), CVM (CVMTypes.cf), PrivNIC (PrivNIC.cf), MultiPrivNIC (MultiPrivNIC.cf), Oracle Enterprise agent (OracleTypes.cf), and Oracle ASM (OracleASMTypes.cf). The files are located in the /etc/VRTSvcs/conf/config directory. These files define the agents that control the resources in the cluster.
- The VCS types file (types.cf) includes all agents bundled with VCS. Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for information about VCS agents.
- The CFS types file (CFSTypes.cf) includes the CFSMount agent and CFSfsckd types. The CFSMount agent mounts and unmounts the shared volume file systems. The CFSfsckd types are defined for cluster file system daemons and do not require user configuration.
- The CVM types file (CVMTypes.cf) includes definitions for the CVMCluster agent, CVMVxconfigd agent, and the CVMVolDg agent. The CVMCluster agent, which is automatically configured during installation, starts CVM in the cluster by automatically importing shared disk groups, controls node membership in the cluster, and defines how nodes communicate the state of volumes. The CVMVxconfigd agent starts and monitors the vxconfigd daemon; this daemon maintains disk and disk group configurations, communicates configuration changes to the kernel, and modifies configuration information stored on disks. The CVMVolDg agent starts the volumes in a shared disk group, sets activation modes, and monitors specified critical volumes.
- The PrivNIC types file (PrivNIC.cf) includes definitions for the PrivNIC agent. The PrivNIC agent queries LLT to count the number of visible nodes on each of the LLT Ethernet interfaces.
- The MultiPrivNIC types file (MultiPrivNIC.cf) includes definitions for the MultiPrivNIC agent. The MultiPrivNIC agent provides resources for UDP/IP support for Oracle RAC 11g cache fusion capabilities.

- The Oracle Enterprise agent types file (OracleTypes.cf) includes definitions for the Oracle agent and the Netlsnr agent. The Oracle agent monitors the resources for an Oracle database and the Netlsnr agent manages the resources for the listener process.
- The Oracle ASM types file (OracleASMTypes.cf) includes definitions for the ASM Instance agent and the ASM DG agent.
The ASM Instance agent monitors the ASM instance, brings the ASM instance online, and takes it offline.
The ASM DG agent mounts, unmounts, and monitors ASM disk groups required for ASM-enabled databases.
- 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 I/O fencing”](#) on page 125.
- The main.cf now includes the cvm service group. This group includes definitions for monitoring the CFS and 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.

Sample main.cf file after SF Oracle RAC installation

The configuration file created on each node is called main.cf and is located at: /etc/VRTSvcs/conf/config/main.cf.

The main.cf file contains the minimum information that defines the cluster and its nodes. Depending upon the VCS configuration type, the main.cf file contains different elements.

Note: Review the main.cf configuration file after the SF Oracle RAC installation and before the Oracle installation.

Note the following important points when working with the main.cf file :

- The cluster definition within the main.cf includes cluster information that was provided during the configuration. This 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, then the main.cf includes the VxSS service group and “SecureClus = 1” cluster attribute.

Review the sample main.cf configuration file below for the file syntax and different file parameters.

Note: Additional examples of different sample main.cf files are located in the appendix. See [“About sample main.cf files”](#) on page 507.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "PrivNIC.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)

system nebula (
)

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)
CFSfsckd vxfsckd (
)

CVMcluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)
```

```
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```

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.
See [“Setting the PATH variable”](#) on page 74.
- 3 Verify LLT operation.
See [“Verifying LLT”](#) on page 136.
- 4 Verify GAB operation.
See [“Verifying GAB”](#) on page 139.
- 5 Verify the cluster operation.
See [“Verifying the cluster”](#) on page 140.

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.

- 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 | more
```

The output on galaxy resembles:

Node	State	Link	Status	Address
*0 galaxy	OPEN			
		en1	UP	08:00:20:93:0E:34
		en2	UP	08:00:20:93:0E:34
1 nebula	OPEN			
		en1	UP	08:00:20:8F:D1:F2
		en2	DOWN	
2	CONNWAIT			

```

en1 DOWN
en2 DOWN
3          CONNWAIT
en1 DOWN
en2 DOWN
.
.
.
31         CONNWAIT
en1 DOWN
/dev/en:2 DOWN
```

Note that the output lists 32 nodes. The command reports the status on the two nodes in the cluster, galaxy and nebula, along with the details for the non-existent nodes.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An 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.

- 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 ... 28 29 30 31
        connects: 0 1
  7      gab        0x7
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
  31     gab        0x1F
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        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 7-1](#) lists the different ports that the software configures for different functions.

Table 7-1 GAB port description

Port	Function
a	GAB
b	I/O fencing
d	Oracle Disk Manager (ODM)
f	Cluster File System (CFS)

Table 7-1 GAB port description (continued)

Port	Function
h	Veritas Cluster Server (VCS: High Availability Daemon)
o	VCSMM driver
v	Cluster Volume Manager (CVM)
w	vxconfigd (module for CVM)

For more information on GAB, refer to the *Veritas Cluster Server User'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 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 User'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:

```
hastatus -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  ClusterService  galaxy                Y        N                ONLINE
B  ClusterService  nebula                Y        N                OFFLINE
```

- 2 Review the command output for the following information:

- The system state

If the value of the system state is RUNNING, VCS is successfully installed and started.

- The ClusterService group state

In the sample output, the group state lists the ClusterService group, which is ONLINE on galaxy and OFFLINE on 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 User's Guide* for information about the system attributes for VCS.

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	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	117
galaxy	ConfigChecksum	10844
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcS/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Fri May 26 17:22:48 2006
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.0.30.0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO

#System	Attribute	Value
galaxy	Limits	
galaxy	LinkHbStatus	en1 UP en2 UP
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	120
galaxy	SourceFile	./main.cf
galaxy	SysInfo	Aix:galaxy,5,2,00023BDA4C00
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

Upgrading to SF Oracle RAC 5.0 MP3

This chapter includes the following topics:

- [Supported upgrade options for SF Oracle RAC 5.0 MP3](#)
- [Overview of tasks for upgrading SF Oracle RAC](#)
- [Preparing to upgrade SF Oracle RAC to 5.0 MP3](#)
- [Upgrading to SF Oracle RAC 5.0 MP3](#)
- [Performing post-upgrade tasks for SF Oracle RAC 5.0 MP3](#)
- [Minimum downtime upgrade from SF Oracle RAC 5.0 to 5.0 MP3](#)
- [Upgrading licenses to SF Oracle RAC level](#)
- [Upgrading the repository database](#)

Supported upgrade options for SF Oracle RAC 5.0 MP3

The following versions of SF Oracle RAC can be upgraded to SF Oracle RAC 5.0 MP3.

Table 8-1 Supported upgrade paths

Current SF Oracle RAC version	Supported upgrade options for SF Oracle RAC 5.0 MP3
SF Oracle RAC version 4.0 MP4, 5.0x	Direct upgrade to SF Oracle RAC 5.0 MP3

Table 8-1 Supported upgrade paths (continued)

Current SF Oracle RAC version	Supported upgrade options for SF Oracle RAC 5.0 MP3
SF Oracle RAC versions 5.0, 5.0 MP1	Minimum downtime upgrade to SF Oracle RAC 5.0 MP3

If you are currently using a version of SF Oracle RAC that is not supported for a direct upgrade, you must upgrade to SF Oracle RAC 5.0 or another supported version before upgrading to 5.0 MP3. SF Oracle RAC 5.0 is provided on a disc in the 5.0 MP3 software package.

To upgrade to SF Oracle RAC 5.0:

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

Caution: SF Oracle RAC software must be at the same version across all nodes in an SF Oracle RAC cluster, in this case 5.0MP3++. All components and point products of SF Oracle RAC must be from the same version of the release in this case 5.0MP3++.

Symantec strongly recommends upgrading all SF Oracle RAC component products to the same version at the same time. In a CVM/CFS environment, dependencies between SF Oracle RAC component products will not be met if you do not upgrade all components to the same version.

Overview of tasks for upgrading SF Oracle RAC

Depending on the upgrade scenario, upgrading involves performing the following tasks:

Table 8-2 Upgrade tasks

Current SF Oracle RAC version	Upgrade tasks for SF Oracle RAC 5.0 MP3
SF Oracle RAC version 4.0 MP4, 5.0x	Upgrade tasks: ■ Consult the Late Breaking News TechNote: http://entsupport.symantec.com/docs/280186.htm ■ Upgrade your operating system if you are using an unsupported version. See “Supported software” on page 38. ■ Prepare to upgrade to SF Oracle RAC 5.0 MP3: See “Preparing to upgrade SF Oracle RAC to 5.0 MP3” on page 147. ■ Upgrade to SF Oracle RAC 5.0 MP3: See “Upgrading to SF Oracle RAC 5.0 MP3” on page 151. ■ Perform post-upgrade tasks: See “Performing post-upgrade tasks for SF Oracle RAC 5.0 MP3” on page 157.
SF Oracle RAC versions 5.0, 5.0 MP1	Minimum downtime upgrade tasks: See “Minimum downtime upgrade from SF Oracle RAC 5.0 to 5.0 MP3” on page 164.

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

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

Preparing to upgrade SF Oracle RAC to 5.0 MP3

Before you install SF Oracle RAC Maintenance Pack 3, you must stop Oracle RAC and the cluster resources.

To stop Oracle RAC and the cluster resources

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

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

- 3 For Oracle RAC 9i, stop all Oracle RAC resources including the database on all nodes.

- Stop `gsd`. On each node, log in as Oracle user and enter:

```
$ $ORACLE_HOME/bin/gsdctl stop
```

- Offline all VCS service groups that contain resources for managing Oracle listener process or the database instances. As root user, enter:

```
# hagrps -offline group_name -any
```

- 4 For Oracle RAC 10g, stop all resources.

- If the database instances are under VCS control, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline group_name -any
```

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

```
$ srvctl stop database -d database_name
```

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

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

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

- 5 For Oracle RAC 10g or Oracle RAC 11g, stop CRS on each node in the cluster. Symantec highly recommends putting CRS under VCS control.

- If CRS is controlled by VCS, log in as superuser on any system in the cluster and enter:

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

- If CRS is not controlled by VCS, log in as superuser on each system in the cluster and enter:

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

This may take a minute or two, because the CRS process takes some time to stop in the background after the command is given.

- 6 Take offline all other VCS groups that depend on VxFS file systems or VxVM disk groups, whether local or CFS.
- 7 Stop all applications that use VxFS or VxVM but that are not under VCS control.
- 8 Make sure that no processes are running which make use of mounted shared file system or shared volumes.
- 9 Take offline all VCS groups that contain CFSSMount and CVMVolDg.

```
# hagrps -offline group -sys galaxy
# hagrps -offline group -sys nebula
```

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

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

- 11 Deport any CVM disk groups that are not under VCS control.

- 12** Set the VCS resource attribute, `AutoStart`, to 0 for all VCS resources that manage the Oracle RAC 10g or Oracle RAC 11g CRS or Oracle RAC database instances. From any system in the cluster, enter:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 0

# hagrps -modify cvm_group AutoStart 0
# haconf -dump -makero
```

- 13** If you are upgrading from SF Oracle RAC 5.0x, run the following command on all nodes in the SF Oracle RAC cluster:

```
# hastop -all
```

- 14** For Oracle RAC 10g or Oracle RAC 11g, comment out the CRS-specific lines in the `/etc/inittab` file to prevent the system from starting the CRS daemons during system startup. On each system, prefix the CRS lines with a colon (":") to comment them out.

For example:

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

- 15** On each system, invoke “`slibclean`” to unload all unused modules:

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

- 16** If you are upgrading from SF Oracle RAC 4.1 or earlier, unload the `qlog` drivers on all nodes:

```
# modinfo | grep qlog
236      1341f18  17e10 319      1  qlog (VxQLOG 4.1_REV-MP2b QuickLog dr)
# modunload -i 236
```

- 17** Stop Cluster File System on all nodes using the following command:

```
# /opt/VRTSvxfs/sbin/fsclustadm cfsdeinit
```

- 18** For an optional port verification step, make sure only ports a, b, d, and o are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port o gen 6b5905 membership 01
```

- 19** For an optional disk group verification step, make sure no disk groups are imported:

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

Before upgrading, make sure your file systems are clean.

To make sure the file system is clean

- 1** Verify the VCS service groups with CVMVolDg and CFSMount are offline.

```
# hagrps -state group
```

- 2** Check and repair each file system:

```
# fsck -V vxfs /dev/vx/dsk/diskgroup/volume
```

Upgrading to SF Oracle RAC 5.0 MP3

To upgrade SF Oracle RAC:

- From SF Oracle RAC 5.0, 5.0 MP1, and 5.0 MP1 Update 1 for AIX 6.1:
See [“Upgrading from SF Oracle RAC 5.0x using the installmp installer”](#) on page 152.
- From SF Oracle RAC 4.0 MP4 to 5.0 MP3 using the `installsfrac` installer.
See [“Upgrading from pre-5.0 versions of SF Oracle RAC using the installsfrac installer”](#) on page 154.

- If you are using an earlier version of SF Oracle RAC, you must use the *Veritas Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide* to install SF Oracle RAC 5.0 before you can upgrade to SF Oracle RAC 5.0 MP3.

Warning: After running the `installmp` script, the filesets are in an applied state. Once the filesets are committed, you cannot roll back the 5.0 Maintenance Pack 3 product to the previously installed release.

Upgrading from SF Oracle RAC 5.0x using the installmp installer

If you have a working SF Oracle RAC 5.0 or 5.0 MP1 cluster you can add the patches using the `installmp` installer.

To upgrade to SF Oracle RAC 5.0 MP3 using the installmp installer

- 1 Insert the disc containing the Veritas SF Oracle RAC 5.0 MP3 software in a disc drive connected to one of the nodes for installation.
- 2 Log in as superuser on one of the systems for installation.
- 3 Mount the software (5.0 MP3) disc.

```
# mkdir /cdrom

# mount -V cdrfs -o ro /dev/cd0 /cdrom
```

- 4 Browse to the directory containing installer program.

```
# cd /cdrom
```

Where `/cdrom` is the software disc mount point.

- 5 Stop VCS:

```
# hastop -all
```

Wait for VCS to stop.

- 6 Unmount all VxFS filesystems.

Make sure all VxFS filesystems are unmounted.

- 7 Start the SF Oracle RAC installer utility.

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

- To install using SSH (default):

```
# ./installmp galaxy nebula
```


- To install using RSH:

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

- 8 Review the output as the installer checks that the systems are ready for installation. At the conclusion of these initial checks, and before installing any software, the installer does the following:

- Creates a log file (installmp-xxxxxx) on each system in the directory /var/tmp.
- Specifies the utilities it uses to communicate with the remote systems; typically this is ssh -x and scp
- Lists the filesets that it will install on each system

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

```
installmp is now ready to install MP3.
Required 5.0 processes that are currently running will be
stopped.
Are you sure you want to install MP3? [y,n,q] (y)
```

Press Enter to continue.

- 9 Note the location of the summary and log files for future reference.
- 10 On all nodes, restore the configuration files from the backups created in an earlier procedure.

See [“To stop Oracle RAC and the cluster resources”](#) on page 148.

- Restore the original main.cf file:

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

- Merge any changes you made to the original types.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

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

- Merge any changes you made to the original OracleTypes.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

```
/etc/VRTSvcs/conf/config/OracleTypes.cf
```

- 11 Restart each system in the cluster.

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

- 12 Make sure that all eight GAB ports are up after the restart.

```
# gabconfig -a
```

- 13 Make sure that all the resources configured under VCS except the CSSD and Oracle resource come online.

Upgrading from pre-5.0 versions of SF Oracle RAC using the installsfrac installer

If you have a working SF Oracle RAC 4.0 MP4 cluster you can upgrade using the `installsfrac` installer.

To upgrade to SF Oracle RAC 5.0 MP3 using the installsfrac installer

- 1 Mount the software disc:

```
# mkdir /cdrom
# mount -V cdrfs -o ro /dev/cd0 /cdrom
```

- 2 Stop VCS:

```
# hastop -all
```

Wait for VCS to stop.

- 3 Unmount all VxFS filesystems.

Make sure all VxFS filesystems are unmounted.

- 4 Start the SF Oracle RAC installer utility. Include the `-rsh` option if you have not set up SSH:

```
# cd /cdrom/dvd1/sfrac
# ./installsfrac
```

- 5** The installer displays information about the cluster, the systems, and the service groups, and prompts you about upgrading:

...

Do you want to upgrade to version 5.0 MP3 on these systems using the current configuration? [y,n,q,?] (y)

Press Enter.

- 6** The installer verifies RSH accessibility and the existing licenses.

...

Do you want to enter another license key for galaxy? [y,n,q] (n)

- 7** If necessary, answer "y" and enter the license keys at the prompt.
- 8** The installer checks existing AIX operating system filesets and reports any missing recommended filesets, prompting you to add them before continuing.
- 9** Install any required or recommended filesets for SF Oracle RAC and answer "y" at the prompt to continue.
- 10** The installer displays the list of all filesets included with SF Oracle RAC and proceeds to stop SF Oracle RAC processes on each system in the cluster. The installer prompts you to confirm whether to proceed with the upgrade.

All SFRAC processes that are currently running must be stopped.
 Are you sure you want to upgrade SFRAC? [y,n,q] (y)

Answer "y" to continue with the upgrade.

- 11** As the installer proceeds, it:
 - Takes the CVM group and other dependent groups offline.
 - Updates GAB and VCS configurations as needed for SF Oracle RAC.
- 12** When the installer is ready to stop the cluster to make updates to the configuration, it prompts for your OK. At the prompt, press Enter.

...

Are you ready to begin the Veritas Storage Foundation for Oracle RAC upgrade at this time? [y,n,q] (y)

- 13** The installer backs up the configuration files, freezes VCS service groups, and uninstalls the 4.0 SF Oracle RAC filesets.

```
Backing up types.cf and main.cf .....Done
Freezing group cvm .....Done
...
Stopping SFRAC: 100%
Shutdown completed successfully on all systems
Uninstalling SFRAC: 100%
Uninstall completed successfully on all systems
Installing SFRAC: 100%
Installation completed successfully on all systems
```

- 14** Note the location of the summary and log files for future reference.

- 15** Before restarting the system, update the configuration files. On all nodes, restore the configuration files from the backups created in an earlier procedure.

See [“To stop Oracle RAC and the cluster resources”](#) on page 148.

- Merge any changes you made to the original types.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

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

- Merge any changes you made to the original OracleTypes.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

```
/etc/VRTSvcs/conf/config/OracleTypes.cf
```

- 16** Restart each system in the cluster.

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

Performing post-upgrade tasks for SF Oracle RAC 5.0 MP3

You must relink Oracle RAC to the Veritas libraries after upgrading to SF Oracle RAC 5.0 MP3.

To create the new skgxn archive library before relinking Oracle RAC 10g or Oracle RAC 11g

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

- Unload any unused modules:

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

- Create the 32-bit SKGXN archive library under /opt/VRTSvcs/rac/lib:

```
# cd /opt/VRTSvcs/rac/lib
# rm libskgxn2.a # cp libvcsmm.so shr_skgxn2.o
# ar -X32 -qv libskgxn2.a shr_skgxn2.o
# rm shr_skgxn2.o
```

- Create the 64-bit SKGXN archive library under /opt/VRTSvcs/rac/lib64:

```
# cd /opt/VRTSvcs/rac/lib64
# rm libskgxn2.a
# cp libvcsmm.so shr_skgxn2.o
# ar -X64 -qv libskgxn2.a shr_skgxn2.o
# rm shr_skgxn2.o
```

- 2 If you are running Oracle RAC 10g R1, copy these SF Oracle RAC SKGXN libraries to \$CRS_HOME:

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

Run the commands on each system in the cluster.

- 3 If you are running Oracle RAC 10gR2 or Oracle RAC 11gR1, perform the following commands on each system in the cluster.

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

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

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

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

Run the commands on each system in the cluster.

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

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

- 4 For Oracle RAC 10g or Oracle RAC 11gR1, copy the clsinfo script to /opt/ORCLcluster:

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

To relink Oracle RAC 10g or Oracle RAC 11g using the installer

1 Invoke the `installsf` installer:

```
# cd /opt/VRTS/install
# ./installsf -configure
```

2 Enter the system names when prompted:

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

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

■ Select the appropriate Oracle version:

- 1) Oracle 9iR2
- 2) Oracle 10gR1
- 3) Oracle 10gR2
- 4) Oracle 11g

■ Select “Relink Oracle” from the menu:

- 1) Install Oracle Clusterware (CRS)
- 2) Install Oracle RDBMS server
- 3) Relink Oracle
- b) [Go to previous menu]

4 In the “Install or Relink Oracle” menu, enter the required information. For example, if you are using Oracle RAC 10g R2:

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

Enter Oracle Patch Level: [b] (0.1) 0.3
```

- 5 Confirm your responses in the verification screen. The installer copies the SF Oracle RAC libraries to \$ORACLE_HOME/lib.

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 10.2

Oracle Patch Level: 0.3

Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct? [y,n,q] (y)
```

- 6 If any VCS service groups were frozen during the upgrade, unfreeze them.

As root user, enter:

```
# haconf -makerw

# hagrpl -unfreeze group_name -persistent

# haconf -dump -makero
```

- 7 As root user, remove the prefix “:” to uncomment the CRS-specific lines in the /etc/inittab file.

- 8 Start CRS on each node in the cluster.

- If CRS is not controlled by VCS, log in as superuser on each system in the cluster and enter:

```
# /etc/init.crs start
```

- If CRS is controlled by VCS, log in as superuser on any system in the cluster and enter:

```
# hares -online cssd-resource -sys galaxy
# hares -online cssd-resource -sys nebula
```

- 9 Start all resources.

- If the Oracle RAC database instances are under VCS control, bring the corresponding VCS service groups online. As superuser, enter:

```
# hagrpl -online oracle_group -sys galaxy
```



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

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

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

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

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

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

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

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# hagrps -modify cvm_group AutoStart 1
# haconf -dump -makero
```

To relink Oracle RAC 9i

- 1 Invoke `installsfrac`.

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

- 2 Enter the system names when prompted:

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

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

- Select the Oracle RAC 9i version:

```
1) Oracle 9iR2
2) Oracle 10gR1
3) Oracle 10gR2
3) Oracle 11gR1
```

- Select “Relink Oracle” from the menu:

```
1) Install Oracle RDBMS server
2) Relink Oracle
b) [Go to previous menu]
```

- 4 In the “Install or Relink Oracle” menu, enter the required information. For example, if you are using Oracle RAC 9i R2:

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

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

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

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Release: 9.2 Oracle Patch Level: 0.7
Oracle Bits: 64 Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct? [y,n,q] (y)
```

- 6 As Oracle user, enter the command:

```
$ $ORACLE_HOME/bin/psdctl start
```

- 7 If any VCS service groups were frozen during the upgrade, unfreeze them.

As root user, enter:

```
# haconf -makerw

# hagrps -unfreeze group_name -persistent

# haconf -dump -makero
```

- 8 Bring online the Oracle RAC resources configured under VCS.

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

- 9 Verify that the Oracle RAC instances are up and running.

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

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

To commit the updated filesets

- ◆ Use:

```
installp -C filset_name
```

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

Upgrading CVM disk group version

To take advantage of the new features in this release, you must upgrade the Veritas Cluster Volume Manager (CVM) disk group version to 140. The default protocol version is 80 and does not need to be upgraded.

To upgrade the disk group version

- 1 Upgrade the disk group version to 140. Enter the following command on the master node:

```
# vxpdg -T 140 upgrade <disk_group_name>
```

- 2 After upgrading CVM in a VCS environment, you should run the command:

```
# vxvcmconfig upgrade
```

If this command is not run, you will see a warning in the engine log file, `/opt/VRTSvcs/log/engine_A.log`.

To verify the cluster protocol version

- ◆ Enter the following command to verify the cluster protocol version:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

Minimum downtime upgrade from SF Oracle RAC 5.0 to 5.0 MP3

This procedure is valid only for upgrading SF Oracle RAC 5.0 to 5.0 MP3. For upgrades from previous versions of SF Oracle RAC:

See [“Supported upgrade options for SF Oracle RAC 5.0 MP3”](#) on page 145.

Caution: There is a potential for dependency problems between product components that no longer match when upgrading part of a cluster at a time. Following the minimum downtime procedures carefully will enable you to avoid these problems.

Upgrading clusters with Oracle RAC 10g or Oracle RAC 11g R1

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

The following assumptions are made for this procedure:

- SF Oracle RAC 5.0 is installed
- Oracle RAC 10g or Oracle RAC 11gR1 is installed
- Oracle RAC binaries are installed on local file systems for each node in the cluster

The minimum downtime methodology involves upgrading half of the nodes in the cluster at a time.

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

To prepare the first half of the Oracle RAC 10g or Oracle RAC 11g R1 cluster for minimum downtime upgrade

- ◆ On the first half of the cluster nodes, for example *galaxy* and *nebula*, back up the configuration files:

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

To upgrade the first half of an Oracle RAC 10g or Oracle RAC 11g R1 cluster

- 1 Log in to a node in the first half of the cluster, for example, *galaxy*, as root user, set the Oracle group and CSSD resource AutoStart 0.

```
# haconf -makerw
# hagr -modify oracle_resource AutoStart 0
# hares -modify cssd_resource AutoStart 0
# haconf -dump -makero
```

- 2 Offline Oracle and all parallel groups except CVM on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*.

- If the Oracle RAC instance is under VCS control:

```
# hagr -offline oracle_group -sys galaxy
# hagr -offline oracle_group -sys nebula
```

- If the Oracle RAC instance is not under VCS control, log in as Oracle user and shut down all the instances on *galaxy* and *nebula*.

For example:

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

- 3 Stop CRS daemon on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*. If the CRS daemon is under VCS control, offline it.

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

- 4 Switch Failover groups from the first half of the cluster, for example, from *galaxy*, to the second half of the cluster, for example to *jupiter* and *mercury*.

```
# hagrps -switch failover_group -to jupiter/
mercury
```

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

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

# umount /mount_point
```

- 6 Offline CVM group:

```
# hagrps -offline cvm-group -sys galaxy
# hagrps -offline cvm-group -sys nebula
```

- 7 Log in to the nodes of the first half of the cluster, for example, *galaxy* and *nebula*, as root user and stop VCS.

```
# hastop -local
```

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

For example:

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

- 9 On the nodes of the first half of the cluster, for example *galaxy* and *nebula*, invoke `slibclean` to unload all unused modules:

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

- 10 On the nodes of the first half of the cluster, for example *galaxy* and *nebula*, stop VEA if it is running.

```
# /opt/VRTSob/bin/vxsvcctl stop
```

- 11 Change to the CD-ROM directory.

```
# cd /cdrom
```

- 12 Start the installer.

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

- 13 Move the ODM startup script to different name on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*:

```
# mv /etc/rc.d/rc2.d/S99odm /etc/rc.d/rc2.d/nostart.S99odm
```

14 Restart each system in the cluster when prompted.

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

15 When the nodes of the first half of the cluster, for example *galaxy* and *nebula*, come up, make sure these nodes join the second half of the cluster, for example *jupiter* and *mercury*, at GAB/LLT/VXFEN level. CVM/CFS/HAD will not come up due to mismatch versions. Verify the ports a, b, and o all have 0123 membership.

For example, enter on node *galaxy*:

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

Enter on *jupiter*:

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

To prepare for relinking the first half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster

1 Perform the following setup steps on the nodes of the first half of the cluster, for example *galaxy* and *nebula*, before bringing up CRS:

- Unload any unused modules:


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

- Create the 32-bit SKGXN archive library under /opt/VRTSvcs/rac/lib:

```
# cd /opt/VRTSvcs/rac/lib
# rm libskgx2.a # cp libvcsmm.so shr_skgxn2.o
# ar -X32 -qv libskgx2.a shr_skgxn2.o
# rm shr_skgxn2.o
```

- Create the 64-bit SKGXN archive library under /opt/VRTSvcs/rac/lib64

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

- 2 If you are running Oracle RAC 10gR1, copy these SF Oracle RAC SKGXN libraries to \$CRS_HOME on the nodes of the first half of the cluster, for example *galaxy* and *nebula*.

```
# cp /opt/VRTSvcs/rac/lib/libskgx2.a \
$CRS_HOME/lib32/libskgx2.a
# cp /opt/VRTSvcs/rac/lib/libskgx2.a \
$CRS_HOME/lib32/libskgxnr.a
# cp /opt/VRTSvcs/rac/lib64/libskgx2.a \
$CRS_HOME/lib/libskgx2.a
# cp /opt/VRTSvcs/rac/lib64/libskgx2.a \
$CRS_HOME/lib/libskgxnr.a
```

- 3 If you are running Oracle RAC 10gR2 or Oracle RAC 11gR1, perform the following commands on the nodes of the first half of the cluster, for example *galaxy* and *nebula*.

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

```
# rm /opt/ORCLcluster/lib/libskgx2.a
# rm /opt/ORCLcluster/lib/libskgx2.so
# rm /opt/ORCLcluster/lib32/libskgx2.a
# rm /opt/ORCLcluster/lib32/libskgx2.so
```

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

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

Run the commands on each system in the cluster.

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

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

- 4 For Oracle RAC 10g or Oracle RAC 11gR1, copy the `clsinfo` script to /opt/ORCLcluster:

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

To relink Oracle RAC 10g or Oracle RAC 11gR1 libraries

- 1 Invoke `installsfrac`:

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

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

- Select the appropriate Oracle RAC version:

```
1) Oracle 9iR2
2) Oracle 10gR1
3) Oracle 10gR2
4) Oracle 11gR1
```

- Select "Relink Oracle" from the menu.

- 1) Install Oracle Clusterware (CRS)
- 2) Install Oracle RDBMS server
- 3) Relink Oracle
- b) [Go to previous menu]

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

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

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

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

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 10.2

Oracle Patch Level: 0.1

Oracle Base: /app/oracle

Switch Applications to 5.0 MP3 cluster
```

When you relink the Oracle libraries for the first half of your cluster you can:

- Switch applications to the upgraded nodes in the first half of the cluster
- Upgrade the second half of your cluster

To switch applications to the upgraded half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster

- 1 On the second half of the cluster nodes, for example *jupiter* and *saturn*, back up the configuration files:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /kernel/drv/vcsmm.conf /kernel/drv/vcsmm.conf.save
```

- 2 Offline Oracle RAC and all parallel groups on *jupiter* and *mercury* except CVM.

- If the Oracle RAC instance is under VCS control:

```
# hagrpl -offline oracle_resource -sys jupiter
# hagrpl -offline oracle_resource -sys mercury
```

- If the Oracle RAC instance is not under VCS control, log in as Oracle user and shut down all the instances on the nodes of the second half of the cluster. For example:

```
$ srvctl stop instance -d database_name -i instance_name \
-n jupiter
$ srvctl stop instance -d database_name -i instance_name \
-n mercury
$ srvctl stop nodeapps -n jupiter
$ srvctl stop nodeapps -n mercury
```

Note: Your downtime begins now.

- 3 Stop the CRS daemon on the nodes of the second half of the cluster, for example *jupiter* and *mercury*. If the CRS daemon is under VCS control, offline it.

```
# hares -offline cssd_resource -sys jupiter
# hares -offline cssd_resource -sys mercury
```

- 4 Offline Failover groups on *jupiter* and *mercury*.

```
# hagrpl -offline failover_group -sys jupiter
# hagrpl -offline failover_group -sys mercury
```

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

```
# mount |grep vxfs

# fuser -c /mount_point

# umount /mount_point
```

- 6 Offline CVM group:

```
# hagr -offline cvm-group -sys galaxy
# hagr -offline cvm-group -sys nebula
```

- 7 Log in on a node in the second half of the cluster, for example *jupiter*, as root user, and stop VCS:

```
# hstop -all
```

- 8 Stop ODM on both *jupiter* and *mercury*:

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

- 9 Restore ODM startup scripts on the first half of the cluster, for example *galaxy* and *nebula*:

```
# mv /etc/rc.d/rc2.d/nostart.S99odm /etc/rc.d/rc2.d/S99odm
```

- 10 Restart ODM on *galaxy* and *nebula*:

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

- 11 On the first half of the cluster, for example on *galaxy* and *nebula*, as root user, remove the prefix ":" to uncomment the CRS-specific lines in the */etc/inittab* file.

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

- 12 On the first half of the cluster, restore the configuration files from the backups created in an earlier procedure.

See [“To prepare the first half of the Oracle RAC 10g or Oracle RAC 11g R1 cluster for minimum downtime upgrade”](#) on page 165.

- Restore the original main.cf file:

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

- Merge any changes you made to the original types.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

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

- Merge any changes you made to the original OracleTypes.cf file and backed up in a previous procedure at:

```
/etc/VRTSvcs/conf/config/OracleType.se.cf.save
```

with the new types.cf file at:

```
/etc/VRTSvcs/conf/config/OracleTypes.cf
```

13 Start VCS on *galaxy* and *nebula*.

```
# hstart
```

14 Start the CRS daemon on nodes in the first half of the cluster, for example *galaxy* and *nebula*:

- If CRS is not configured under VCS control, log in as superuser on *galaxy* and *nebula* and enter:

```
# /etc/init.crs start
```

- If CRS is configured under VCS control, log in as superuser and enter the following command for each system in the cluster:

```
# hares -online cssd_resource -sys galaxy  
# hares -online cssd_resource -sys nebula
```

15 Bring Oracle RAC online on *galaxy* and *nebula*.

- If the Oracle RAC instance is under VCS control:

```
# hagr -online oracle_resource -sys galaxy  
# hagr -online oracle_resource -sys nebula
```

- If the Oracle RAC instance is not under VCS control, log in as Oracle user and start all the instances on *galaxy* and *nebula*. For example:

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

- 16 Bring online all the applications on the second half of the cluster, for example on *galaxy* and *nebula* which were offline before upgrade.

Note: Your downtime ends now. You have Oracle RAC instances running on the first half of your cluster, for example one Oracle RAC instance running on *galaxy* and one instance running on node *nebula*.

To upgrade the second half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster

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

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

- 3 Log in to the nodes on the second half of the cluster, for example *jupiter* and *mercury*, invoke `slibclean` to unload all unused modules:

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

- 4 Stop VEA if it is running:

```
# /opt/VRTSob/bin/vxsvcctl stop
```

5 Upgrade SF Oracle RAC on *jupiter* and *mercury*.

```
# cd /cdrom
For ssh
# ./installmp jupiter
               mercury

For rsh:

# ./installmp -rsh jupiter mercury
```

6 Restart each system in the cluster when prompted.

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

7 When *jupiter* and *mercury* come back online, they will join *galaxy* and *nebula* to form a four-node SF Oracle RAC 5.0 MP3 cluster.

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0d membership 0123
Port b gen 5c3d13 membership 0123
Port d gen 5c3d0e membership 0123
Port f gen 5c3d33 membership 0123
Port h gen 5c3d2d membership 0123
Port o gen 5c3d12 membership 0123
Port v gen 5c3d2f membership 0123
Port w gen 5c3d31 membership 0123
```

8 Prepare for relinking after Oracle RAC 10g or Oracle RAC 11gR1 upgrade.

See [“To prepare for relinking the first half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster”](#) on page 168.

9 Relink the Oracle RAC libraries on the second half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster:

See [“To relink Oracle RAC 10g or Oracle RAC 11gR1 libraries”](#) on page 170.

To bring up applications on the second half of the Oracle RAC 10g or Oracle RAC 11gR1 cluster

- 1 On the nodes of the second half of the cluster, for example, *jupiter* and *mercury*, as root user, remove the prefix ":" to uncomment the CRS-specific lines in the `/etc/inittab` file.

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

- 2 Start the CRS daemon on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*:

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

```
# /etc/init.crs start
```

- If CRS is configured under VCS control, log in as superuser and enter the following command for each system in the cluster:

```
# hares -online cssd_resource -sys jupiter
# hares -online cssd_resource -sys mercury
```

- 3 Bring Oracle RAC online on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*.

- If the Oracle RAC instance is under VCS control:

```
# hagrps -online oracle_resource \
-sys jupiter
# hagrps -online oracle_resource \
-sys mercury
```

- If the Oracle RAC instance is not under VCS control, Login as Oracle user and start all the instances on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*. For example:

```
$ srvctl start instance -d database_name -i instance_name \
-n jupiter
$ srvctl start instance -d database_name -i instance_name \
-n mercury
```

- 4 Bring up all the applications which were offline before the upgrade on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*.
- 5 Change the Oracle group and CSSD AutoStart attribute to 1:

```
# haconf -makerw
# hares -modify oracle_resource AutoStart 1
# hares -modify cssd_resource AutoStart 1
# haconf -dump -makero
```

To commit the updated filesets

- ◆ Use:

```
installp -C filset_name
```

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

If you are using the repository database, you must upgrade it after upgrading SF Oracle RAC.

See [“Upgrading the repository database”](#) on page 189.

Upgrading clusters with Oracle RAC 9iR2

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

The following assumptions are made for this procedure:

- SF Oracle RAC 5.0 is installed
- Oracle RAC 9iR2 is installed
- Oracle RAC binaries are installed on separate nodes for each node in the cluster

The minimum downtime methodology involves upgrading half of the nodes in the cluster at a time.

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

To prepare the first half of the Oracle RAC 9iR2 cluster for minimum downtime upgrade

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

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

To upgrade the first half of the Oracle RAC 9iR2 cluster

- 1 Set the Oracle group to AutoStart 0.

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

- 2 Offline all parallel groups on the nodes of the first half of the cluster, for example, *galaxy* and *nebula* except CVM. If the Oracle RAC instance is under VCS control, offline it.

```
# hagr -offline group_name -sys galaxy
# hagr -offline group_name -sys nebula
```

- 3 Stop the GSD daemon:

```
$ gsdctl stop
```

- 4 Switch Failover groups from the nodes of the first half of the cluster, for example, *galaxy* to the other nodes, for example, *jupiter* and *mercury*.

```
# hagr -switch failover_group -to jupiter/
mercury
```

- 5 Unmount all the vxfs file systems which are not under VCS control.

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

# fuser -m /mount_point

# umount /mount_point
```

6 Offline CVM group:

```
# hagrps -offline cvm_group -sys galaxy  
  
# hagrps -offline cvm_group -sys nebula
```

7 Stop VCS on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*

```
# hastop -local
```

8 Stop VEA if it is running:

```
# /opt/VRTSob/bin/vxsvcctl status  
  
# /opt/VRTSob/bin/vxsvcctl stop
```

9 Change directories to the CD-ROM directory, start the installer:

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

10 Move the ODM startup script to different name on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*:

```
# mv /etc/rc.d/rc2.d/S99odm /etc/rc.d/rc2.d/nostart.S99odm
```

11 Restart each system in the cluster when prompted.

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

12 Make sure the nodes of the first half of the cluster, for example, *galaxy* and *nebula*, join the nodes of the second half of the cluster, for example, *jupiter* and *mercury*, at GAB/LLT/VXFEN level, however CVM/CFS/HAD will not come up due to mismatch versions. Verify the ports a, b, o all have 0123 membership.

For example: Enter on node *galaxy*:

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

Enter on node *jupiter*

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0b membership 0123
Port b gen 5c3d10 membership 0123
Port d gen 5c3d0c membership ; 2
Port d gen 5c3d0c visible 01
Port f gen 5c3d15 membership ; 23
Port f gen 5c3d15 visible 01
Port h gen 5c3d2a membership ; 23
Port h gen 5c3d2a visible 01
Port o gen 5c3d0f membership 0123
Port v gen 5c3d12 membership ; 23
Port v gen 5c3d12 visible 01
Port w gen 5c3d13 membership ; 23
Port w gen 5c3d13 visible 01
```

To relink Oracle RAC 9iR2 libraries

1 Invoke `installsfrac`:

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

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

■ Select the Oracle RAC 9iR2 version:

- 1) Oracle 9iR2
- 2) Oracle 10gR1
- 3) Oracle 10gR2
- 4) Oracle 11gR1

■ Select “Relink Oracle” from the menu:

- 1) Install Oracle RDBMS server
- 2) Relink Oracle
- b) [Go to previous menu]

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

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

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

- 4 Confirm your responses in the verification screen. The installer copies the SF Oracle RAC 5.0 MP3 Oracle RAC libraries to /opt/ORCLcluster, where it expects libskgxn.

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Release: 9.2

Oracle Patch Level: 0.7

Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct? [y,n,q] (y)
```

To switch applications to the upgraded half of the Oracle RAC 9i cluster

- 1 Offline all Parallel groups on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*, except CVM. If the Oracle RAC instance is under VCS control, offline it.

Note: Your downtime begins now.

```
# hagrps -offline group_name -sys jupiter
# hagrps -offline group_name -sys mercury
```

- 2 Stop the GSD daemon.

```
$ gsdctl stop
```

- 3 Offline Failover groups on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*.

```
# hagrps -offline failover_group -sys jupiter  
# hagrps -offline failover_group -sys mercury
```

- 4 Unmount all the vxfs file systems which are not under VCS control.

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

- 5 Offline the CVM group.

```
# hagrps -offline cvm_group -sys galaxy  
# hagrps -offline cvm_group -sys nebula
```

- 6 Stop VCS on the nodes of the second half of the cluster, for example, *jupiter* and *mercury*.

```
# hastop -all
```

- 7 Stop ODM on both *jupiter* and *mercury*:

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

- 8 Restore ODM startup scripts on the first half of the cluster, for example *galaxy* and *nebula*:

```
# mv /etc/rc.d/rc2.d/nostart.S99odm /etc/rc.d/rc2.d/S99odm
```

- 9 Restart ODM on *galaxy* and *nebula*:

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

- 10 On the first half of the cluster, restore the configuration files from the backups created in an earlier procedure.

See [“To prepare the first half of the Oracle RAC 9iR2 cluster for minimum downtime upgrade”](#) on page 179.

- Restore the original main.cf file:


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

- Merge any changes you made to the original types.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

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

- Merge any changes you made to the original OracleTypes.cf file and backed up in a previous procedure at:

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

with the new types.cf file at:

```
/etc/VRTSvcs/conf/config/OracleTypes.cf
```

- 11 Start VCS on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*.

```
# hastart
```

- 12 Start GSD. Bring Oracle RAC online on the nodes of the second half of the cluster, for example, *galaxy* and *nebula*.

```
$ gsdctl start
```

```
# hagrpl -online oracle_group -sys galaxy/nebula
```

Note: Your downtime ends now. You have Oracle RAC instances running on the nodes of the first half of the cluster, for example, *galaxy* and *nebula*.

To upgrade the second half of the Oracle RAC 9iR2 cluster

- 1 Log in to the nodes of the second half of the cluster, for example, *jupiter* and *mercury*.
- 2 Stop VEA if it is running:

```
# /opt/VRTSob/bin/vxsvcctrl status
```

```
# /opt/VRTSob/bin/vxsvcctrl stop
```

- 3 Change directories to the CD-ROM directory and start the installer:

```
# cd /cdrom # ./installmp [-rsh] jupiter mercury
```

- 4 Restart each system in the cluster when prompted.

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

When the nodes of the second half of the cluster, for example, *jupiter* and *mercury* come back online, they will join the nodes of the first half of the cluster, for example, *galaxy* and *nebula* to form a four-node SF Oracle RAC 5.0 MP3 cluster.

- 5 Relink the Oracle RAC libraries on the second half of the Oracle RAC 9iR2 cluster:

See “[To relink Oracle RAC 9iR2 libraries](#)” on page 182.

- 6 Bring online the Oracle RAC resources configured under VCS.

```
# hares -online listener_res -sys galaxy
```

```
# hares -online listener_res -sys nebula
```

```
# hagrps -online group_name -sys galaxy
```

```
# hagrps -online group_name -sys nebula
```

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

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

To commit the updated filesets

- ◆ Use:

```
installmp -C filset_name
```

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

If you are using the repository database, you must upgrade it after upgrading SF Oracle RAC.

See [“Upgrading the repository database”](#) on page 189.

Upgrading licenses to SF Oracle RAC level

You can upgrade from other Storage Foundation 5.0 MP3 products to SF Oracle RAC 5.0. You can upgrade the following SF products to SF Oracle RAC:

- Storage Foundation for Oracle
- Storage Foundation for Oracle High Availability
- Storage Foundation Cluster File System

Review the following procedures that outline how to perform these upgrades.

Upgrading to SF Oracle RAC from Storage Foundation for Oracle

If you have Storage Foundation for Oracle installed, use the following procedure to add the license for SF Oracle RAC and implement it.

To upgrade from Storage Foundation for Oracle to SF Oracle RAC

- 1 Install the license for SF Oracle RAC.

```
# installsfrac -license
```

- 2 Shutdown and restart all nodes.

```
# shutdown -r
```

- 3 Use the installer to configure SF Oracle RAC.

See [“Configuring SF Oracle RAC components”](#) on page 95.

Upgrading to SF Oracle RAC from Storage Foundation HA for Oracle

If you have Storage Foundation for Oracle HA installed, use the following procedure to add the license for SF Oracle RAC and implement it.

To upgrade from Storage Foundation HA for Oracle to SF Oracle RAC

- 1 Install the license for SF Oracle RAC.

```
# installsfrac -license
```

- 2 Configure the vcsmm driver.

```
# echo `sbin/vcsmmconfig -c &' > /etc/vcsmmtab
```

Enter this command on all systems in the cluster.

- 3 Shutdown and restart all nodes.

```
# shutdown -r
```

- 4 Configure I/O fencing.

See [“Setting up I/O fencing”](#) on page 125.

- 5 Enter the command to configure CFS.

```
# /opt/VRTSvxfs/cfs/bin/cfscluster config
```

Upgrading to SF Oracle RAC from Storage Foundation Cluster File System

To upgrade from SF CFS to SF Oracle RAC, perform the following steps.

To upgrade from Storage Foundation Cluster File System to SF Oracle RAC

- 1 Install the license for SF Oracle RAC.

```
# installsfrac -license
```

- 2 Create /etc/vcsmmtab (Veritas Cluster Server Membership Manager tab) on all the nodes in the cluster using the following command:

```
/sbin/vcsmmconfig -c &
```

This tab will now automatically configure VCSMM every time the system is rebooted.

3 Shutdown and restart all nodes.

```
# shutdown -r
```

The following steps in this procedure are required to set up a database that is to be managed by SF Oracle RAC. For additional information about this process, please refer to *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

4 Configure SF Oracle RAC per your requirements.

5 Change the permissions, modes, and ownership of the following files using the following commands:

```
# chmod 750 /opt/VRTSdbed  
# chown oracle:dba /opt/VRTSdbed
```

6 Create the repository database.

Refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for information about setting up the repository database using the Veritas sfua_db_config script.

Upgrading the repository database

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

- Change permissions
- Upgrade the repository database

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

To change permissions

- 1 Change permissions for the following directory:

```
# chmod 750 /opt/VRTSdbed
```

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

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

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

To upgrade the repository in a cluster environment

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

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

This command:

- Removes the repository information from the VCS configuration
- Unmounts the repository
- Deports the disk group
- Takes the volumes offline

- 3 Reconfigure the repository:

```
# vxdg import sfua_rep
# vxvol -g dbed_rep startall
```

- 4 Mount the repository file system manually.

```
# mount -V vxfs /dev/vx/dsk/sfua_rep/dbed_rep /sfua_rep
```

- 5 Run the repository upgrade command again without any options:

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

The following is required to configure the repository for SF Oracle RAC:

- A mount point of already mounted Veritas Volume on a shared storage, with Veritas File system
- A public NIC used by each system in the cluster

- A Virtual IP address and netmask

Oracle RAC 10g installation in an SF Oracle RAC environment

- [Preparing to install Oracle RAC 10g](#)
- [Installing Oracle RAC 10g](#)
- [Configuring VCS service groups for Oracle RAC 10g](#)
- [Upgrading to Oracle 10g and migrating the database](#)

Preparing to install Oracle RAC 10g

This chapter includes the following topics:

- [About preparing to install Oracle RAC 10g](#)
- [Pre-installation tasks](#)


About preparing to install Oracle RAC 10g

After setting up Veritas Storage Foundation for Oracle RAC 10g, prepare to install Oracle RAC 10g software. You can install the Oracle RAC 10g software on shared storage or locally on each node.

Note: Review your Oracle installation manuals and the appropriate Oracle support websites, before installing Oracle RAC 10g.

[Figure 9-1](#) displays the required steps to prepare for an Oracle RAC 10g installation. Each step is discussed in detail in this chapter.

Figure 9-1 Preparing to install Oracle RAC 10g flowchart



Pre-installation tasks

Performing pre-installation operations involves both manual and automated tasks from the SF Oracle RAC configuration program. Before installing Oracle RAC 10g you must perform the following tasks:

- Setting the system parameters
- Identifying storage for Oracle RAC 10g components
- Launch the SF Oracle RAC configuration program
- Creating Oracle user and groups
- Setting Oracle Clusterware oracle user capabilities
- Configuring shell limits for Oracle user
- Configuring kernel parameter for processes per user
- Tuning the Virtual Memory Manager (VMM)
- Verifying the user "nobody" exists
- Configuring private IP addresses for CRS
- Identifying the public virtual IP addresses for use by Oracle
- Creating Oracle disk groups, volumes, and mount points
 - Preparing \$CRS_HOME and \$ORACLE_HOME on each node
 - Creating OCR and VOTE-disk volumes
- Verifying SF Oracle RAC resources on all nodes
- Verifying GAB port membership

Note: When running the SF Oracle RAC configuration program, you can perform all preinstallation tasks or choose specific tasks to perform. You can also perform preinstallation tasks sequentially at one time, or you can access the program to perform individual tasks at a later time.

Note: Several preinstallation tasks must be manually performed.

Setting the system parameters

Set the system parameters, including the OS kernel and networking parameters.

Refer to the latest Oracle documentation for information about the system parameters and procedures. Restart the nodes for the new values to take effect.

Identifying storage for Oracle RAC 10g components

The SF Oracle RAC configuration program provides options to perform the following tasks:

- Oracle RAC 10g preinstallation tasks
- Oracle RAC 10g CRS installation
- Oracle RAC 10g binaries installation
- Link Veritas libraries with Oracle

The configuration program prompts you for information while installing Oracle RAC 10g in an SF Oracle RAC environment. The program provides default values for some information, which you can choose to use.

Note: The examples in this chapter assume a two-node cluster, with the nodes galaxy and nebula.

The following information is required for the configuration:

- Information to create Oracle user and Inventory group as primary group:
See [“Creating Oracle user and groups”](#) on page 201.
- Information to configure private IP addresses for CRS:
See [“Configuring private IP addresses for CRS”](#) on page 210.
- Information for \$CRS_HOME on each system:
See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 218.
- Information to create OCR and Voting Disk directories on CFS:
See [“Creating OCR and VOTE-disk volumes”](#) on page 221.

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

Launch the SF Oracle RAC configuration program

The following procedure describes the Oracle RAC 10g pre-installation tasks using the SF Oracle RAC configuration program and menu.

To perform pre-installation tasks in the configuration menu

- 1 If the SF Oracle RAC configuration program is not currently running, launch the SF Oracle RAC configuration program.

As root user on any one system, enter:

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

By default, the installsfrac program uses SSH for remote communication.

To use RSH, specify the -rsh option with the installsfrac program.

- 2 If you have not previously installed SF Oracle RAC, then enter the system names, separated by spaces.
- 3 The SF Oracle RAC configuration program displays the Symantec copyright information, as well as the location of the installsfrac logs

Note: If necessary, the installsfrac logs can be used for troubleshooting purposes.

- 4 If you have previously installed SF Oracle RAC, then information from the SFRAC configuration files currently on the system appears. This information will include the cluster name, cluster ID, systems for the cluster, and cluster service groups.

If you have not previously installed SF Oracle RAC, then you will be prompted to enter this information.

You will be prompted whether you wish to configure SFRAC from these systems.

```
Do you want to configure SFRAC from these systems? [y,n,q] (y)
```

- 5 The SF Oracle RAC configuration program then performs an initial system check of the ssh or rsh communication, SF Oracle RAC installation, and architecture on the nodes.
- 6 The SF Oracle RAC configuration program then checks your system license on the nodes. You will be asked whether you want to enter another license key on each system.

If not, then enter "n" to proceed.

7 The SF Oracle RAC configuration program then displays the following information about its INSTRUCTIONS ONLY steps:

- * Do not perform any action
- * Require you to understand the instructions and act
- * May require you to refer to the documentation

Press enter to proceed.

The following additional information is provided:

- * Keep root shells open on all systems to run INSTRUCTIONS ONLY steps
- * Follow the steps in order
- * Navigation keys for all screens: q to quit, b for previous menu

8 The configuration program menu is then displayed:

- 1) Check systems for SFRAC ** INSTRUCTIONS ONLY **
- 2) Configure SFRAC
- 3) Prepare to install Oracle
- 4) Install or relink Oracle

From this menu, you can choose to perform all installation tasks, or to perform specific tasks.

From the configuration program menu, select **Prepare to install Oracle**

9 After selecting the **Prepare to install Oracle** option, the following menu is presented:

This step helps you complete the sysadmin tasks before installing Oracle software.

- 1) [Perform all the following tasks]
 - 2) Create userid and groupid for Oracle
 - 3) Configure private IP addresses for CRS (Oracle 10g only)
 - 4) Configure CRS and UDP private address for failover (Oracle 11g)
 - 5) Create disk groups, volumes, and files systems for Oracle **INSTRUCTIONS ONLY**
- b) [Go to previous menu]

Proceed to select the appropriate options for your installation and follow the steps to pre-install Oracle RAC 10g.

For detailed information about option 2 (Create userid and groupid for Oracle RAC 10g): See [“Creating Oracle user and groups”](#) on page 201.

For detailed information about option 3 (Configure private IP addresses for CRS, Oracle RAC 10g only): See [“Configuring private IP addresses for CRS”](#) on page 210.

For detailed information about option 5 (Create disk groups, volumes, and file systems for Oracle RAC 10g): See [“Creating Oracle disk groups, volumes, and mount points”](#) on page 216.

Creating Oracle user and groups

You must create the groups oinstall (the Oracle Inventory group) and dba, and the user oracle, assigning the primary group for oracle to be oinstall and the secondary group for oracle to be dba.

Refer to the latest Oracle documentation for additional information.

Table 9-1 Required information to create Oracle user and Inventory group as primary group

Required information	Example
Oracle user name	oracle Default is oracle.
Oracle user id	100
Oracle user home directory	/home/oracle
Oracle group name	oinstall Default is oinstall.
Oracle group id	1000

Table 9-2 Required information to create OS DBA group as secondary group

Required information	Example
Oracle group name	dba Default is dba.
Oracle group id	1001

You can create Oracle users and groups using the SF Oracle RAC configuration program. As discussed above, refer to the Oracle documentation for information about creating the oinstall (Oracle Inventory), dba group, and the oracle user. After reviewing the Oracle documentation, proceed to create these Oracle groups and user.

Note: Make sure that the user and group IDs are unused on all the nodes in the cluster.

To create Oracle user and groups

- 1 Access the SF Oracle RAC configuration program if you are not currently using it.

See [“To perform pre-installation tasks in the configuration menu”](#) on page 199.

- 2 Select **Create userid and groupid for Oracle** from the configuration menu.

The path to this configuration menu option is:

```
Main Menu > Prepare to install Oracle >  
Create userid and groupid for Oracle
```

- 3 Provide the required information.

The configuration program assigns the same values on all nodes. The text within parentheses is the default, which you may select by pressing Enter.

```
Enter Oracle UNIX user name: (oracle)  
Enter Oracle UNIX user id: 100  
Enter Oracle UNIX user home dir: /home/oracle  
Enter Oracle UNIX group name: (oinstall)  
Enter Oracle UNIX group id: 1000  
User oracle does not exist on any node. Do you want to create it  
with the information provided [y,n,q] (y)
```

The set of Oracle user IDs and group IDs in each cluster configuration must be the same. Furthermore, if you are configuring GCO then the user IDs and group IDs of all nodes of both the primary and secondary clusters must be the same.

- 4 The installer verifies that the specified `userid` does not exist on any of the systems in the cluster and then creates it. Enter `y` to create the oracle user with the information provided.

- 5 Enter the information to create secondary groups, "dba" and "oper".

The SF Oracle RAC configuration program displays the following message:

```
Do you want to create another secondary group for Oracle user? [y,n,q] (y)

Enter Oracle UNIX secondary group name: [b] dba
Enter Oracle UNIX secondary group id (numerical): [b] 1001

Group dba does not exist on any node. Do you want to create it
with the information provided [y,n,q] (y) y
    Creating group dba on galaxy ..... Done
    Adding Oracle user (oracle) to group (dba) on galaxy .. Done
    Creating group dba on nebula ..... Done
    Adding Oracle user (oracle) to group (dba) on nebula .. Done
```

- 6 Next, you are prompted to perform SSH or RSH verification. Enter "y" to perform verification. Review as the configuration program verifies access for Oracle user on all the nodes. The Oracle installation process requires SSH or RSH permission to be set for the Oracle user. If the SSH or RSH verification fails on any nodes, enable SSH or RSH access for those nodes. To enable rsh access for the nodes, leave the installer at this prompt and proceed to set up access. Open a new session to set up access. You can return to this session after setting up oracle user equivalence. See [“Setting up Oracle user equivalence for RSH and RCP”](#) on page 206.

7 Edit the .profile file to enable paths to \$CRS_HOME, \$ORACLE_HOME, and \$ORACLE_BASE on each node.

On each node, set the proper environment variables (the examples are for ksh):

```
$ cat .profile
set -o vi

PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:\
/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]          # This is at Shell startup.  In normal
then echo "$MAILMSG"      # operation, the Shell checks
fi                        # periodically.

ORACLE_HOME=/app/oracle/orahome; export ORACLE_HOME
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data; export ORA_NLS33
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib; export CLASSPATH
DISPLAY=YOUR.DISPLAY.IP; export DISPLAY
CRS_HOME=/app/crshome; export CRS_HOME
ORA_CRS_HOME=/app/crshome; export ORA_CRS_HOME

LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$CRS_HOME/lib:\
$CRS_HOME/lib32:/usr/lib:/opt/VRTSdbms3/lib; export LIBPATH
PATH=$PATH:$ORACLE_HOME/bin:$CRS_HOME/bin:/opt/VRTS/bin:\
/opt/VRTSdbed/bin; export PATH

ORACLE_SID=PROD; export ORACLE_SID

umask 022

set -o vi
$
```

8 After editing the .profile file as oracle user, enter the following where /usr/oracle is user home:

```
$ . /usr/oracle/.profile
```

Setting up Oracle user equivalence for RSH and RCP

To set up Oracle user equivalence for RSH and RCP

- 1 Access the SF Oracle RAC configuration program (even if you are currently using it) and open another terminal session as `oracle` user.
- 2 As `oracle` user on each system, edit `/etc/hosts.equiv` file and add entries similar to the following:

```
galaxy oracle
nebula oracle
```

- 3 On each system, set the password for the "oracle" user:

```
[root@galaxy /]# passwd oracle
Changing password for "oracle"
oracle's New password:
Re-enter oracle's new password:
```

- 4 On each system, login as user "oracle" and change the password.

```
[root@galaxy /]# su - oracle
$ passwd
Changing password for "oracle"
oracle's New password:
Re-enter oracle's new password:
```

- 5 On each system, as user "oracle", verify "RSH" access:

```
$ rsh galaxy date
Mon Apr 24 10:02:45 PDT 2006
$ rsh nebula date
Mon Apr 24 10:02:45 PDT 2006
```

You can now create the secondary groups for Oracle.

See [“Creating Oracle user and groups”](#) on page 201.

To verify RSH access for "oracle" user

- 1 Return to the installer session.
See [“Creating Oracle user and groups”](#) on page 201.
- 2 At the installer prompt, answer "y" to verify "RSH" accessibility.
- 3 Quit the installation program.

Setting Oracle Clusterware oracle user capabilities

For an Oracle RAC 10g Release 2 installation (beginning with version 10.2.0.4), ensure that the Oracle Clusterware oracle user has the following capabilities:

CAP_NUMA_ATTACH, CAP_BYPASS_RAC_VMM, and CAP_PROPAGATE.

Setting Oracle Clusterware oracle user capabilities

- 1 To check existing capabilities, enter the following command as `root`; in this example, the Oracle Clusterware oracle user is `oracle`:

```
# /usr/sbin/luser -a capabilities oracle
```

- 2 To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser \  
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE oracle
```

- 3 The `ncargs` attribute controls the maximum command line length, including the environment variables on the AIX nodes.

Change this attribute to a larger value before invoking the Oracle Installer.

The following is an example of the command to set this attribute:

```
# /usr/sbin/chdev -l sys0 -a ncargs=128
```

Configuring shell limits for Oracle user

Oracle recommends setting shell limits. Refer to your Oracle manual and the Oracle documentation website for the latest information about setting shell limits before proceeding with the following procedure.

[Table 9-3](#) lists the values for the shell soft limits that you must verify whether these match.

Note: The parameter and shell limit values shown in this section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

Table 9-3 Shell limits for Oracle user

Shell Limit	Recommended Value
Soft FILE size	-1

Table 9-3 Shell limits for Oracle user *(continued)*

Shell Limit	Recommended Value
Soft CPU time	-1
Soft DATA segment	-1
Soft STACK size	-1

If the values do not match those of the table, change them.

To configure shell limits

- 1 Use SMIT to verify the current value. For example:

smitty chuser
- 2 In the field, enter the user name: oracle
- 3 Scroll down the list. Verify the current values against the table above, editing them if necessary.
- 4 Press F10 to exit when you are done.
- 5 Repeat the steps above and verify that the values are the same on all cluster nodes.

Configuring kernel parameter for processes per user

Review the latest Oracle documentation for information about all OS kernel parameters that need to be configured. The following procedure is used to configure the maximum processes per user.

Oracle RAC 10g requires that the AIX configurable kernel parameter that specifies the maximum number of processes per user is greater than 2048. The value must be the same on each of the cluster nodes. Furthermore, if you are configuring GCO make sure this value is the same on all nodes in both primary and secondary clusters.

To configure maximum processes per user

- 1 Check the current value by entering the command:

```
# smitty chgsys
```

- 2 Examine the value of the kernel parameter:

```
Maximum number of PROCESSES allowed per user
```

If necessary, change the value to be 2048 or greater.

- 3 Press F10 to exit when you are done.
- 4 Use the previous steps to check that the value is the same on each cluster node.

Tuning the Virtual Memory Manager (VMM)

Perform the following tasks to tune the Virtual Memory Manager (VMM).

Tuning the Virtual Memory Manager (VMM)

- 1 Use the AIX System Management Interface Tool (SMIT) to access the following selections:
 - smitty > Performance & Resource Scheduling> Tuning Kernel & Network
 - smitty> Parameters> Tuning Virtual Memory Manager, File System and Logical Volume Manager Params>Change/Show Parameters for Next Boot
- 2 Identify each of the following parameters and set it with the value listed below:

```
lru_file_repage=0
minperm%=3
maxperm%=90
maxclient%=90
strict_maxclient=1
strict_maxperm=0
```

Verifying the user "nobody" exists

Verify the user "nobody" exists on each system in the cluster:

```
# id nobody
uid=4294967294 (nobody) gid=4294967294 (nobody)
```

Note: Make sure that the uid and gid are the same across the nodes in your cluster.

Configuring private IP addresses for CRS

The CRS daemon requires a private IP address on each node to enable communications and heartbeating.

Determine a private NIC device for which LLT is configured. Look at the file /etc/llttab. For example, if a device is used as an LLT interconnect on one system, you can configure an available IP address for it.

Table 9-4 Required information to configure private IP addresses for CRS

Required information	Example
Private IP address for each node	■ On galaxy: 192.168.12.1 ■ On nebula: 192.168.12.2 The PrivNIC resource does not accept leading zeros in the private IP addresses. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula. Private IP address formats similar to the following are not acceptable and will fail: ■ 192.168.12.01 ■ 192.168.12.02

Table 9-4 Required information to configure private IP addresses for CRS
(continued)

Required information	Example
Private network interfaces for each node	You have to choose an LLT device as a device for the CRS heartbeat. For example, if the LLT devices are en1, en2: ■ On galaxy: en1, en2 ■ On nebula: en1, en2 Then the PrivNIC device names will be as follows: Device@galaxy= { en1 = 0, en2 = 1 } Device@nebula= {en1 = 0, en2 = 1 } If aggregated device names are configured under LLT, then the aggregated names must be used in PrivNIC agent. For example, if LLT device names are: ■ On galaxy: aggr1 ■ On nebula: aggr1 Then the Device Attribute for the PrivNIC agent would be as follows: Device@galaxy = { aggr1 = 0 } Device@nebula = { aggr1 = 0 }
Netmask for the cluster	255.255.255.0

After confirming the values, the configuration program adds a new section in the VCS configuration file (main.cf) for the PrivNIC resource in the CVM group. Symantec recommends adding all llt private interfaces to the PrivNIC resource.

Perform the tasks described in the following three procedures to configure private IP addresses for CRS:

- Add private IP addresses to the /etc/hosts file.
- Configure the private IP addresses for CRS using the SF Oracle RAC configuration menu.
- Verify the private NIC address configuration

To add private IP addresses to /etc/hosts

- 1 Log in to each system as root
- 2 Using vi or another text editor, add the following entries to the /etc/hosts file:

```
192.168.12.1    galaxy_priv
192.168.12.2    nebula_priv
```

To configure private IP addresses for CRS

- 1 Access the SF Oracle RAC configuration program if you are not currently using it.

See [“Launch the SF Oracle RAC configuration program”](#) on page 198.

For example, enter the following commands:

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

- 2 Select the option to configure private IP addresses for CRS (Oracle RAC 10g only) from the configuration menu.

The path to this configuration menu option is:

```
Main Menu > Prepare to install Oracle >
Configure private IP addresses for CRS (Oracle 10g only)
```

3 Enter the private IP address information for each host.

CRS requires the private NICs on all nodes to be the same; therefore, make sure `/etc/llttab` uses the same `en*` devices on all nodes in the cluster. If the private nics are not the same, CRS will not configure correctly.

```
Enter the private IP for galaxy: [b] 192.168.12.1
```

```
Enter the NIC 1 for private network for galaxy (x if done): [b]  
en1
```

```
Enter the NIC 2 for private network for galaxy (x if done): [b]  
en2
```

```
Enter the NIC 3 for private network for galaxy (x if done): [b] x
```

```
Enter the private IP for nebula: [b] 192.168.12.2
```

```
Enter the NIC 1 for private network for nebula (x if done): [b]  
en1
```

```
Enter the NIC 2 for private network for nebula (x if done): [b]  
en2
```

4 Enter the netmask for the cluster.

```
Enter the netmask for private network: [b] 255.255.255.0
```

5 Review and confirm the private IP address information for CRS.

CRS private IP configuration information verification

System Name: galaxy
CRS IP address: 192.168.12.1

Interfaces: en1 en2

System Name: nebula
CRS IP address: 192.168.12.2

Interfaces: en1 en2

Netmask: 255.255.255.0

Is this information correct? [y,n,q,b] (y)

After confirming the values, the utility adds a new section in the VCS configuration file (main.cf) for the PrivNIC resource in the CVM group.

See “[PrivNIC agent](#)” on page 631.

- 6 After the private IP address is configured, you can exit the installer by entering **q** or continue with your configuration.

To verify the private NIC address configuration

- 1 Access and view the VCS main.cf located in the following directory:
/etc/VRTSvcs/conf/config
- 2 Verify that the PrivNIC resource, ora_priv, shows up in VCS main.cf:

```
PrivNIC ora_priv (
    Critical = 0
    Device @galaxy = { en1 = 0, en2 = 1 }
    Device @nebula = { en1 = 0, en2 = 1 }
    Address @galaxy = "192.168.12.1"
    Address @nebula = "192.168.12.2"
    NetMask = "255.255.255.0"
)
```

- 3 Verify that the "ora_priv" resource is online on all systems in the cluster by issuing the following command:

```
# hares -state ora_priv
# Resource      Attribute      System      Value
ora_priv        State          galaxy      ONLINE
ora_priv        State          nebula      ONLINE
```

- 4 On each system, check the output of "ifconfig"

```
# ifconfig -a
```

- 5 From each system, ping the private IP addresses:

```
# ping nebula_priv
PING 192.168.12.2 (192.168.12.2): 56 data bytes
64 bytes from 192.168.12.2: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 192.168.12.2: icmp_seq=1 ttl=255 time=0 ms
# ping galaxy_priv
PING 192.168.12.1 (192.168.12.1): 56 data bytes
64 bytes from 192.168.12.1: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 192.168.12.1: icmp_seq=1 ttl=255 time=0 ms
```

Identifying the public virtual IP addresses for use by Oracle

Before starting the Oracle installation, you must identify the public virtual IP addresses for each node. Oracle requires one virtual public network address for the Oracle listener process on each node. Public virtual IP addresses are used by client applications to connect to the Oracle database. Either VCS or CRS can manage the virtual public IPs.

Symantec recommends that the public virtual IP name be in this format: hostA-vip

Note: The symbolic name for the nodes must be different than the host's current fully qualified domain name (FQDN).

An IP address and an associated host name should be registered in the domain name service (DNS) for each public network interface.

To identify the public virtual IP addresses for use by Oracle

- 1 Open a second session for this procedure.
- 2 Obtain one virtual IP address per node from your network administrator.

For example, for galaxy:

```
10.10.11.1
```

For example, for nebula:

```
10.10.11.2
```

- 3 On the second session and using either vi or another text editor, add an entry for the virtual IP address and virtual public name in the /etc/hosts file for all nodes. Each node must have a separate public virtual IP address.

For example:

```
10.10.11.1 galaxy-vip
10.10.11.2 nebula-vip
```

- 4 Add the public virtual IP addresses to DNS.

Note: Contact your network administrator for assistance with registering with the DNS.

- 5 After following the above steps, the following entries should appear in the /etc/hosts files for all the nodes:

```
192.168.1.1    galaxy_priv
10.182.13.92  galaxy-vip
192.168.1.2    nebula_priv
10.182.13.93  nebula-vip
```

Creating Oracle disk groups, volumes, and mount points

To create disk groups, volumes, and mount points for Oracle, review the following guidelines.

Before you install the Oracle Cluster Ready Services (CRS) and Oracle RAC 10g binaries, you must create storage space for these installations. You need to provide storage for the following directories and files:

- The home directories, CRS_HOME for CRS and ORACLE_HOME for the Oracle binaries.
See “[Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node](#)” on page 218.
- The CRS files for Oracle Cluster Registry (OCR) and the VOTE-disk. The files can reside in volumes on raw device, or in directories in a cluster file system.
See “[Creating OCR and VOTE-disk volumes](#)” on page 221.

You will create CVM volumes or a directory on CFS for database file storage later in the installation process.

Local or a cluster file system on shared storage

Oracle binaries (ORACLE_HOME) can either be installed locally on each system or on a cluster file system on shared storage. Before proceeding with an installation, consider the following information about each type of Oracle binary installation:


- Local installations provide a comfort level using traditional installation methods and the possibility of improved protection against a single point of failure.
- CFS installations provide a single Oracle installation to manage, regardless of number of nodes. This scenario offers a necessary reduction in storage requirements and easy addition of nodes.

Select the location based on your high availability requirements.

Note: Symantec and Oracle generally recommend using local installations.

[Figure 9-2](#) below displays the installation process for CRS and the Oracle binaries.

Figure 9-2 CRS and Oracle binaries installation process



Preparing \$CRS_HOME and \$ORACLE_HOME on each node

\$CRS_HOME and \$ORACLE_HOME must be prepared on each node for the Oracle installation.

Follow the directions for the following procedures to prepare \$CRS_HOME and \$ORACLE_HOME on each node:

- Identify the directories required for \$CRS_HOME:
See [Table 9-5](#) on page 219.
- Create a file system on local storage for Oracle/CRS binaries (/app):
See [“To create a file system on local storage for Oracle/CRS binaries \(/app\)”](#) on page 219.
- Prepare \$CRS_HOME on each node:
See [“To prepare \\$CRS_HOME on each node”](#) on page 220.
- Prepare \$ORACLE_BASE on each node:
See [“To prepare \\$ORACLE_BASE on each node”](#) on page 220.
- Prepare \$ORACLE_HOME on each node:
See [“To prepare \\$ORACLE_HOME on each node”](#) on page 221.

Identify the directories required for installing Oracle and CRS software.

Table 9-5 Required information \$CRS_HOME on each system

Required information	Example
Disk on each node for Oracle and CRS for Oracle and CRS binaries	Disk_1
Disk group on each local storage	■ On galaxy: orabindg_galaxy ■ On nebula: orabindg_nebula
Volume for each local disk group	■ On galaxy: orabinvol_galaxy ■ On nebula: orabinvol_nebula
File system on each local volume	■ On galaxy: /dev/vx/rdisk/orabindg_galaxy/orabinvol_galaxy ■ On nebula: /dev/vx/rdisk/orabindg_nebula/orabinvol_nebula
Mount point for each local file system	/app
Directory to install CRS	/app

To create a file system on local storage for Oracle/CRS binaries (/app)

- 1 As root user, first create a VxVM local diskgroup, orabindg_hostname:

```
# vxdg init orabindg_galaxy Disk_1
```

- 2 Create a volume, orabinvol_hostname:

```
# vxassist -g orabindg_galaxy make orabinvol_galaxy 12G
```

- 3 Create directory, /app

```
# mkdir /app
```

- 4 Create a filesystem with this volume, orabinvol_hostname

```
# mkfs -V vxfs /dev/vx/rdisk/orabindg_galaxy/orabinvol_galaxy
```

5 Mount /app:

```
# mount -V vxfs /dev/vx/dsk/orabindg_galaxy/orabinvol_galaxy  
/app
```

6 Add an entry for this filesystem. For example:

In /etc/filesystems add:

```
/app:  
dev = dev/vx/dsk/orabindg_galaxy/orabinvol_galaxy  
vfs = vxfs  
mount = true  
options = rw  
account = false
```

To prepare \$CRS_HOME on each node

- 1** On each system, log in as "root".
- 2** Create the directory for CRS_HOME:

```
# mkdir -p /app/crshome
```

- 3** Change ownership and permissions

```
# chown -R oracle:oinstall /app/crshome
```

```
# chmod -R 744 /app/crshome
```

To prepare \$ORACLE_BASE on each node

- 1** On each system, log in as "root".
- 2** Create the directory for ORACLE_BASE:

```
# mkdir -p /app/oracle
```

- 3** Change ownership and permissions

```
# chown -R oracle:oinstall /app/oracle
```

```
# chmod -R 744 /app/oracle
```

To prepare \$ORACLE_HOME on each node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE_HOME:

```
# mkdir -p /app/oracle/orahome
```

- 3 Change ownership and permissions

```
# chown -R oracle:oinstall /app/oracle/orahome

# chmod -R 744 /app/oracle/orahome
```

Creating OCR and VOTE-disk volumes

The CRS installation requires predefined locations for the Oracle Cluster Registry (OCR) and VOTE-disk components.

Note: Whether you create volumes or file system directories, you can add them to the VCS configuration to make them highly available.

Review the following information and perform the procedures described below to create OCR and VOTE-disk volumes and file systems:

- Information to create OCR and Voting Disk directories on CFS:
See [Table 9-6](#) on page 221.
- Information to create a filesystem for OCR and VOTE disks (/ocrvote):
See [“To create a filesystem for OCR and VOTE disks \(/ocrvote\)”](#) on page 223.

Identify the directories required for installing Oracle and CRS software.

Table 9-6 Information to create OCR and Voting Disk directories on CFS

Required information	Example
Shared disk	HDS0_10,HDS0_21
Shared disk group	ocrvotedg
Volume for shared disk group	ocrvotevol
File system on shared volume	/dev/vx/rdsk/ocrvotedg/ocrvotevol

Table 9-6 Information to create OCR and Voting Disk directories on CFS
(continued)

Required information	Example
Mount point for shared file system	/ocrvote
Directory for OCR files	ocr
Directory for Voting Disk files	vote
Local disk on each node	HDS0_11
Disk group for each local disk	■ On galaxy: orabindg_galaxy ■ On nebula: orabindg_nebula
Volume for each local disk group	■ On galaxy: orabinvol_galaxy ■ On nebula: orabinvol_nebula
File system on each local volume	■ On galaxy: /dev/vx/rdisk/orabindg_galaxy/ or abinvol_galaxy ■ On nebula: /dev/vx/rdisk/orabindg_nebula/ or abinvol_nebula
Directory for Oracle base (ORACLE_BASE)	/app/oracle
Directory to install Oracle (ORACLE_HOME):	/app/oracle/orahome

To create a filesystem for OCR and VOTE disks (/ocrvote)

- 1** As root user, from the CVM master, create a shared VxVM diskgroup

```
# vxdg -s init ocrvotedg HDS0_10 HDS0_21
```

Note: To determine the CVM master, issue the `vxctl -c mode` command.

- 2** As root user, from the CVM master, create a mirrored volume, ocrvotevol:

```
# vxassist -g ocrvotedg make ocrvotevol 1G nmirrors=2
```

- 3** As root user, from CVM master, create a filesystem with the volume, ocrvotevol.

```
# mkfs -V vxfs /dev/vx/dsk/ocrvotedg/ocrvotevol
```

- 4** On each system, create a directory, /ocrvote:

```
# mkdir /ocrvote
```

- 5** On each system, mount /ocrvote

```
# mount -V vxfs -o cluster /dev/vx/dsk/ocrvotedg/ocrvotevol \  
/ocrvote
```

- 6** As root user, from any system, change permissions on /ocrvote

```
# chown -R oracle:oinstall /ocrvote
```

- 7** The following steps add the CFSSMount and CVMVolDg resources corresponding to this "/ocrvote" filesystem to the "cvm" group in VCS configuration.

Enter the following command to change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 8 Enter the following commands to set up the file systems:

```
# hares -add ocrvote_mnt CFSMount cvm

# hares -modify ocrvote_mnt Critical 0

# hares -modify ocrvote_mnt MountPoint "/ocrvote"

# hares -modify ocrvote_mnt BlockDevice \
"/dev/vx/dsk/ocrvotedg/ocrvotevol"

# hares -modify ocrvote_mnt MountOpt mincache=direct
```

- 9 Enter the following commands to set up the volumes:

```
# hares -add ocrvote_voldg CVMVolDg cvm

# hares -modify ocrvote_voldg CVMDiskGroup ocrvotedg

# hares -modify ocrvote_voldg CVMVolume -add ocrvotevol

# hares -modify ocrvote_voldg CVMActivation sw
```

- 10 Enter the following commands to link parent to the child:

```
# hares -link ocrvote_mnt ocrvote_voldg

# hares -link ocrvote_mnt vxfsckd

# hares -link ocrvote_voldg cvm_clus
```

- 11 Enter the following commands to tell the system to come up enabled:

```
# hares -modify ocrvote_voldg Enabled 1

# hares -modify ocrvote_mnt Enabled 1

# haconf -dump -makero
```


- 12** Verify the stanzas corresponding to the `ocrvote_mnt` and `ocrvote_voldg` resources in `main.cf`:

```
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  
)  
  
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
ocrvote_voldg requires cvm_clus
```

- 13** Verify that the VCS resources, `ocrvote_mnt`, and `ocrvote_voldg`, are ONLINE on all systems in the cluster.

```
# hares -state ocrvote_mnt  
  
# hares -state ocrvote_voldg
```

- 14** If the `voldg` is manually online, the `CVMVoldg` resources under VCS will not report as online. VCS must be stopped and restarted.

The following procedure can also be used to create a filesystem for OCR and VOTE disks (/ocrvote) by editing the `main.cf` file for Oracle RAC 10g.

Note: The following procedure requires a restart on all the nodes.

To create a filesystem for OCR and VOTE disks (/ocrvote) by editing the main.cf file

- 1 Log in to one system as root.
- 2 Save your existing configuration to prevent any changes while you modify main.cf:

```
# haconf -dump makero
```

If your configuration is not writable, a warning appears: "Cluster not writable".

You may safely ignore the warning.

- 3 Make sure that VCS is not running while you edit the main.cf by using the hastop command.

This command stops the VCS engine on all systems and leaves the resources available.

```
# hastop -all -force
```

- 4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
```

```
# cp main.cf main.orig
```

- 5 Using vi or another text editor, edit the main.cf file, modifying the cvm service group.

Specifically, edit the stanzas corresponding to the ocrvote_mnt and ocrvote_voldg resources in main.cf:

```
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
)

ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
```

- 6 Save and close the main.cf file by verifying the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

- 7 Use the following command to start the VCS engine on one system:

```
# hstart
```

- 8 Enter the following command:

```
# hastatus
```

- 9 When "LOCAL_BUILD" is listed in the message column, start VCS on the other system with the following command:

```
# hstart
```

- 10 Verify that the service group resources are brought online. On one system, enter the following command:

```
# hagrps -display
```

- 11 Enter the following command to check the status of the cvm group:

```
# hagrps - state
```

- 12 Enter the following command to check the status of resources:

```
# hagrps - state
```

Restart VCS on all the nodes.

Verifying SF Oracle RAC resources on all nodes

With the configuration file in place on each node, verify that the resources you configured come online.

To verify VCS, CVM, and CFS on all nodes

- 1 To verify that the resources are online, enter the following command on each node:

```
# hastatus
```

- 2 If VCS is not running, then use the following procedure to start VCS, CVM, and CFS on all the nodes.

To start VCS, CVM, and CFS on all nodes

- ◆ Start VCS, CVM, and CFS from each node by issuing the following command:

```
# hstart
```

If VCS is not running, then verify GAB port membership by entering the `gabconfig -a` command

Verifying GAB port membership

After starting VCS, CVM, and CFS on each node, verify GAB port membership by running 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 v gen    ada416 membership 01
```

```
Port w gen    ada418 membership 01
```


Installing Oracle RAC 10g

This chapter includes the following topics:

- [About installing Oracle RAC 10g](#)
- [Installing Oracle Clusterware \(CRS\)](#)
- [Installing Oracle RAC 10g database binaries](#)
- [Verifying the Oracle RAC 10g CRS and database installation](#)
- [Completing Oracle RAC 10g post-installation tasks](#)


About installing Oracle RAC 10g

After completing the Oracle RAC 10g pre-installation tasks, proceed to install Oracle RAC 10g software locally on each node. For supported versions of Oracle: See [“Supported software”](#) on page 38.

Note: Review your Oracle installation manuals and the appropriate Oracle support websites, before installing Oracle RAC 10g.

[Figure 10-1](#) displays the steps required for a Oracle RAC 10g installation. Each step is discussed in detail in this chapter.

Figure 10-1 Oracle RAC 10g installation flowchart



Installing Oracle RAC 10g in an SF Oracle RAC environment involves the following tasks:

- [Installing Oracle Clusterware \(CRS\)](#)
- [Installing Oracle RAC 10g database binaries](#)
- [Verifying the Oracle RAC 10g CRS and database installation](#)
- [Completing Oracle RAC 10g post-installation tasks](#)

Installing Oracle Clusterware (CRS)

If you selected Install Oracle Clusterware (CRS) in the SF Oracle RAC configuration program, then after you set the Oracle RAC 10g environment variables, the program launches the Oracle utility to install CRS.

The CRS software is installed on each node in the location created for it.

See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 218.

To install CRS

- 1 Launch the installsfrac program. As root user on any one node in the cluster, enter:

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

By default, the installsfrac utility uses ssh for remote communication. However, rsh can be used in place of ssh by using the "rsh" option with the installsfrac utility. For an example of using the "rsh" option with the installsfrac utility:

```
# ./installsfrac -rsh -configure
```

- 2 The installer displays the cluster name, cluster ID, systems, and service groups. When prompted whether to configure SFRAC from the displayed systems, press enter (or "y") to proceed.
- 3 Press Enter to continue after the initial system check is successfully completed.
- 4 Press Enter to continue when the licenses are successfully verified.
- 5 Navigate to the Main Menu and select Install or relink Oracle. Some choices for installing and configuring may vary, depending on the operating system you are running.
- 6 In the Choose Oracle version menu, select the appropriate version of Oracle (10gR1 or 10gR2).
- 7 In the Choose task menu, select the task Install Oracle Clusterware (CRS).
- 8 In the Set DISPLAY dialog, enter the value for DISPLAY.

This value should refer to the address of the X11 server where you want OUI for CRS installation displayed.
- 9 Enter Oracle user and group information:

- In the Oracle Username dialog, enter Oracle Unix User Account when prompted. The installer checks for the user on all systems.
- In the Oracle Groupname dialog, enter Oracle Inventory group when prompted. The installer checks for group existence on all systems.
- Press Enter to continue.

10 As you run the installer, be prepared with the following information required for the installation and configuration of the CRS component:

Specify File Locations dialog box	■ The installation destination (\$CRS_HOME) on each node. For example, /app/crshome. ■ The path to \$CRS_HOME.
Cluster Configuration dialog box	The names of the nodes and their associated host names. In the example CRS installation, the nodes and their associated host names are galaxy and nebula.
Private Interconnect Enforcement dialog box	The private NIC IP addresses you have set up for each node. For example, lan1 on subnet 1.0.0.0.
Oracle Cluster Registry dialog box	The name of a file in the OCR directory or raw volume. For example, a file named ocr_file in the directory /ocrvote/ocr.
Voting Disk dialog box	The name of a file in the Vote-disk directory or raw volume. For example, a file named vote_file in the directory /ocrvote/vote.

- 11 In the CRS Install Image dialog, enter the absolute path of CRS install image when prompted. The installer validates the CRS installer. Press Enter to continue.
- 12 In the Oracle base directory dialog, enter the location of base of the Oracle directory structure for OFA compliant databases. The installer validates the Oracle Base Directory. If the directory does not exist , installer prompts for the creation of oracle base directory on all nodes. Choose option 'y' to create oracle base directory on all nodes.
- 13 Press Enter to continue.

- 14 In the CRS Home directory dialog, enter absolute path of CRS home directory when prompted. Installer validates the CRS home directory. If the directory does not exist, installer prompts for the creation of the directory on all nodes. Choose option "y" to create oracle base directory on all nodes.
- 15 Press Enter to continue.
- 16 In the Patch Level dialog, specify the patch level of Oracle CRS/Database Software being installed or relinked. Accept the default when installing base version of Oracle RAC 10gR1 (10.1.0.2) or Oracle RAC 10gR2 (10.2.0.1)
- 17 The Installer prints the CRS installation information for verification. If the information displayed by the installer is correct choose option "y" otherwise choose option "n".

Example:

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Installation Path:
/cdrom/Oracle10g/cluserware/Disk1
Oracle Clusterware (CRS) Home: /app/oracle/crshome
Oracle Release: 10.2

Oracle Patch Level: 0.1

Oracle Base: /app/oracle
Is this information correct? [y,n,q] (y)
```

- 18 Press Enter to continue.
- 19 The installer now validates the inputs and performs certain setup steps, after which it prompts you to run rootpre.sh on all cluster nodes.
Do not respond to the prompt yet, but proceed to step 20.
- 20 On each system, log in as root in another terminal session.
 - On each system, change the directory to the location of the rootpre.sh script, and invoke the script.
 - After invoking rootpre.sh on all nodes, return to the installsfrac prompt in step 19

- Enter "y" to indicate that the script has been invoked on all systems in the cluster.

21 Installer now invokes the OUI for Oracle CRS Software Installation.

```
Invoking OUI to install Oracle CRS. Please wait .. ok
```

```
Now install Oracle CRS software from the OUI. Refer to  
Oracle documentation for installation instructions.
```

```
NOTE: Near the end of the installation, Oracle CRS  
Software installer will prompt you to run the root.sh  
script.
```

```
Press <RETURN> here when prompted to run root.sh
```

Note: Do not press Return yet, but wait for the prompt to do so.

To install Oracle 10.2.0.1 CRS using the Oracle interface

- 1 If the "Specify Inventory Directory and Credentials" window is presented, enter the location of Oracle inventory and the AIX group name.
- 2 When the Oracle Universal Installer (OUI) for Oracle CRS Installation displays, navigate to the "Specify Home Details" window. Change OUIHome1 to CRS_HOME and specify the path for \$CRS_HOME.
- 3 Make sure that all Product-Specific Prerequisite Checks are successful
- 4 Make sure that all systems in the cluster appear on the "Specify Cluster Configuration" screen. Modify the private name and virtual name as required.
- 5 When the "Specify Network Interface Usage" window presents a list of interface devices for which IP addresses are configured, verify that the public interface and private interface match those you have configured for Oracle RAC 10g. If devices are shown on the list that are not part of the configuration for Oracle RAC 10g, be sure they are labeled with "Do Not Use".
- 6 On the "Oracle Cluster Registry" window:
 - Choose External Redundancy
 - Set "OCR Location" to "/ocrvote/ocr"
- 7 In the "Voting Disk" window:
 - Choose External Redundancy
 - Set "Voting Disk Location" to "/ocrvote/vote"

- 8 Confirm the settings on the final screen.
Ensure that /tmp has enough disk space to be able to accommodate large temporary files that are created during the installation.
Select "Install" to start the CRS installation.
- 9 When Oracle CRS binaries have been installed, the OUI prompts you to invoke two scripts, namely, orainstRoot.sh and root.sh.
Do not run root.sh yet, but return to the installsfrac session from the step in the previous procedure where you were prompted to run rootpre.sh.
See ["To install CRS"](#) on page 233.
- 10 In the installsfrac session, press "Return". The installsfrac utility now verifies the CRS installation and copies the SF Oracle RAC libraries to \$CRS_HOME.
- 11 Next, login as "root" on each system and invoke the two scripts, orainstRoot.sh and root.sh.

Note: You must NOT run these scripts simultaneously on all systems.

- 12 Return to the OUI and indicate that you have run root.sh on all systems in the cluster.
- 13 Exit the OUI.
- 14 If the installation and setup is successful, then root.sh should be able to start the CRS daemons on each node. To verify if CRS is up and running, enter:

```
# /app/crshome/bin/crs_stat -t
```

Name	Type	Target	State	Host

ora....axy.gsd	application	ONLINE	ONLINE	nebula
ora....axy.ons	application	ONLINE	ONLINE	nebula
ora....axy.vip	application	ONLINE	ONLINE	nebula
ora....ula.gsd	application	ONLINE	ONLINE	galaxy
ora....ula.ons	application	ONLINE	ONLINE	galaxy
ora....ula.vip	application	ONLINE	ONLINE	galaxy

- 15 If the `crs_stat` command fails with message "failed to contact crs daemons", it means the CRS installation was unsuccessful.

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

- 16 Oracle CRS installation is now complete. Add a resource to the VCS configuration to monitor the CSSD daemon:

See "[CSSD agent](#)" on page 638.

Installing Oracle RAC 10g database binaries

After installing the CRS component, proceed to install the Oracle database software.

Review the preliminary procedures to install the Oracle RAC 10g software in an SF Oracle RAC environment:

See "[Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node](#)" on page 218.

Note: Symantec recommends that you install Oracle binaries locally on each node.

When installing the database software for Oracle 10.1.0.2 (10gR1 base installation), you must copy the installation image to a directory that is writable by the root user.

To install Oracle RAC 10g database binaries

- 1 If the `installsfrac` program is not already started and running for previous procedures, then launch the `installsfrac` program.

As root user on any one node in the cluster, enter:

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

By default, the `installsfrac` utility uses `ssh` for remote communication. However, `rsh` can be used in place of `ssh` by using the "`rsh`" option with the `installsfrac` utility.

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

- 2 Navigate to the Main Menu and select **Install and Relink Oracle**.

- 3 In the Choose Oracle version menu, select the appropriate version of Oracle (10gR2).
- 4 In the Choose task menu, select **Install Oracle RDBMS server**.
- 5 In the Set DISPLAY dialog, enter the value for DISPLAY.

This value should refer to the address of the X11 server where you want OUI for CRS installation displayed..
- 6 Enter Oracle user and group information.
 - In the Oracle Username dialog, enter Oracle Unix User Account when prompted. The installer checks for the user on all systems.
 - In the Oracle Groupname dialog, enter Oracle Inventory group when prompted. The installer checks for group existence on all systems.
 - Press Enter to continue.
- 7 In the Database Install Image dialog, specify the location of Oracle Database Software install image.

Example: /cdrom/Oracle10g/database/Disk1
- 8 In the Oracle base directory dialog, specify the base of the Oracle directory structure for OFA compliant databases.

Example: /app/oracle
- 9 In the CRS Home directory dialog, specify the directory containing Oracle CRS Software.

Example: /app/crshome
- 10 In the Database home directory dialog, specify the directory containing Oracle Database Software.

Example: /app/oracle/orahome
- 11 In the Patch Level dialog, specify the patch level of Oracle CRS/Database Software being installed or relinked. Accept the default when installing base version of Oracle RAC 10gR1 (10.1.0.2) or Oracle RAC 10gR2 (10.2.0.1)

- 12** The installer prints the oracle environment information for verification. If the information displayed by the installer is correct choose option "y" otherwise choose option "n."

For example:

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 10.2
Oracle Patch Level: 0.1
Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
RDBMS Installation Path: /cdrom/Oracle10g/database/Disk1
Is this information correct? [y,n,q] (y)
```

Press Enter to continue.

The SF Oracle RAC installer invokes the Oracle User Interface (OUI) for Database Software Installation.

- 13** The installer now validates the inputs and performs certain setup steps, after which it prompts you to run `rootpre.sh` on all cluster nodes.

```
Verifying CRS on galaxy ..... Running
Verifying CRS on nebula ..... Running
Comparing Oracle UNIX user id on all systems ..... Done
Comparing Oracle group id on all systems ..... Done
```

ORACLE DB INSTALLATION

For installation to proceed, Oracle requires you to run the `rootpre.sh` script located under
`/cdrom/Oracle10g/database/Disk1/rootpre.`

Run the following commands on all nodes:

```
# cd /cdrom/Oracle10g/database/Disk1/rootpre
# ./rootpre.sh
```

The following output will be displayed:

```
Has rootpre.sh been run by root? [y,n,q] y
```

At this point, run `rootpre.sh` on all nodes, then type 'y' return.

- 14** The installer now brings up the OUI for Oracle RDBMS Software Installation.

```
Invoking OUI to install Oracle Software. Please wait .. ok
```

```
Now install Oracle database software from the OUI. Refer to
Oracle documentation for installation instructions.
```

```
NOTE: Near the end of the installation, Oracle Database
Software installer will prompt you to run the root.sh
script.
```

```
Press <RETURN> here when prompted to run root.sh
```

Note: Do not press Return yet, but wait for the prompt to do so.

- 15 When the OUI for Oracle Database Software installation displays, follow the instructions and respond to the prompts to install Oracle RAC 10g software. When you are prompted for the names of the systems for installation, select all nodes. In the "Select Database Configuration" window, choose the option, "Do not create a starter database."

- 16 Confirm the settings on the final OUI window.

Ensure sufficient disk space on /tmp, so that large temporary files can be created in the /tmp area during this installation.

Select "Install" to start the Oracle database software installation.

- 17 When the Oracle database binaries have been installed, the OUI prompts you to invoke the root.sh script.

Do not run root.sh yet.

Return to the installsfrac session from step 14.

- 18 In the installsfrac session, press Return. The installsfrac utility now verifies the database software installation, copies the SF Oracle RAC libraries to \$ORACLE_HOME, and relinks Oracle on each node in the cluster.

```
Verifying binaries in /app/oracle/orahome on galaxy ...ok
Verifying binaries in /app/oracle/orahome on nebula ...ok
Copying SFRAC libskgxn on galaxy .....ok
Copying SFRAC libskgxn on nebula .....ok
Copying SFRAC ODM library on galaxy .....ok
Copying SFRAC ODM library on nebula .....ok
Copying SFRAC libskgxp on galaxy .....ok
Copying SFRAC libskgxp on nebula .....ok
Relinking Oracle on galaxy .....ok
Relinking Oracle on nebula .....ok
Oracle post-installation setup is now complete.
```

You may now invoke root.sh to complete Oracle installation

Press <RETURN> here after running root.sh on all nodes:

- 19 As root user on each system, invoke the root.sh script:

```
# cd /app/oracle/orahome
# ./root.sh
```

- 20** After the `root.sh` script is run on all nodes, hit return as required by the previously displayed directions:

Press <RETURN> here after running `root.sh` on all nodes

- 21** Return to the OUI and indicate you have run `root.sh`.
- 22** Exit the OUI.
- 23** After completing installation of Oracle RAC 10g software, you must install the appropriate patches.

See [“Adding Oracle RAC 10g patchsets”](#) on page 244.

Verifying the Oracle RAC 10g CRS and database installation

To verify that the installations of the Oracle CRS and Oracle RAC 10g have succeeded, issue the command described below from any node in the cluster. The output should show processes running on all nodes, as in the following example:

```
# /app/crshome/bin/crs_stat
NAME=ora.galaxy.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy

NAME=ora.galaxy.gsd
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy

NAME=ora.galaxy.ons
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy

NAME=ora.nebula.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on nebula

NAME=ora.nebula.gsd
TYPE=application
```

```
TARGET=ONLINE
STATE=ONLINE on nebula

NAME=ora.nebula.ons
TYPE=application
TARGET=ONLINE
STATE=ONLINE on nebula
```

Completing Oracle RAC 10g post-installation tasks

After installing the Oracle software, complete these tasks:

- [Adding Oracle RAC 10g patchsets](#)
- [Relinking the SF Oracle RAC libraries to Oracle RAC 10g](#)
- [Upgrading Oracle databases](#)
- [Creating the Oracle RAC 10g database](#)
- [Configuring the Oracle RAC 10g service group in a VCS configuration](#)

Adding Oracle RAC 10g patchsets

Before installing Oracle RAC 10g patch or patchset software, perform the following tasks:

- Review the latest information on supported Oracle patches and patchsets:
<http://entsupport.symantec.com/docs/280186>
- You must have installed Oracle RAC 10g R1 (10.1.0.2) or Oracle RAC 10g R2 (10.2.0.1) software.
- Review the patch or patchset notes that accompany the patch or patchset for instructions on installing them and performing the post-installation operations.

Use the following procedure to add Oracle RAC 10g R1 or Oracle RAC R2 patchsets to your node if you have installed Oracle, but have not yet configured Oracle in your cluster.

Applying Oracle 10.1.0.5 Patchset on AIX

Applying Oracle 10.1.0.5 patchset involves the following tasks:

- Install 10.1.0.5 CRS
- Install 10.1.0.5 Database Software
- Relink Oracle

To install 10.1.0.5 CRS

- 1 Stop Oracle instances on all systems in the cluster.
- 2 Stop CRS on all system in the cluster.
- 3 Log in as root and change permissions for \$CRS_HOME.

```
# chown -R oracle:oinstall /app/crshome
```

- 4 Refer to Oracle 10.1.0.5 Patch note and apply the patch to CRS binaries.
- 5 Near the end of the CRS patch installation, OUI will prompt you to invoke root10105.sh on all systems in the cluster. Log in as root user on each system in the cluster and run the following commands:

- Change directory to \$CRS_HOME/install/patch10105/lib

```
# cd $CRS_HOME/install/patch10105/lib
```

- Check if the 64-bit libskgxn2.a and libskgxn2.so exist. If they do, copy SF Oracle RAC SKGXM libraries in their place.

```
# ls libskgxn2.a libskgxn2.so
```

```
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a libskgxn2.a
```

```
# cp /opt/VRTSvcs/rac/lib64/libvcsmm.so libskgxn2.so
```

- Check if the 64-bit libskgxp10.a exists. If it does, copy SF Oracle RAC SKGXP library in its place.

```
# ls libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a libskgxp10.a
```

- Change directory to \$CRS_HOME/install/patch10105/lib32

```
# cd $CRS_HOME/install/patch10105/lib32
```

- Check if 32-bit libskgxn2.a and libskgxn2.so exist. If they do, copy SF Oracle RAC SKGXM libraries in their place.

```
# ls libskgxp10.a libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a libskgxp10.a
```

- Check if 32-bit libskgxp10.a exists. If it does, copy SF Oracle RAC SKGXP library in its place.

```
# ls libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.a libskgxp10.a
```

- 6 Now invoke the script, root10105.sh. This script copies everything from under \$CRS_HOME/install/patch10105 directory to their corresponding location under \$CRS_HOME. It then starts the CRS daemons. You must run this script on each system in the cluster (in the same order as specified by Oracle installer).

```
# cd $CRS_HOME/install
```

```
# ./root10105.sh
```

- 7 Verify that CRS daemons are up on systems in the cluster. The command, crs_stat, should list all the CRS resources.

```
# $CRS_HOME/bin/crs_stat -t
```

To install 10.1.0.5 Oracle database Software

- 1 Refer to Oracle 10.1.0.5 Patch note and apply the patch to Oracle binaries.
- 2 Run root.sh script at the end of the patch installation as instructed by Oracle.
- 3 Relink the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

Applying Oracle RAC 10g Release 2 Patchset on AIX

Applying Oracle RAC 10g Release 2 patchset involves the following tasks:

- Install 10.2.0.3 CRS
- Install 10.2.0.3 Database Software
- Relink Oracle

To install 10.2.0.3 CRS

- 1 Stop Oracle instances on all systems in the cluster.
- 2 Stop CRS on all system in the cluster.

- 3** Log in as root and change permissions for \$CRS_HOME.

```
# chown -R oracle:oinstall /app/crshome
```

- 4** Refer to Oracle 10.2.0.3 Patch note and apply the patch to CRS binaries.

- 5** Near the end of the CRS patch installation, OUI will prompt you to invoke root102.sh on all systems in the cluster. Log in as root user on each system in the cluster and run the following commands:

- Change directory to \$CRS_HOME/install/patch102/lib

```
# cd $CRS_HOME/install/patch102/lib
```

- Check if the 64-bit libskgxp10.a and libskgxp10.a libraries exist. If they do, copy SF Oracle RAC SKGXP libraries in their place.

```
# ls libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a libskgxp10.a
```

- Change directory to \$CRS_HOME/install/patch102/lib32

```
# cd $CRS_HOME/install/patch102/lib32
```

- Check if 32-bit libskgxp10.a and libskgxp10.a libraries exist. If they do, copy SF Oracle RAC SKGXP library in their place.

```
# ls libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.a libskgxp10.a
```

```
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.a libskgxp10.a
```

- 6
- Now invoke the script, root102.sh. This script copies everything from under \$CRS_HOME/install/patch102 directory to their corresponding location under \$CRS_HOME. It then starts the CRS daemons. You must run this script on each system in the cluster (in the same order as specified by Oracle installer).

```
# cd $CRS_HOME/install
# ./root102.sh
```

- 7
- Verify that CRS daemons are up on systems in the cluster. The command, crs_stat, should list all the CRS resources.

```
# $CRS_HOME/bin/crs_stat -t
```

To install 10.2.0.3 Oracle database Software

- 1
- Make sure CRS is running on all the nodes.
- 2
- Stop the node applications by issuing the following commands:

```
# /crshome/crs/bin/srvctl stop nodeapps -n galaxy
# /crshome/crs/bin/srvctl stop nodeapps -n nebula
```

Note: CRS should be running on all nodes, but all objects should be offline.

- 3
- Run the following command and make sure that your CRS displayed output is similar to the following output:

```
# /crshome/crs/bin/crs_stat -t
```

Name	Type	Target	State	Host
-----	-----	-----	-----	-----
ora....31.inst	application	OFFLINE	OFFLINE	
ora....32.inst	application	OFFLINE	OFFLINE	
ora.DBS3.db	application	OFFLINE	OFFLINE	
ora....XY.lsnr	application	OFFLINE	OFFLINE	
ora...axy.gsd	application	OFFLINE	OFFLINE	
ora...axy.ons	application	OFFLINE	OFFLINE	
ora...axy.vip	application	OFFLINE	OFFLINE	
ora....LA.lsnr	application	OFFLINE	OFFLINE	
ora....ula.gsd	application	OFFLINE	OFFLINE	
ora....ula.ons	application	OFFLINE	OFFLINE	
ora....ula.vip	application	OFFLINE	OFFLINE	

- 4 Refer to Oracle 10.2.0.3 Patch note and apply the patch to Oracle binaries.
- 5 Run root.sh script at the end of the patch installation as instructed by Oracle.
- 6 Relink the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

Relinking the SF Oracle RAC libraries to Oracle RAC 10g

If you added or upgraded the Oracle patches, you must relink the SF Oracle RAC libraries to Oracle. You must link Oracle with the Veritas IPC library and enable Oracle to use the Veritas ODM and cluster membership (VCSMM) libraries. You can relink the libraries using the SF Oracle RAC configuration program.

Note: Symantec recommends you to relink the SF Oracle RAC libraries only after completing all the required patch additions.

To relink the SF Oracle RAC libraries to Oracle on AIX

- 1 If the installsfrac program is not already started and running for previous procedures, then launch the installsfrac program.

As root user on any one node in the cluster, enter:

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

By default, the installsfrac program uses ssh for remote communication. However, RSH can be used in instead of SSH by using the "-rsh" option with the installsfrac utility.

```
# cd /opt/VRTS/install
# installsfrac -rsh -configure
```

- 2 Select 'Install or relink Oracle'.
- 3 Select the version of Oracle software (Oracle RAC 10gR1 or Oracle RAC 10gR2).
- 4 Select 'Relink Oracle'.
- 5 In the Oracle user name screen, enter the Oracle UNIX user name when prompted. The installer checks for the user on all systems.
- 6 In the Oracle Group name screen, enter the name of the Oracle inventory group when prompted. The installer checks for the existence of this group on all systems.

- 7 In the Oracle base directory screen, specify the base of the Oracle directory structure for OFA compliant databases.
- 8 In the Database Home Directory screen, specify the location of \$ORACLE_HOME.
- 9 In the Patch Level screen, specify the last two digits of the patch level of Oracle database software being installed or relinked. For example, enter 0.5 for 10.1.0.5; enter 0.2 for 10.2.0.2; enter 0.3 for 10.2.0.3.
- 10 The installer utility prints the entered information for verification. Answer "y" if the information displayed is correct, answer "n" otherwise.

For example:

```
Oracle environment information verification
Oracle Unix user: oracle
Oracle Unix Group: dba
Oracle Release: 10.2
Patch Level: 0.3
Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct: [y,n,q] (y)
```

- 11 Installer now copies 5.0 SF Oracle RAC libraries to ORACLE_HOME and then relinks Oracle.

```
Verifying binaries in /app/oracle/orahome on galaxy ...ok
Verifying binaries in /app/oracle/orahome on nebula ...ok
Copying SFRAC libskgxn on galaxy .....ok
Copying SFRAC libskgxn on nebula .....ok
Copying SFRAC ODM library on galaxy .....ok
Copying SFRAC ODM library on nebula .....ok
Copying SFRAC libskgxp on galaxy .....ok
Copying SFRAC libskgxp on nebula .....ok
Relinking Oracle on galaxy .....ok
Relinking Oracle on nebula .....ok
```

Oracle Relinking is now complete.

- 12 After relinking is complete, enter "q" to quit the SF Oracle RAC installer.

Upgrading Oracle databases

If your configuration already has a database and you are upgrading from 10.2.0.1 to 10.2.0.3, then you need to run the dbua (Database Upgrade Assistant) or manually upgrade the database.

To migrate existing Oracle databases to the latest Oracle patch level:

See the README.html file downloaded with the patch. The file is located in the /oracle/patch directory.

Creating the Oracle RAC 10g database

To set up a new Oracle RAC 10g database:

- Create the database on shared storage.
- Use your own tools or review the guidelines on using the Oracle dbca (Database Creation Assistant) tool.
- For instructions on installing the Oracle RAC 10g database:
Refer to your latest Oracle documentation.
- To create a database on shared raw VxVM volumes or shared VxFS file systems:
See [“About creating a test database”](#) on page 617.

Note: If your configuration will support replicating global clusters using SF Oracle RAC, then the Oracle RAC 10g database should not be set up on a GCO secondary site. Follow the above steps only when setting up Oracle on a primary GCO site.

Configuring the Oracle RAC 10g service group in a VCS configuration

After installing Oracle RAC 10g and creating a database, proceed to modify the VCS configuration file. Review the sample VCS configurations and details on configuring service groups in an Oracle RAC 10g environment

.See [“Configuring VCS resources for Oracle”](#) on page 260.

Configuring VCS service groups for Oracle RAC 10g

This chapter includes the following topics:


- [About VCS service group for Oracle RAC 10g dependencies](#)
- [Configuring VCS resources for Oracle](#)
- [Location of VCS log files](#)

About VCS service group for Oracle RAC 10g dependencies

A VCS service group is a collection of resources working together to provide application services to clients. A VCS service group typically includes multiple resources that are both hardware and software based. For example, a resource may be a physical component such as a disk or network interface card, or a software component such as Oracle RAC 10g or a Web server, or a configuration component such as an IP address or mounted file system.

[Figure 11-1](#) displays the steps required to configure the VCS service groups for an Oracle RAC 10g installation. Each step is discussed in detail in this chapter.

Figure 11-1 VCS service group configuration steps



VCS service group dependencies with or without the VCS Oracle agent

VCS service group dependencies are based on whether you use the VCS Oracle agent or not. The following figures illustrate the dependencies.

- In a configuration *with* the VCS Oracle agent, VCS controls the Oracle database. This is the preferred configuration. An online local firm dependency exists between the Oracle group and the CVM group.

Note: In a setup with multiple databases, the configuration with the VCS Oracle agent must be used.

For additional details on service group dependencies, refer to the *Veritas Cluster Server User's Guide*.

[Figure 11-2](#) displays a schematic of a configuration with the VCS Oracle agent, and the dependencies between the VCS service groups and resources. In this figure, the name inside the circular shape is the "resource type". The name just outside of the circular shape is the "resource name".

- In a configuration *without* the VCS Oracle agent, CRS controls the database.

Note: The configuration without the VCS Oracle agent may be used only in single database setups.

Figure 11-3 displays a schematic of a configuration without the VCS Oracle agent, and the dependencies between the VCS service groups and resources. In this figure, the name inside the circular shape is the "resource type". The name just outside of the circular shape is the "resource name".

Figure 11-2 Configuration with the VCS Oracle agent (Preferred configuration)




Table 11-1 defines the CVM service group resources displayed in Figure 11-2.

Figure 11-3 Configuration without the VCS Oracle agent




Table 11-1 defines the CVM service group resources displayed in Figure 11-3.

Table 11-1 CVM service group resource definitions

Resource name	Resource type	Definitions
cssd	Application	The cssd resource monitors the Oracle RAC 10g clusterware processes. The purpose of the cssd resource is to ensure that the dependency of cssd on the OCR and VOTE resources and the PrivNIC resource (optional) are satisfied. For additional information: See “CSSDagent” on page 638.

Table 11-1 CVM service group resource definitions (*continued*)

Resource name	Resource type	Definitions
ora_priv	PrivNIC	The PrivNIC resource can be used to maintain a private IP address that is locally highly available on LLT ethernet interfaces. Such private IP addresses are required by the CRS daemons in Oracle RAC 10g to provide communication between nodes. For additional information: See “PrivNIC agent” on page 631.
vxfsckd	CFSfsckd	The Veritas file system check daemon (vxfsckd) is responsible for ensuring file system consistency when a node fails that was a primary node for a shared file system. Cluster file system check daemon (CFSfsckd) is responsible for ensuring file system consistency in a cluster when a node fails that was a primary node for a shared file system.
cvm_clus	CVMCluster	The CVMCluster agent controls system membership on the cluster port associated with Veritas Volume manager (VxVM). For additional information: See “CVMCluster agent” on page 623.

Table 11-1 CVM service group resource definitions (*continued*)

Resource name	Resource type	Definitions
cvm_vxconfigd	CVMVxconfigd	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 configuration information stored on disks. CVMVxconfigd must be present in the CVM service group. For additional information: See “CVMVxconfigd agent” on page 625.
ocrvote_voldg	CVMVolDg	The CVMVolDg agent represents and controls CVM disk groups and the CVM volumes within the disk groups. For additional information: See “CVMVolDg agent” on page 627.
ora_db	Oracle	Oracle database agent. Oracle application is the resource type. Note: This resource is only available for the configuration with the VCS Oracle agent.

Table 11-1 CVM service group resource definitions (*continued*)

Resource name	Resource type	Definitions
crs.ora.mnt	CFS_Mount	The Cluster Ready Service Oracle Mount agent (crs_ora_mnt) brings online, takes offline, and monitors an Oracle cluster file system mount point. The CFSSMount agent brings online, takes offline, and monitors a cluster file system mount point. Note: This resource is only available for the configuration with the VCS Oracle agent. For additional information: See “CFSSMount agent” on page 629.
oradata_mnt	CFSSMount	The CFSSMount agent brings online, takes offline, and monitors a cluster file system mount point. For additional information: See “CFSSMount agent” on page 629.
oradata_voldg	CVMVolDg	The CVMVolDg agent represents and controls CVM disk groups and the CVM volumes within the disk groups. For additional information: See “CVMVolDg agent” on page 627.

Automatic database starting prevention

The CRS and Oracle agent may attempt to start the instance at the same time if the database mount is available. To prevent automatic database starting, you must

change the Management policy for the database (automatic, manual) to MANUAL using SRVCTL command. The command will change the AUTO_START attribute of CRS database and instance resources.

To prevent automatic database starting

- ◆ Depending on the status of the database:
 - If the database is already registered and running:
 To change it to manual, execute the following commands as an Oracle user:

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

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

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

Configuring VCS resources for Oracle

VCS resources for Oracle can be configured using the following methods:

- By editing the VCS configuration file, main.cf, to define the service groups:
 See [“Editing the main.cf file to configure the CVM and Oracle service groups”](#) on page 260.
- By configuring the CVM and Oracle service groups using the CLI:
 See [“Configuring the CVM and Oracle service groups using the CLI”](#) on page 262.

Editing the main.cf file to configure the CVM and Oracle service groups

This section describes how to manually edit the `main.cf` file to configure the VCS and Oracle service groups.

To configure VCS service groups for Oracle RAC 10g

- 1 Log in to one system as `root`.
- 2 Save your existing configuration to prevent any changes while you modify `main.cf`:

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

If the configuration is not writable, a warning appears: "Cluster not writable." You may safely ignore the warning.

- 3 Make sure VCS is not running while you edit `main.cf` by using the `hastop` command to stop the VCS engine on all systems and leave the resources available:

```
# hastop -all -force
```

- 4 Make a backup copy of the `main.cf` file:

```
# cd /etc/VRTSvcs/conf/config
```

```
# cp main.cf main.orig
```

- 5 Using `vi` or another text editor, edit the `main.cf` file, modifying the `vcs` service groups and creating `Oracle` service groups. Use the following `main.cf` samples as a guideline:

See [“Sample main.cf for Oracle RAC 10g with Oracle agent”](#) on page 511.

See [“Sample main.cf for Oracle RAC 10g without Oracle agent”](#) on page 508.

Modifying the VCS configuration

For additional information and instructions on modifying the VCS configuration by editing the `main.cf` file, refer to the *Veritas Cluster Server User's Guide*.

Saving and checking the configuration

After you finish configuring the CVM and Oracle service groups by editing the `main.cf` file, verify the new configuration.

To save and check the configuration

- 1 Save and close the main.cf file.
- 2 Verify that the syntax of the file `/etc/VRTSvcs/conf/config/main.cf`:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

- 3 Use the following command to start the VCS engine on one system:

```
# hstart
```

- 4 Enter the following command:

```
# hstatus
```

- 5 When "LOCAL_BUILD" is listed in the message column, start VCS on the other system with the following command:

```
# hstart
```

- 6 Verify that the service group resources are brought online. On one system, enter the following command:

```
# hagr -display
```

To verify the state of newly added resources

- 1 Use `hagr -state` to check status of groups.
- 2 Use `hares -state` to check status of resources.

A restart is now required to make sure that CRS and Oracle RAC database instances use Symantec libraries.

Configuring the CVM and Oracle service groups using the CLI

For detailed information and instructions on modifying the VCS configuration by using the command line interface, refer to the *VERITAS Cluster Server User's Guide*.

Location of VCS log files

On all cluster nodes, look at the following log files for any errors or status messages:

```
/var/VRTSvcs/log/engine_A.log
```

When large amounts of data are written, multiple log files may be required, such as engine_B.log, engine_C.log, and so on. The engine_A.log contains the most recent data.

Upgrading to Oracle 10g and migrating the database

This chapter includes the following topics:

- [About upgrading and migrating to Oracle RAC 10g software](#)
- [Overview of upgrade and migration tasks](#)
- [Migrating from Oracle RAC 9i R2 to Oracle RAC 10g](#)
- [Migrating from Oracle RAC 10gR1 to Oracle RAC 10gR2](#)

About upgrading and migrating to Oracle RAC 10g software

SF Oracle RAC 5.0 MP3 supports the following Oracle RAC software upgrades:

- Oracle RAC 9iR2 to Oracle RAC 10gR2
- Oracle RAC 10gR1 to Oracle RAC 10gR2

Overview of upgrade and migration tasks

If you have not already installed and configured Veritas Storage Foundation for Oracle RAC 5.0 MP3, you will need to do so during the preparation for migration procedures.

The upgrade and migration procedures are specific for the OS, current version of Oracle RAC, and target version of Oracle RAC you are using. The high-level order of procedures:

- Upgrade OS if required

See [“Supported operating systems”](#) on page 39.

- Upgrade to Veritas Storage Foundation for Oracle RAC 5.0 MP3 if you have not already upgraded

See [“Supported upgrade options for SF Oracle RAC 5.0 MP3”](#) on page 145.

- Install target version of Oracle RAC, including patchsets if required
- Migrate the database
- Relink the Symantec libraries

Migrating from Oracle RAC 9i R2 to Oracle RAC 10g

The migration procedure assumes that the beginning configuration includes the following components, and that these components are up and running on the cluster nodes:

- Supported operating system
- SF Oracle RAC 5.0 MP3
- Oracle RAC 9i R2

The migration procedure consists of the following tasks:

- [Performing pre-upgrade tasks](#)
- [Installing Oracle RAC 10g](#)
- [Migrating an existing Oracle RAC 9i R2 database to Oracle RAC 10g](#)
- [Performing post-upgrade tasks](#)

Performing pre-upgrade tasks

The migration procedures assume that you have installed Oracle RAC 9i on your cluster nodes.

Before you upgrade and migrate the Oracle RAC 9i software, do the following:

- Review the configuration of VCS service groups.
See [“About VCS service group for Oracle RAC 10g dependencies”](#) on page 253.
- Review the example main.cf file.
See [“About sample main.cf files”](#) on page 507.
- Configure Oracle RAC 10g prerequisites.
See [“About preparing to install Oracle RAC 10g”](#) on page 195.

Perform the following pre-upgrade tasks before migrating from Oracle RAC 9i R2 to Oracle RAC 10g.

To perform pre-upgrade tasks

- 1 Upgrade OS and install any patches, if required. For details:
See Oracle metalink documentation
- 2 Take a hot or cold backup of the existing Oracle RAC 9iR2 database.
- 3 Take a backup of the existing Oracle Home and Central Inventory.
- 4 Freeze the VCS service groups that contain resources to monitor Oracle RAC database instances and/or listener processes. As root user, enter:

```
# haconf -makerw

# hagr -freeze oracle_group -persistent

# haconf -dump -makero
```
- 5 Shutdown Oracle RAC database instances and listener processes on each node in the cluster.
- 6 Stop the Global Services Daemon, GSD. As the Oracle user, enter:

```
$ $ORACLE_HOME/bin/gsdctl stop
```
- 7 Rename the SRVM configuration location file. As the Oracle user, enter:

```
$ cd /var/opt/oracle

$ mv srvConfig.loc srvConfig.loc_backup
```
- 8 Configure Private IP addresses for CRS communication.
See [“Configuring private IP addresses for CRS”](#) on page 210.
- 9 Create home directory for CRS binaries and a new home directory for Oracle RAC 10g R2 binaries.
See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 218.
- 10 Create diskgroup, volume and filesystem for OCR and vote disk.
See [“Creating OCR and VOTE-disk volumes”](#) on page 221.
- 11 Add a CFSMount and CVMVolDg resource to VCS configuration for monitoring the shared diskgroup and filesystem.

Installing Oracle RAC 10g

Install the Oracle RAC 10g binaries.

- Install Clusterware for Oracle RAC 10g.
See [“Installing Oracle Clusterware \(CRS\)”](#) on page 233.
- Install Oracle RAC 10g binaries.
See [“Installing Oracle RAC 10g database binaries”](#) on page 238.
- Add required Oracle RAC patches or patchsets.
See [“Adding Oracle RAC 10g patchsets”](#) on page 244.
- Complete post-installation tasks.
See [“Completing Oracle RAC 10g post-installation tasks”](#) on page 244.

To install Oracle RAC 10g

- 1 Install Oracle RAC 10g Clusterware software.
See [“Installing Oracle Clusterware \(CRS\)”](#) on page 233.
- 2 Verify that the CRS daemons are running. As user Oracle, enter:

```
$ $CRS_HOME/bin/crs_stat -t
```
- 3 If `crs_stat -t` does not show VIP objects, run VIPCA to create VIP objects.
- 4 Add the "cssd" resource (for monitoring the CRS daemons) to VCS configuration.
- 5 Install Oracle RAC 10g RDBMS software. For detailed instructions:
See [“Installing Oracle RAC 10g database binaries”](#) on page 238.
- 6 Install the latest patchset. For detailed instructions:
See [“Adding Oracle RAC 10g patchsets”](#) on page 244.
- 7 Complete post-installation tasks.
See [“Completing Oracle RAC 10g post-installation tasks”](#) on page 244.

Migrating an existing Oracle RAC 9i R2 database to Oracle RAC 10g

To migrate an existing database from Oracle RAC 9i R2 database to Oracle RAC 10g, upgrade the database to Oracle RAC 10g.

For details, see Oracle metalink documentation.

Performing post-upgrade tasks

To perform post-upgrade tasks

- 1 Unfreeze the VCS service groups that were frozen earlier.

See [“Performing pre-upgrade tasks”](#) on page 266.

As root user, enter:

```
# haconf -makerw

# hagrps -unfreeze oracle_group -persistent

# haconf -dump -makero
```

- 2 Relink Symantec libraries. For details:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

Migrating from Oracle RAC 10gR1 to Oracle RAC 10gR2

The migration procedure assumes that the following configuration is up and running on the cluster nodes:

- Supported operating system
- SF Oracle RAC 5.0 MP3
- Oracle RAC 10g R1

Tasks for Migration:

- [Performing pre-upgrade tasks](#)
- [Installing Oracle RAC 10gR2](#)
- [Performing post-upgrade tasks](#)

After upgrading the database, make sure Oracle RAC is linked to Veritas libraries for Oracle RAC 10g R2.

Performing pre-upgrade tasks

Complete the pre-upgrade tasks before migrating from Oracle RAC 10gR1 to Oracle RAC 10gR2.

To perform pre-upgrade tasks

- 1 Upgrade the OS and install any patches, if required.
For details, see Oracle metalink documentation.
- 2 Take a hot or cold backup of the existing database
- 3 Take a backup of the existing Oracle home and central inventory.
- 4 Shutdown the Oracle RAC instance.

- If Oracle RAC is under VCS control, freeze the Oracle group:

```
# haconf -makerw

# hagrps -freeze oracle_group -persistent

# haconf -dump -makero
```

- Use Oracle commands to shutdown Oracle.

- 5 Shut down CRS.

- If CRS is under VCS

```
# haconf -makerw

# hagrps -freeze cssd_group -persistent

# haconf -dump -makero
```

- Stop CRS.

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

Installing Oracle RAC 10gR2

After completing the pre-upgrade tasks, complete the upgrade procedure.

To upgrade from Oracle RAC 10g R1 to Oracle RAC 10gR2

- 1 Run pre-update.

```
# /cdrom/CRS/upgrade/preupdate.sh
```

- 2 Install the Oracle RAC 10gR2 CRS.

See [“Installing Oracle Clusterware \(CRS\)”](#) on page 233.

- 3 As root user run rootupgrade as directed by CRS installer
`# $ORA_CRS_HOME/install/rootupgrade`
- 4 Make sure 10gR2 CRS is running.
 - To list the version of CRS software installed
`# $ORA_CRS_HOME/bin/crsctl query crs softwareversion`
 - To list the Oracle RAC CRS software operating version:
`# $ORA_CRS_HOME/bin/crsctl query crs activeversion`
- 5 Install the Oracle RAC 10g R2 RDBMS.
 See [“Installing Oracle RAC 10g database binaries”](#) on page 238.

Migrating the existing Oracle RAC 10g R1 database to Oracle RAC 10g R2

Upgrade the database to Oracle RAC 10g R2.
 For details, see Oracle metalink documentation.

Performing post-upgrade tasks

To perform post-upgrade tasks

- 1 Unfreeze the VCS service groups that were frozen earlier.
 See [“Performing pre-upgrade tasks”](#) on page 269.
 As root user, enter:
`# haconf -makerw`
`# hagr -unfreeze cssd_group -persistent`
`# hagr -unfreeze oracle_group -persistent`
`# haconf -dump -makero`
- 2 Relink Symantec libraries. For details:
 See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

Oracle RAC 11g installation in an SF Oracle RAC environment

- [Preparing to install Oracle RAC 11g](#)
- [Installing Oracle RAC 11g](#)
- [Installing Oracle RAC 11g manually](#)
- [Configuring VCS service groups for Oracle RAC 11g](#)
- [Upgrading to Oracle RAC 11g and migrating the database](#)

Preparing to install Oracle RAC 11g

This chapter includes the following topics:

- [About preparing to install Oracle RAC 11g](#)
- [Performing pre-installation operations](#)


About preparing to install Oracle RAC 11g

After setting up Veritas Storage Foundation for Oracle RAC, prepare to install the Oracle RAC 11g software. You can install the Oracle RAC 11g software on shared storage or locally on each node.

Review your Oracle installation manuals and the appropriate Oracle support websites, before installing Oracle RAC 11g.

[Figure 13-1](#) displays the steps required to prepare for an Oracle RAC 11g installation. Each step is discussed in detail in this chapter.

Figure 13-1 Preparing to install Oracle RAC 11g flowchart



Performing pre-installation operations

Performing pre-installation operations involves both manual and automated tasks from the SF Oracle RAC configuration program. Before installing Oracle RAC 11g, the following tasks must be performed:

- Setting the system parameters
- Identify storage for the Oracle RAC 11g components
- Launching the SF Oracle RAC configuration program and menu
 - Creating Oracle user and groups
 - Configuring CRS and UDP private IP addresses for failover
- Setting Oracle Clusterware oracle user capabilities
- Configuring shell limits
- Configuring the user process parameter
- Configuring Networks Tuning Parameters
- Tuning the Virtual Memory Manager (VMM)
- Verifying the user "nobody" exists
- Identifying the public virtual IP addresses for use by Oracle
- Creating Oracle disk groups, volumes, and mount points
 - Preparing \$CRS_HOME and \$ORACLE_HOME on each node
 - Creating OCR and VOTE-disk volumes and file systems
- Verifying SF Oracle RAC resources on all nodes
- Verifying GAB port membership

Note: When running the SF Oracle RAC configuration program, you can perform all preinstallation tasks or choose specific tasks to perform. You can also perform preinstallation tasks sequentially at one time, or you can access the program to perform individual tasks at a later time.

Note: Several preinstallation tasks must be manually performed.

Setting the system parameters

Set the system parameters, including the OS kernel and networking parameters.

Refer to the latest Oracle documentation for information about the system parameters and procedures. Restart the nodes for the new values to take effect.

Identify storage for the Oracle RAC 11g components

The SF Oracle RAC configuration program provides options to perform the following Oracle RAC 11g installation tasks:

- Oracle RAC 11g CRS installation
- Oracle RAC 11g database installation
- Relinking SF Oracle RAC libraries to Oracle RAC 11g

The SF Oracle RAC configuration program prompts you for information while installing Oracle RAC 11g in an SF Oracle RAC environment. The program provides default values for some information, which you can accept.

Note: The examples in this chapter assume a two-node cluster, with the nodes galaxy and nebula.

The following information is required for the configuration:

- Information to create Oracle user and Inventory group as primary group:
[Table 13-1](#)
- Information to create OS DBA as secondary group:
[Table 13-2](#)
- Information to configure private IP addresses for CRS:
[Table 13-4](#)
- Information for \$CRS_HOME on each system:
[Table 13-5](#)
- Information to create OCR and Voting Disk directories on CFS:
[Table 13-6](#)

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

Launching the SF Oracle RAC configuration program and menu

The following procedure describes the Oracle RAC 11g pre-installation tasks using the SF Oracle RAC configuration program and menu.

To perform pre-installation tasks in the configuration menu

- 1 If the SF Oracle RAC configuration program is not currently running, launch the SF Oracle RAC configuration program.

As root user on any one system, enter:

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

By default, the installsfrac program uses SSH for remote communication.

To use RSH, specify the -rsh option with the installsfrac program.

- 2 If you have not previously installed SF Oracle RAC, then enter the system names, separated by spaces.
- 3 The SF Oracle RAC configuration program displays the Symantec copyright information, as well as the location of the installsfrac logs.

Note: If necessary, the installsfrac logs can be used for troubleshooting purposes.

- 4 You will be prompted whether you wish to configure SFRAC from these systems.

```
Do you want to configure SFRAC from these systems? [y,n,q] (y)
```

- 5 The SF Oracle RAC configuration program then performs an initial system check of the SSH or RSH communication, SF Oracle RAC installation, and architecture on the nodes.
- 6 The SF Oracle RAC configuration program then checks your system license on the nodes. You will be asked whether you want to enter another license key on each system.

If not, then enter "n" to proceed.

7 The SF Oracle RAC configuration program then displays the following information about its INSTRUCTIONS ONLY steps:

- * Do not perform any action
- * Require you to understand the instructions and act
- * May require you to refer to the documentation

Press enter to proceed.

The following additional information is provided:

- * Keep root shells open on all systems to run INSTRUCTIONS ONLY steps
- * Follow the steps in order
- * Navigation keys for all screens: q to quit, b for previous menu

8 The configuration program menu is then displayed:

- 1) Check systems for SFRAC ** INSTRUCTIONS ONLY **
- 2) Configure SFRAC
- 3) Prepare to install Oracle
- 4) Install or relink Oracle

From this menu, you can choose to perform all installation tasks, or to perform specific tasks.

From the configuration program menu, select **Prepare to install Oracle**.

9 After selecting the **Prepare to install Oracle** option, the following menu is presented:

This step helps you complete the sysadmin tasks before installing Oracle software.

- 1) [Perform all the following tasks]
- 2) Create userid and groupid for Oracle
- 3) Configure private IP addresses for CRS (Oracle 10g only)
- 4) Configure CRS and UDP private address for failover (Oracle 11g)
- 5) Create disk groups, volumes, and files systems for Oracle **INSTRUCTIONS ONLY**
- b) [Go to previous menu]

Proceed to select the appropriate options for your installation and follow the steps to pre-install Oracle RAC 11g.

For detailed information about option 2 (Create userid and groupid for Oracle RAC 11g):

See [“Creating Oracle user and groups”](#) on page 281.

For detailed information about option 4 (Configure CRS and UDP private address for failover (Oracle RAC 11g):

See [“Configuring CRS and UDP private IP addresses for failover”](#) on page 291.

For detailed information about option 5 (Create disk groups, volumes, and file systems for Oracle):

See [“Creating Oracle disk groups, volumes, and mount points”](#) on page 302.

Creating Oracle user and groups

You must create the Oracle groups oinstall (the Oracle Inventory group) and dba. You must create the Oracle user oracle. You must assign the primary group for Oracle to be oinstall and the secondary group for Oracle to be dba.

Refer to the latest Oracle documentation for additional information.

Table 13-1 displays examples of the information required to create an Oracle user and inventory group as the primary group.

Table 13-1 Required information to create Oracle user and Inventory group as primary group

Required information	Example
Oracle user name	oracle Default is oracle.
Oracle user id	100
Oracle user home directory	/home/oracle
Oracle group name	oinstall Default is oinstall.
Oracle group id	1000

Table 13-2 displays examples of the information required to create OS DBA group as the secondary group.

Table 13-2 Required information to create OS DBA group as secondary group

Required information	Example
Oracle group name	dba Default is dba.
Oracle group id	1001

You can create Oracle users and groups using the SF Oracle RAC configuration program. As discussed above, refer to the Oracle documentation for information about creating the oinstall (Oracle inventory), dba group, and oracle user. After reviewing the Oracle documentation, proceed to create these Oracle groups and user.

Note: Make sure that the user and group IDs are unused on all the nodes in the cluster.

To create Oracle user and groups

- 1 Access the SF Oracle RAC configuration program if you are not currently using it.

See [“Launching the SF Oracle RAC configuration program and menu”](#) on page 278.

- 2 Select **Create userid and groupid for Oracle** from the configuration menu.
- 3 Provide the required information for the configuration program.

The first step in this configuration process creates the user and group ids on all cluster nodes for Oracle.

The configuration program assigns the same values on all nodes. The text within parentheses is the default, which you may select by pressing Enter.

```
Enter Oracle UNIX user name: [b] (oracle)
Enter Oracle UNIX user id (numerical): [b] 100
Enter Oracle UNIX user home dir: [b] /home/oracle
Enter Oracle UNIX group name: [b] (oinstall)
Enter Oracle UNIX group id: [b] 1000
User oracle does not exist on any node. Do you want to create it
with the information provided [y,n,q] (y)
```

Note: The set of Oracle user IDs and group IDs in each cluster configuration must be the same.

- 4 The installer verifies that the specified `userid` does not exist on any of the systems in the cluster and then creates it. Enter `y` to create the oracle user with the information provided.

5 Next, enter the information to create secondary group, "dba".

The SF Oracle RAC configuration program displays the following message:

```
This step creates secondary group ids on all cluster nodes.
Oracle requires some other special groups for identifying
operating system accounts that have database admin (SYSDBA)
privileges and for those accounts that have limited sets
of database admin (SYSOPER) privileges. You should create
such groups on all systems and Oracle user should be part
of these groups.
```

The following options are then presented. Enter the appropriate values for your installation or accept the default values.

```
Do you want to create another secondary group for Oracle user? [y,n,q] (y)
```

```
Enter Oracle UNIX secondary group name: [b] dba
```

```
Enter Oracle UNIX secondary group id (numerical): [b] 1001
```

```
Group dba does not exist on any node. Do you want to create it
with the information provided [y,n,q] (y) y
```

```
Creating group dba on galaxy ..... Done
Adding Oracle user (oracle) to group (dba) on galaxy .. Done
Creating group dba on nebula ..... Done
Adding Oracle user (oracle) to group (dba) on nebula .. Done
```

- 6 Next, you are prompted to perform SSH or RSH verification. Enter "y" to perform verification. Review as the configuration program verifies access for Oracle user on all the nodes. The Oracle installation process requires SSH or RSH permission to be set for the Oracle user. If the SSH or RSH verification fails on any nodes, enable SSH or RSH access for those nodes. To enable RSH access for the nodes, leave the installer at this prompt and proceed to set up access. Open a new session to set up access. You can return to this session after setting up oracle user equivalence.

See [“Setting up Oracle user equivalence for RSH and RCP”](#) on page 286.

- 7** Edit the Oracle user .profile file to enable paths to \$CRS_HOME, \$ORACLE_HOME, and \$ORACLE_BASE on each node.

On each node, set the proper environment variables (the following examples are for ksh):

```
$ cat .profile
set -o vi

PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:\
/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]          # This is at Shell startup.  In normal
then echo "$MAILMSG"      # operation, the Shell checks
fi                        # periodically.

ORACLE_HOME=/app/oracle/orahome; export ORACLE_HOME
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data; export ORA_NLS33
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib; export CLASSPATH
DISPLAY=YOUR.DISPLAY.IP; export DISPLAY
CRS_HOME=/app/crshome; export CRS_HOME
ORA_CRS_HOME=/app/crshome; export ORA_CRS_HOME

LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$CRS_HOME/lib:\
$CRS_HOME/lib32:/usr/lib:/opt/VRTSdbms3/lib; export LIBPATH
PATH=$PATH:$ORACLE_HOME/bin:$CRS_HOME/bin:/opt/VRTS/bin:\
/opt/VRTSdbed/bin; export PATH

ORACLE_SID=PROD; export ORACLE_SID

umask 022

set -o vi
$
```

Setting up Oracle user equivalence for RSH and RCP

This section discusses how to set up Oracle user equivalence for RSH and RCP.

To set up Oracle user equivalence for RSH and RCP

- 1 Open another terminal session as `oracle` user.
- 2 As `oracle` user on each system, edit `$HOME/.rhosts` file and add entries similar to the following:

```
galaxy oracle
nebula oracle
```

- 3 On each system, set the password for the "oracle" user:

```
[root@galaxy /]# passwd oracle

Changing password for "oracle"
oracle's New password:
Re-enter oracle's new password:
```

- 4 On each system, login as user "oracle" and change the password.

```
[root@galaxy /]# su - oracle

$ passwd

Changing password for "oracle"
oracle's New password:
Re-enter oracle's new password:
```

- 5 On each system, as user "oracle", verify "rsh" access:

```
$ rsh galaxy date

Mon Apr 24 10:02:45 PDT 2006

$ rsh nebula date

Mon Apr 24 10:02:45 PDT 2006
```

You can now create the secondary groups for Oracle.

See [“Creating Oracle user and groups”](#) on page 281.

To verify RSH access for "oracle" user

- 1 Return to the installer session.
See [“Creating Oracle user and groups”](#) on page 281.
- 2 At the installer prompt, answer "y" to verify "RSH" accessibility.
- 3 Quit the installation program.

Setting Oracle Clusterware oracle user capabilities

Ensure that the Oracle Clusterware oracle user has the following capabilities:

CAP_NUMA_ATTACH, CAP_BYPASS_RAC_VMM, and CAP_PROPAGATE.

Setting Oracle Clusterware oracle user capabilities

- 1 To check existing capabilities, enter the following command as `root`; in this example, the Oracle Clusterware oracle user is `oracle`:

```
# /usr/sbin/luser -a capabilities oracle
```

- 2 To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser \  
capabilities=CAP_NUMA_ATTACH,CAP_BYPASS_RAC_VMM,CAP_PROPAGATE oracle
```

- 3 The `ncargs` attribute controls the maximum command line length, including the environment variables on the AIX nodes.

Change this attribute to a larger value before invoking the Oracle Installer.

The following is an example of the command to set this attribute:

```
# /usr/sbin/chdev -l sys0 -a ncargs=128
```

Configuring shell limits

Oracle recommends setting shell limits. Refer to your Oracle manual and the Oracle documentation website for the latest information about setting shell limits before proceeding with the following procedure.

[Table 13-3](#) displays the shell limits you must increase to improve the performance of the software.

Table 13-3 Shell limits

Shell limit	Item in limit.conf	Hard limit
Maximum number of open file descriptors	nofile	65536
Maximum number of processes available to a single user	maxuproc	16384

To configure shell Limits

- 1 To improve the performance of the software, you must increase the following shell limits.

Add the following lines to the `/etc/security/limits` file:

```
default:
fsize = -1
core = -1
cpu = -1
data = 512000
rss = 512000
stack = 512000
nofiles = 2000
```

- 2 Enter the following command to list the current setting for the maximum number of process allowed by the Oracle software user:

```
/usr/sbin/lsattr -E -l sys0 -a maxuproc
```

If necessary, change the maxuproc setting using the following command:

```
/usr/sbin/chdev -l sys0 -a maxuproc=16384
```

- 3 Repeat this procedure on all other nodes in the cluster.

Configuring the user process parameter

Review the latest Oracle documentation for information about all OS kernel parameters that need to be configured and the required values. The following procedure is used to configure the maximum processes per user.

To configure maximum processes per user

- 1 Check the current value by entering the command:

```
# smitty chgsys
```

- 2 Examine the value of the kernel parameter:

```
Maximum number of PROCESSES allowed per user
```

If necessary, change the value per your Oracle documentation. The value must be the same on each of the cluster nodes.

- 3 Press F10 to exit when you are done.
- 4 Use the previous steps to check that the value is the same on each cluster node.

Configuring Networks Tuning Parameters

Perform the following tasks to configure the networks tuning parameters.

Configure the Networks Tuning Parameters

- 1 Use the AIX System Management Interface Tool (SMIT) to access the following selections:
 - smitty> Performance & Resource Scheduling> Tuning Kernel & Network
 - smitty> Parameters> Tuning Network Option Parameters> Change/Show Parameters for Next Boot
- 2 Identify each of the following parameters and set it with the value listed below:

```
udp_sendspace=65536
udp_recvspace=655360
tcp_sendspace=65536
tcp_recvspace=65536
rfc1323=1
sb_max=1310720
ipqmaxlen=512
```

Tuning the Virtual Memory Manager (VMM)

Perform the following tasks to tune the Virtual Memory Manager (VMM).

Tuning the Virtual Memory Manager (VMM)

- 1 Use the AIX System Management Interface Tool (SMIT) to access the following selections:
 - smitty > Performance & Resource Scheduling> Tuning Kernel & Network
 - smitty> Parameters> Tuning Virtual Memory Manager, File System and Logical Volume Manager Params> Change/Show Parameters for Next Boot
- 2 Identify each of the following parameters and set it with the value listed below:

```
lru_file_repage=0
minperm%=3
maxperm%=90
maxclient%=90
strict_maxclient=1
strict_maxperm=0
```

Verifying the user "nobody" exists

Verify the user "nobody" exists on each system in the cluster:

```
#id nobody
uid=4294967294 (nobody) gid=4294967294 (nobody)
```

Note: Make sure that the uid and gid are the same across the nodes in your cluster.

Configuring CRS and UDP private IP addresses for failover

Oracle CRS requires an IP address on each node for heartbeat. Oracle RDBMS requires an IP address for sending cache-fusion data through UDP protocol.

In the SF Oracle RAC cluster, there are redundant private links for sending llt communication and heartbeat. Symantec recommends that all CRS and UDP cache-fusion links be llt links. Having a CRS heartbeat and llt heartbeats on the same set of links provides a single view of the cluster to both VCS and CRS. Having a different set of links causes unnecessary downtime and side-effects.

Configuring the multiple private links for UDP cache-fusion provides additional bandwidth. The MultiPrivNIC agent relies on llt to manage all of the private IP

addresses that are needed for CRS and UDP cache-fusion. This agent will provide failover capability for these IP addresses in case of link failures.

You can configure CRS and UDP private addresses for failover using either the SFRAC for Oracle installer or manually. To configure CRS and UDP private addresses for failover using the SFRAC for Oracle installer, perform the tasks described in the following procedures:

- Configure the private IP addresses for CRS and UDP cache-fusion using the SF Oracle RAC configuration menu.
See [“To configure private IP addresses for CRS and UDP cache-fusion”](#) on page 294.
- Add the private IP addresses to the `/etc/hosts` file.
See [“To add private IP addresses to the `/etc/hosts` file”](#) on page 297.
- Verify the private IP addresses.
See [“To verify the private IP addresses”](#) on page 298.

To configure the private IP addresses for CRS using the SF Oracle RAC configuration menu, the following procedures must be followed:

- To proceed to create a private IP address on each node, first determine the private NIC device for which LLT is configured. If a private NIC device is used as an LLT interconnect on one system, you can configure an available IP address for it.
- Access the file `/etc/llttab` on each node and review the private NIC device.

To configure CRS and UDP private addresses for failover manually, perform the tasks described in the following section:

See [“Manually configuring MultiPrivNIC for UDP cache fusion”](#) on page 299.

Table 13-4 Required information to configure the private IP addresses for CRS (examples)

Required information	Example
Private network interfaces for each node	You have to choose an LLT device as a device for the CRS heartbeat. For example, if the LLT devices are en1, en2: ■ On galaxy: en1, en2 ■ On nebula: en1, en2 Then the MultiPrivNIC device names will be as follows: Device@galaxy= { en1 = 0, en2 = 1 } Device@nebula= { en1 = 0, en2 = 1 } If aggregated device names are configured under LLT, then the aggregated names must be used in MultiPrivNIC agent. For example, if LLT device names are: ■ On galaxy: aggr1 ■ On nebula: aggr1 Then the Device Attribute for the MultiPrivNIC agent would be as follows: Device@galaxy = { aggr1 = 0 } Device@nebula = { aggr1 = 0 }

Table 13-4 Required information to configure the private IP addresses for CRS (examples) *(continued)*

Required information	Example
Private IP addresses for each node	On galaxy: ■ NIC 1 (en1): 192.168.1.1 (CRS) ■ NIC 1 (en1): 192.168.1.2 (UDP 1 to use for cluster_interconnects) ■ NIC 2 (en2): 192.168.2.1 (UDP 2 to use for cluster_interconnects) On nebula: ■ NIC 1 (en1): 192.168.1.3 (CRS) ■ NIC 1 (en1): 192.168.1.4 (UDP 1 to use for cluster_interconnects) ■ NIC 2 (en2): 192.168.2.2 (UDP 2 to use for cluster_interconnects) One private IP address is to be used for CRS and the other two private IP addresses are to be used for cluster_interconnects (UDP). To avoid routing issues, ensure that the private IP addresses assigned to the first NIC (CRS and UDP 1) are using the same subnet. Additionally, ensure that the third IP address (UDP 2) for the second NIC is on a different subnet than the first NIC.
Netmask for the cluster	255.255.255.0

Symantec recommends adding all llt private interfaces to the MultiPrivNIC resource, as described in the following procedure.

After confirming the values, the SF Oracle RAC configuration program will add a new section in the VCS configuration file (main.cf) for the MultiPrivNIC resource in the CVM group.

To configure private IP addresses for CRS and UDP cache-fusion

- 1 Access the SF Oracle RAC configuration program if you are not currently using it.

See [“Launching the SF Oracle RAC configuration program and menu”](#) on page 278.
- 2 Select option 4, **Configure CRS and UDP private IP addresses for Failover (Oracle 11g)** from the configuration menu.

- 3** After selecting option 4 from the SF Oracle RAC configuration program, the following information is displayed:

This step will configure private IP addresses using MultiPrivNIC resource of VCS.

- * You must use all LLT links for MultiPrivNIC.
 Non-LLT links can not be used
- * IP addresses used here must be manually added to /etc/hosts on all nodes
- * UDP IP addresses must be added to the oracle init file as cluster_interconnect parameter
 This step is NOT required if you will only install Oracle 9i.

For information about the procedure to add UDP IP addresses to the oracle init file:

See [“Adding UDP IP addresses to the Oracle init file”](#) on page 334.

- 4** The SF Oracle RAC configuration program then proceeds to discover and display NIC devices on each of the nodes. You are prompted to select the NIC devices on each node.

```
Discovering NICs on galaxy .... discovered en0 en1 en2 en3
Enter NIC 0 for galaxy (x if done): [b] en1
Enter NIC 1 for galaxy (x if done): [b] en2
Enter NIC 2 for galaxy (x if done): [b] x
```

```
Discovering NICs on nebula .... discovered en0 en1 en2 en3
Enter NIC 0 for nebula (x if done): [b] (en1)
Enter NIC 1 for nebula (x if done): [b] (en2)
Enter NIC 2 for nebula (x if done): [b] x
```

5 Enter the private IP address information for each host. You will be prompted for this information.

```
Enter IP address 1 for galaxy for en1 (x if done): [b] 192.168.1.1
Enter IP address 2 for galaxy for en1 (x if done): [b] 192.168.1.2
Enter IP address 3 for galaxy for en1 (x if done): [b] x
```

```
Enter IP address 1 for galaxy for en2 (x if done): [b] 192.168.2.1
Enter IP address 3 for galaxy for en2 (x if done): [b] x
```

```
Enter IP address 1 for nebula for en1 (x if done): [b] 192.168.1.3
Enter IP address 2 for nebula for en1 (x if done): [b] 192.168.1.4
Enter IP address 3 for nebula for en1 (x if done): [b] x
```

```
Enter IP address 1 for nebula for en2 (x if done): [b] 192.168.2.2
Enter IP address 3 for nebula for en2 (x if done): [b] x
```

6 Enter the netmask for the private network.

```
Enter the netmask for private network: [b] 255.255.255.0
```


- 7** The SF Oracle RAC configuration program then displays the configured parameters for the private IP address information for CRS.

Review and confirm the private IP address information for CRS.

```
Private IP configuration information verification
System: galaxy
      Private Interfaces : en1 en2
      Private IPs on en1: 192.168.1.1 192.168.1.2
      Private IPs on en2: 192.168.2.1
System: nebula
      Private Interfaces : en1 en2
      Private IPs on en1: 192.168.1.3 192.168.1.4
      Private IPs on en2: 192.168.2.2

Netmask: 255.255.255.0

Is this information correct? [y,n,q,b] (y)

      Changing configuration to read-write mode.....success
      Adding new multi_priv resource.....success
      Saving configuration.....success
```

- 8** After confirming the values, the SF Oracle RAC configuration program adds a new section in the VCS configuration file (main.cf) for the MultiPrivNIC resource in the CVM group.

See “[MultiPrivNIC agent](#)” on page 635.

- 9** After the private IP address is configured, you can exit the installer by entering **q** or continue with your configuration.

To add private IP addresses to the /etc/hosts file

- 1** Log in to each system as root.
- 2** Using vi or another text editor, add the following entries to the /etc/hosts file on each system:

```
192.168.1.1    galaxy_priv
192.168.1.3    nebula_priv
```

To verify the private IP addresses

- 1 Enter the `hastatus` command.

Review the `hastatus` command output to ensure that the `multi_priv` resource is installed and running on each node.

- 2 Verify that the "multi_priv" resource is online on all systems in the cluster, by entering the following command:

```
# hares -state multi_priv
```

# Resource	Attribute	System	Value
multi_priv	State	galaxy	ONLINE
multi_priv	State	nebula	ONLINE

- 3 Access and view the VCS `main.cf` located in the following directory:
`/etc/VRTSvcs/conf/config`.

Enter the following command:

```
# more /etc/VRTSvcs/conf/config/main.cf
```

- 4 Verify that the MultiPrivNIC resource, `multi_priv`, shows up in VCS `main.cf`:

```
MultiPrivNIC multi_priv (
    Critical = 0
    Device @galaxy = { en1 = 0, en2 = 1 }
    Device @nebula = { en1 = 0, en2 = 1 }
    Address @galaxy = { "192.168.1.1" =0, "192.168.1.2" =0, "192.168.2.1" =1 }
    Address @nebula = { "192.168.1.3" =0, "192.168.1.4" =0, "192.168.2.2" =1 }
    NetMask = "255.255.255.0"
)
```

- 5 On each system, check the output of "ifconfig"

```
# ifconfig -a
```

- 6 From each system, ping the private IP addresses.

For example, on the nebula system enter the following command:

```
# ping galaxy_priv  
  
# galaxy_priv is alive
```

On the galaxy system, enter the following command:

```
# ping nebula_priv  
  
# nebula_priv is alive
```

Manually configuring MultiPrivNIC for UDP cache fusion

This section discusses the procedure for manually configuring MultiPrivNIC for UDP cache fusion. The Veritas MultiPrivNIC agent provides resources for UDP/IP support for Oracle's cache fusion capabilities.

Note: If the SF Oracle RAC cluster is running, then proceed to use the `installsfrac` program to configure MultiPrivNIC for UDP cache fusion. If the SF Oracle RAC cluster is not running, then use the following manual procedure to configure MultiPrivNIC for UDP cache fusion.

Manual MultiPrivNIC for UDP cache fusion configuration

- 1 Save the existing configuration by typing the following command:

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

- 2 Make a backup copy of the `main.cf` file by typing the following command:

```
# cd /etc/VRTSvcs/conf/config  
  
# cp main.cf main.orig
```

- 3 On each node, use vi or another text editor to edit the main.cf file located in the /etc/VRTSvcs/conf/config directory. Modify the main.cf file by adding MultiPrivNIC to the group cvm as displayed below. Use the excerpts from a sample main.cf below as an example for your configuration.

To view a complete sample main.cf:

See [“Sample main.cf for Oracle RAC 11g with Oracle agent”](#) on page 525.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
.
.
.

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)
.
.
.

MultiPrivNIC multi_priv (
    Critical = 0
    Device@galaxy = {en1 = 0, en2 = 1}
    Device@nebula = {en1 = 0, en2 = 1}
    Address@galaxy = {"192.168.1.1" = 0, "192.168.1.2" = 0, "192.168.2.1" = 1}
    Address@nebula = {"192.168.1.3" = 0, "192.168.1.4" = 0, "192.168.2.2" = 1}
    NetMask = "255.255.255.0"
)
.
.
.
.
```

- 4 Save and close the file.
- 5 Verify the syntax of the file `/etc/VRTSvcs/conf/config/main.cf` by typing the following command:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

Identifying the public virtual IP addresses for use by Oracle

Before starting the Oracle installation, you must identify the public virtual IP addresses for each node. 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 recommends that the public virtual IP name be in this format : hostA-vip

Identify the public virtual IP address for use by Oracle

- 1 Open a second session for this procedure.
- 2 Obtain one virtual IP address per node from your network administrator.

For example, for galaxy:

```
10.182.13.92
```

For example, for nebula:

```
10.182.13.93
```

- 3 On the second session and using vi or another text editor, add an entry for the virtual IP address and virtual public name in the `/etc/hosts` file. Each node must have a separate public virtual IP address.

The following is an example of an entry:

```
10.182.13.92 galaxy-vip
10.182.13.93 nebula-vip
```

This procedure must be performed on all nodes.

- 4 Add the public virtual IP addresses to DNS.

Note: Contact your network administrator for assistance with registering with DNS.

- 5 After following the above steps, the following entries should appear in the /etc/hosts files for all nodes:

```
192.168.12.1      galaxy_priv
10.182.13.92     galaxy-vip
192.168.12.4     nebula_priv
10.182.13.93     nebula-vip
```

Creating Oracle disk groups, volumes, and mount points

To create disk groups, volumes, and mount points for Oracle, review the following guidelines.

Before you install the Oracle Cluster Ready Services (CRS) and Oracle RAC 11g binaries, you must create storage space for these installations. You need to provide storage for the following directories and files:

- The home directories, CRS_HOME for CRS and ORACLE_HOME, for Oracle binaries.
See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 303.
- The CRS files for Oracle Cluster Registry (OCR) and the VOTE-disk. The files can reside in volumes on raw device, or in directories in a cluster file system.
See [“Creating OCR and VOTE-disk volumes and file systems”](#) on page 306.

You will create CVM volumes or a directory on CFS for database file storage later in the installation process:

See [“Creating the Oracle RAC 11g database”](#) on page 334.

Local or a cluster file system on shared storage

Oracle binaries (ORACLE_HOME) can either be installed locally on each system or on a cluster file system on shared storage. Before proceeding with an installation, consider the following information about each type of Oracle binary installation:


- Local installations provide a comfort level using traditional installation methods and the possibility of improved protection against a single point of failure.
- CFS installations provide a single Oracle installation to manage, regardless of number of nodes. This scenario offers a necessary reduction in storage requirements and easy addition of nodes.

Select the location based on your high availability requirements.

Note: Symantec and Oracle generally recommend using local installations.

Figure 13-2 displays the installation process for CRS and the Oracle binaries.

Figure 13-2 CRS and Oracle binaries installation process



Preparing \$CRS_HOME and \$ORACLE_HOME on each node

Follow the directions for the following procedures to prepare \$CRS_HOME and \$ORACLE_HOME on each node.

- Identify the directories required for \$CRS_HOME:
See [Table 13-5](#) on page 304.
- Create a file system on local storage for Oracle/CRS binaries:

See [“To create a file system on local storage for Oracle/CRS binaries \(/app\)”](#) on page 305.

- Prepare \$CRS_HOME on each node:
See [“To create a file system on local storage for Oracle/CRS binaries \(/app\)”](#) on page 305.
- Prepare \$ORACLE_BASE on each node:
See [“To prepare \\$ORACLE_BASE on each node”](#) on page 306.
- Prepare \$ORACLE_HOME on each node"
See [“To prepare \\$ORACLE_HOME on each node”](#) on page 306.

Note: Open a second session to enter the commands required for the following procedures.

Identify the directories required for installing Oracle and CRS software.

Table 13-5 Required information \$CRS_HOME on each system

Required information	Example
Disk for each node that contains the Oracle software, CRS for Oracle, and the CRS binaries	Disk_1
Disk group on each local storage	■ On galaxy: orabindg_galaxy ■ On nebula: orabindg_nebula
Volume for each local disk group	■ On galaxy: orabinvol_galaxy ■ On nebula: orabinvol_nebula
File system on each local volume	■ On galaxy: /dev/vx/rdsk/orabindg_galaxy/orabinvol_galaxy ■ On nebula: /dev/vx/rdsk/orabindg_galaxy/orabinvol_nebula
Mount point for each local file system	/app (non-shared)
Directory to install CRS	/app
CRS home directory	/app/crshome

To create a file system on local storage for Oracle/CRS binaries (/app)

- 1** As root user, first create a VxVM local diskgroup, *orabindg_hostname*:

```
# vxdg init orabindg_galaxy Disk_1
```

- 2** Create a volume, *orabinvol_hostname*:

```
# vxassist -g orabindg_galaxy make orabinvol_galaxy 12G
```

- 3** Create directory, /app

```
# mkdir /app
```

- 4** Create a filesystem with this volume, *orabinvol_hostname*

```
# mkfs -V vxfs /dev/vx/rdisk/orabindg_galaxy/orabinvol_galaxy
```

- 5** Mount /app:

```
# mount -V vxfs /dev/vx/dsk/orabindg_galaxy/orabinvol_galaxy  
/app
```

- 6** Add an entry for this filesystem. For example:

In /etc/filesystems add:

```
/app:  
dev = dev/vx/dsk/orabindg_galaxy/orabinvol_galaxy  
vfs = vxfs  
mount = true  
options = rw  
account = false
```

To prepare \$CRS_HOME on each node

- 1 On each system, log in as "root".
- 2 Create the directory for CRS_HOME:

```
# mkdir -p /app/crshome
```

- 3 Change ownership and permissions:

```
# chown -R oracle:oinstall /app/crshome
```

```
# chmod -R 744 /app/crshome
```

To prepare \$ORACLE_BASE on each node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE_BASE:

```
# mkdir -p /app/oracle
```

- 3 Change ownership and permissions:

```
# chown -R oracle:oinstall /app/oracle
```

```
# chmod -R 744 /app/oracle
```

To prepare \$ORACLE_HOME on each node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE_HOME:

```
# mkdir -p /app/oracle/orahome
```

- 3 Change ownership and permissions:

```
# chown -R oracle:oinstall /app/oracle/orahome
```

```
# chmod -R 744 /app/oracle/orahome
```

Creating OCR and VOTE-disk volumes and file systems

The CRS installation requires predefined locations for the Oracle Cluster Registry (OCR) and VOTE-disk components.

Note: Whether you create volumes or file system directories, you can add them to the VCS configuration to make them highly available.

Review the following information and perform the procedures described below to create OCR and VOTE-disk volumes and file systems:

- Information to create OCR and Voting Disk directories on CFS See [Table 13-6](#) on page 307.
- [To create a filesystem for OCR and VOTE disks \(/ocrvote\) using the CLI, or](#)
- [To create a filesystem for OCR and VOTE disks \(/ocrvote\) by editing the main.cf file](#)

Identify the directories required for installing Oracle and CRS software.

Table 13-6 Information to create OCR and Voting Disk directories on CFS

Required information	Example
Shared disk	HDS0_10, HDS0_21
Shared disk group	ocrvotedg
Volume for shared disk group	ocrvotevol
File system on shared volume	/dev/vx/rdisk/ocrvotedg/ocrvotevol
Mount point for shared file system	/ocrvote
Directory for OCR files	ocr
Directory for Voting Disk files	vote
Local disk on each node	HDS0_11
Disk group for each local disk	■ On galaxy: orabindg_galaxy ■ On nebula: orabindg_nebula
Volume for each local disk group	■ On galaxy: orabinvol_galaxy ■ On nebula: orabinvol_nebula

Table 13-6 Information to create OCR and Voting Disk directories on CFS
(continued)

Required information	Example
File system on each local volume	■ On galaxy: /dev/vx/rdisk/orabindg_galaxy/ or abinvol_galaxy ■ On nebula: /dev/vx/rdisk/orabindg_nebula/ or abinvol_nebula
Directory for Oracle base (ORACLE_BASE)	/app/oracle
Directory to install Oracle (ORACLE_HOME):	/app/oracle/orahome

The following are the two procedures that can be used to create a filesystem for the OCR and VOTE disks (/ocrvote). One procedure uses the CLI and the other procedure uses a process where the main.cf file is edited.

Use only one of the two procedures to perform this task to create a filesystem.

To create a filesystem for OCR and VOTE disks (/ocrvote) using the CLI

- 1
- As root user, from the CVM master, create a shared VxVM diskgroup by issuing the following command:

```
# vxpdg -s init ocrvotedg HDS0_10 HDS0_21
```

Note: To determine the CVM master, use the vxctl - c mode command.

- 2
- As root user, from the CVM master, create a mirrored volume, ocrvotevol:

```
# vxassist -g ocrvotedg make ocrvotevol 1G
```

- 3
- As root user, from CVM master, create a filesystem with the volume, ocrvotevol.

```
# mkfs -V vxfs /dev/vx/rdisk/ocrvotedg/ocrvotevol
```

- 4 On each system, create a directory, /ocrvote:

```
# mkdir /ocrvote
```

- 5 On each system, mount /ocrvote

```
# mount -V vxfs -o cluster /dev/vx/dsk/ocrvotedg/ocrvotevol  
/ocrvote
```

- 6 As root user, from any system, change permissions on /ocrvote

```
# chown -R oracle:oinstall /ocrvote
```

- 7 The following steps add the CFSMount and CVMVolDg resources corresponding to this "/ocrvote" filesystem to the "cvm" group in VCS configuration. Before proceeding with these steps, enter the following command to change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 8 Enter the following commands to set up the file system:

```
# hares -add ocrvote_mnt CFSMount cvm  
  
# hares -modify ocrvote_mnt Critical 0  
  
# hares -modify ocrvote_mnt MountPoint "/ocrvote"  
  
# hares -modify ocrvote_mnt BlockDevice \  
"/dev/vx/dsk/ocrvotedg/ocrvotevol"  
  
# hares -modify ocrvote_mnt MountOpt mincache=direct
```

- 9 Enter the following commands to set up the volumes:

```
# hares -add ocrvote_voldg CVMVolDg cvm  
  
# hares -modify ocrvote_voldg CVMDiskGroup ocrvotedg  
  
# hares -modify ocrvote_voldg CVMVolume -add ocrvotevol  
  
# hares -modify ocrvote_voldg CVMActivation sw
```

10 Enter the following commands to link the parent and child:

```
# hares -link ocrvote_mnt ocrvote_voldg

# hares -link ocrvote_mnt vxfsckd

# hares -link ocrvote_voldg cvm_clus
```

11 Enter the following commands to begin enabling the system:

```
# hares -modify ocrvote_voldg Enabled 1

# hares -modify ocrvote_mnt Enabled 1

# haconf -dump -makero
```

12 Verify the stanzas corresponding to the ocrvote_mnt and ocrvote_voldg resources in main.cf:

```
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
)

ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
```

- 13 Verify that the VCS resources, ocrvote_mnt, and ocrvote_voldg, are ONLINE on all systems in the cluster.

```
# hares -state ocrvote_mnt  
  
# hares -state ocrvote_voldg
```

- 14 If the voldg is manually onlined, the CVMVolDg resources under VCS will not report as online. VCS must be stopped and restarted.

The following procedure can be used to create a filesystem for OCR and VOTE disks (/ocrvote) by editing the main.cf file.

Note: The following procedure requires a restart on all the nodes.

To create a filesystem for OCR and VOTE disks (/ocrvote) by editing the main.cf file

- 1 Log in to one system as root.
- 2 Save your existing configuration to prevent any changes while you modify main.cf:

```
# haconf -dump makero
```

If your configuration is not writable, a warning appears: "Cluster not writable". You may safely ignore the warning.

- 3 Make sure that VCS is not running while you edit the main.cf by using the `hastop` command.

This command stops the VCS engine on all systems and leaves the resources available.

```
# hastop -all -force
```

- 4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config  
  
# cp main.cf main.orig
```

- 5 Using vi or another text editor, edit the main.cf file, modifying the cvm service group.

Specifically, edit the stanzas corresponding to the ocrvote_mnt and ocrvote_voldg resources in main.cf:

```
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
)

ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
```

- 6 Save and close the main.cf file by verifying the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

- 7 Use the following command to start the VCS engine on one system:

```
# hstart
```

- 8 Enter the following command:

```
# hastatus
```

- 9 When "LOCAL_BUILD" is listed in the message column, start VCS on the other system with the following command:

```
# hstart
```


- 10 Verify that the service group resources are brought online. On one system, enter the following command:

```
# hagr -display
```

- 11 Enter the following command to check the status of the cvm group:

```
# hagr -state
```

- 12 Enter the following command to check the status of resources:

```
# hagr -state
```

- 13 Restart VCS on all the nodes.

Verifying SF Oracle RAC resources on all nodes

With the configuration file in place on each node, verify that the resources you configured come online.

To verify VCS, CVM, and CFS on all nodes

- 1 To verify that the resources are online enter the following command on each node:

```
# hastatus
```

- 2 If VCS is not running, then use the following procedure to start VCS, CVM, and CFS on all the nodes.

To start VCS, CVM, and CFS on all nodes

- ◆ Start VCS, CVM, and CFS from each node by issuing the following command:

```
# hstart
```

Verifying GAB port membership

Verify GAB port membership by running the `gabconfig -a` command.

The following is an example of this command and output:

```
# 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 v gen      ada416 membership 01
Port w gen      ada418 membership 01
```

Installing Oracle RAC 11g

This chapter includes the following topics:

- [About installing Oracle RAC 11g](#)
- [Installing Oracle Clusterware \(CRS\) using the `installsrac` command](#)
- [Installing Oracle Clusterware \(CRS\)](#)
- [Installing Oracle RAC 11g database binaries](#)
- [Verifying the Oracle RAC 11g CRS and database installation](#)
- [Completing Oracle RAC 11g post-installation tasks](#)


About installing Oracle RAC 11g

After completing the Oracle RAC 11g pre-installation tasks, proceed to install Oracle 11g software locally on each node. SF Oracle RAC supports Oracle 11.1 version or greater.

Review your Oracle installation manuals and the appropriate Oracle support websites, before installing Oracle RAC 11g.

[Figure 14-1](#) displays the steps required for an Oracle RAC 11g installation. Each step is discussed in detail in this chapter.

Figure 14-1 Oracle RAC 11g installation steps



Installing Oracle RAC 11g in an SF Oracle RAC environment involves these tasks:

- Installing Oracle Clusterware (CRS) using the `installsrac` command
- Installing Oracle Clusterware (CRS)
- Installing Oracle RAC 11g database binaries

- [Verifying the Oracle RAC 11g CRS and database installation](#)
- [Completing Oracle RAC 11g post-installation tasks](#)

Installing Oracle Clusterware (CRS) using the `installsfrac` command

Review the following procedure to set the environment variables and invoke the Oracle installer using the SF Oracle RAC configuration program.

Note: You must perform this procedure before installing Oracle 11g CRS and the Oracle 11g binaries.

Depending on what you want to install, do one of the following:

- If you want to install Oracle 11g CRS, use the procedure to invoke Oracle 11g CRS.
- If you want to install Oracle 11g binaries, invoke the Oracle 11g binaries installation.

Note: The benefit of installing Oracle CRS using the `installsfrac` command is that the Veritas libraries will automatically be linked with Oracle RAC.

To set the Oracle environment variables and invoke the Oracle installer

- 1 Launch the SF Oracle RAC configuration program.

```
# ./installsfrac -configure
```

By default, the `installsfrac` utility uses SSH for remote communication. However, RSH can be used in place of SSH by using the "rsh" option with the `installsfrac` utility.

```
# ./installsfrac -rsh -configure
```

The `installsfrac` utility proceeds to perform certain setup steps and validates your inputs. After which it prompts you to run the `rootpre.sh` on all cluster nodes.

After launching the SF Oracle RAC configuration program, the following general information and options are displayed. Select option 4.

```
Main Menu
```

```
. . . . .
```

```
* Keep root shells open on all systems to run
  INSTRUCTIONS ONLY steps
* Follow the steps in order
* Navigation keys for all screens: q to quit,
  b for previous menu
. ....

1) Check systems for SFRAC**INSTRUCTIONS ONLY**
2) Configure SFRAC
3) Prepare to install Oracle
4) Install or relink Oracle

Select an option [1-4,q] 4
```

- 2 The next configuration screen prompts you to check the supported SFRAC matrix for supported Oracle releases and patches. After checking the SFRAC matrix, proceed to the next step by pressing [RETURN].
- 3 Select an Oracle version for your installation. Select option 4 to install Oracle 11gR1.
- 4 Choose the Oracle task. Select option 1 to install Oracle Clusterware (CRS).
- 5 Set the DISPLAY environment variable that is required for the Oracle Universal Installer (OUI). Enter an IP address in the specified format.

For example:

```
Enter DISPLAY environment variable: [b] 10.20.12.150:0.0
```

- 6 Enter the Oracle UNIX user name. Select the default option (oracle).
The Oracle UNIX user name was previously set up:
See [“Creating Oracle user and groups”](#) on page 281.
- 7 Enter Oracle UNIX group name. Select the default option (oinstall).
The Oracle UNIX group name was previously set up:
See [“Creating Oracle user and groups”](#) on page 281.
- 8 Enter the absolute path of the CRS install image. After entering the absolute path, press [Return] to proceed.

For example:

```
/cdrom/Oracle11g/clusterware
```

- 9 Enter the location of the Oracle base directory.

In this example, enter /app/oracle.

10 Enter the absolute path of CRS Home directory.

In this example, enter /app/crshome.

If the CRS home directory you specified does not exist, then the SF Oracle RAC configuration program provides the following guidelines and commands to create one.

This step will create a directory on all cluster nodes.

```
mkdir -p /app/crshome
chown -R oracle:oinstall /app/crshome
chmod -R 744 /app/crshome
```

Run above command on all systems? [y,n,q, b]

If necessary, run the above commands and then enter the absolute path.
Proceed to the next step.

11 Enter the patch level. For example:

```
Main Menu > Install or relink Oracle >
Install Oracle Clusterware (CRS) > Patch Level
```

```
*****
Specifies the patch level of Oracle CRS/Database Software
being installed or relinked. For example:
Enter 0.2 for 10.2.0.2. Enter 0.7 for 9.2.0.7. .
*****
```

Enter Oracle Patch level for 11.1: [b] (0.6)

12 Confirm your configuration information.

Veritas Storage Foundation for Oracle RAC 5.0MP3 Configuration Program

```
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Installation Path:
```

```
/cdrom/Oracle11g/clusterware
```

```
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 11.1
```

```
Oracle Patch Level: 0.6  
Oracle Base: /app/oracle
```

```
Is this information correct? [y,n,q] (y)
```

After confirming the configuration information by entering "y", the SF Oracle RAC configuration program performs the following tasks:

- Copies several libraries
- Verifies fileset prerequisites on the nodes
- Checks the monitor for the appropriate display colors
- Starts the Oracle Clusterware (CRS) installer

The SF Oracle RAC installer invokes the Oracle Universal Installer (OUI) for the Oracle CRS Installer.

The following is an example of the SF Oracle RAC configuration program output:

- 13** For installation to proceed, Oracle requires you to run the rootpre.sh script located at:

```
/cdrom/Oracle/clusterware
```

Run the following commands on all nodes:

```
# cd /cdrom/Oracle/clusterware  
# ./rootpre.sh
```

- 14** Answer "y" or yes, when prompted with the question whether you have run rootpre.sh.

```
Has rootpre.sh been run by root? [y,n,q] y
```

- 15** The following information is presented:

```
Invoking OUI to install Oracle CRS. Please wait..... ok
```

```
Now install Oracle CRS software from the OUI. Refer to Oracle  
documentation for installation instructions.
```

```
NOTE: Near the end of the installation, Oracle CRS installer  
will prompt you to run the root.sh script.
```

```
Press <RETURN> here when prompted to run root.sh:
```


Proceed to the next section to install the Oracle CRS software from the OUI. When prompted, you will return to this step in this procedure to run root.sh.

Installing Oracle Clusterware (CRS)

The SFRAC for Oracle configuration program launches the Oracle Universal Installer (OUI). The OUI installs CRS.

Note: You must set the Oracle 11g environment variables before starting the Oracle 11g CRS installation process.

The CRS software is installed on each node in the location created for it.

See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 303.

To install CRS for an Oracle 11g installation

- 1 Set the Oracle 11g environment variables as described in the previous section using the SF Oracle RAC configuration program.

The SF Oracle RAC configuration program then launches the Oracle Universal Installer (OUI) to install CRS.

- 2 Proceed to follow the directions from the Oracle Universal Installer GUI.

As you run the OUI, be prepared with the following information required for the installation and configuration of the Oracle 11g CRS component.

Full path of the inventory directory	/app/oracle/oraInventory
Operating System group name	oinstall
Enter a name for the installation	OraCRS11g_home
Full path where you want to install the product.	/app/crshome

- 3 You will be prompted for Product-Specific Prerequisite Checks.

The Oracle installer verifies that your environment meets all of the minimum requirements for installing and configuring. You must manually verify and confirm the items that are flagged with warnings and items that require manual checks.

- 4 The Oracle Universal Installer displays the cluster and the nodes to be managed by the Oracle Clusterware.

Verify the displayed information.

For each node verify the name for the public IP address, the name for the private interconnect, and the name for the virtual IP address on the node.

- 5 The Oracle Universal Installer displays the node's network interfaces.

You will need to identify the planned use for each interface: Public, Private, or Do Not use.

Private interfaces are used by Oracle Clusterware for inter-node traffic.

See [“Configuring CRS and UDP private IP addresses for failover”](#) on page 291.

- 6 You will need to specify the Oracle Cluster Registry (OCR) location.

The Oracle Cluster Registry (OCR) stores cluster and database configuration information. You need to specify a cluster file system or a shared raw device containing at least 256 MB of free space that is accessible from all of the nodes in the cluster.

For example:

- OCR location: /ocrvote/ocr

External redundancy must be selected for this option.

Note: Mirroring is performed by Volume Manager.

- 7 You will need to specify the Voting Disk location

The Oracle Clusterware voting disk contains cluster membership information and arbitrates cluster ownership among the nodes of your cluster in the event of network failures. You must specify a cluster file system file or a shared raw device that is accessible by the same name from all of the nodes in the cluster. The Oracle Universal Installer requires at least 256 MB of free space for the voting disk that it creates.

For example:

- Voting Disk Location: /ocrvote/vote

External redundancy must be selected for this option.

Note: Mirroring is performed by Volume Manager.

- 8 The Oracle Universal Installer then presents a summary for the Oracle Clusterware and begins the Oracle Clusterware installation.
- 9 At the end of the Oracle Clusterware installation, you are prompted to invoke `root.sh` to complete the installation. For example:

```
Verifying binaries in /app/crshome on galaxy ..... ok
Verifying binaries in /app/crshome on nebula ..... ok
Copying SFRAC libskgxn on galaxy ..... ok
Copying SFRAC libskgxn on nebula ..... ok
Copying SFRAC ODM library on galaxy ..... ok
Copying SFRAC ODM library on nebula ..... ok
```

```
CRS post-installation setup is now complete.
You may now invoke root.sh to complete CRS installation.
```

- 10 The following scripts must be executed on your cluster nodes:
 - `/app/oracle/oraInventory/orainstRoot.sh`
 - `/app/crshome/root.sh`
- 11 To execute the configuration scripts, perform the following steps:
 - Return to the last step in the previous section to run the `root.sh`.
 - Press RETURN after running `root.sh` on all nodes.
- 12 At this point, the Oracle Universal Installer informs you that the CRS installation was successful.
- 13 Proceed to install the Oracle Database:
 - See [“Installing Oracle RAC 11g database binaries”](#) on page 323.

Installing Oracle RAC 11g database binaries

After installing the CRS component, install the Oracle database software.

Review the procedure to install the Oracle RAC 11g software in an SF Oracle RAC environment. Symantec recommends you to install Oracle binaries locally on each node.

See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 303.

For specific information on how to install Oracle, refer to your Oracle documentation.

To install Oracle database binaries for Oracle RAC 11g

- 1 Return to the SF Oracle RAC configuration program, navigate to the main menu, and select option 2 or the **Install Oracle RDBMS server** option.
- 2 Set the DISPLAY environment variable that is required for the Oracle Universal Installer. Enter an IP address in the specified format.
- 3 Enter Oracle UNIX user name. Select the default option (oracle).

The Oracle UNIX user name was previously set up:

See [“Creating Oracle user and groups”](#) on page 281.

- 4 Enter Oracle UNIX group name. Select the default option (oinstall).

The Oracle UNIX group name was previously set up:

See [“Creating Oracle user and groups”](#) on page 281.

- 5 Enter absolute path of the Database install image.

For example:

```
/cdrom/Oracle11g/database
```

- 6 Enter the location of the Oracle base directory.

In this example, enter /app/oracle.

- 7 Enter absolute path of CRS Home directory.

In this example, enter /app/crshome.

- 8 Enter absolute path of Database Home directory.

In this example, enter /app/oracle/orahome.

If the directory does not exist, the SF Oracle RAC will prompt you to create it and will create it for you by running a set of displayed commands.

Note: The absolute path entered above is only used for a clean installation of Oracle RAC 11g. If upgrading from an earlier version of Oracle RAC, then a different absolute path must be used for this step.

9 Enter the patch level. For example:

```
Main Menu > Install or relink Oracle >
Install Oracle Clusterware (CRS) > Patch Level

*****
Specifies the patch level of Oracle CRS/Database Software
being installed or relinked. For example:
Enter 0.2 for 10.2.0.2. Enter 0.7 for 9.2.0.7. .
*****

Enter Oracle Patch level for 11.1: [b] (0.6)
```

- 10** The SF Oracle RAC configuration program then displays the configuration information for verification. Confirm your configuration information.

After reviewing and approving this information the Oracle Universal Installer (OUI) is invoked.

```
Oracle environment information verification
```

```
Oracle Unix User: oracle
```

```
Oracle Unix Group: oinstall
```

```
Oracle Clusterware (CRS) Installation Path:
```

```
/cdrom/Oracle11g/clusterware
```

```
Oracle Clusterware (CRS) Home: /app/crshome
```

```
Oracle Release 11.1
```

```
Oracle Patch Level: 0.6
```

```
Oracle Base: /app/oracle
```

```
Oracle Home: /app/oracle/orahome
```

```
RDBMS Installation path:
```

```
/cdrom/Oracle11g/database
```

```
Is this information correct? [y,n,q] (y)
```

```
Verifying CRS on galaxy ..... Running
```

```
Verifying CRS on nebula ..... Running
```

```
Comparing Oracle UNIX user id on all systems ..... Done
```

```
Comparing Oracle group id on all systems ..... Done
```

```
ORACLE DB INSTALLATION
```

```
Invoking OUI to install Oracle RDBMS Software. Please wait..... ok
```

```
Now install Oracle database software from the OUI. Refer to Oracle  
documentation for installation instructions.
```

11 The following information is also presented to you:

Near the end of the installation, Oracle Database Software installer will prompt you to run the root.sh script

RETURN here when prompted to run root.sh:

Press RETURN here when prompted to run root.sh.

Proceed to to install the Oracle DB software from the OUI.

12 The SF Oracle RAC installer invokes the Oracle Universal Installer (OUI) for Database Software Installation.

Note: Consult your Oracle documentation for any required information, and for information about installing the Oracle database software using the Oracle interface.

13 As you run the OUI, be prepared with the following information required for the installation and configuration of the Oracle database software:

Specify File Locations dialog box

- Oracle Home location: The installation destination(\$ORACLE_HOME). Depending on where you install Oracle 11g binaries, this location is either on shared storage or an identical location on each of the local cluster nodes. For example, /app/oracle/orahome.
- Oracle Base location: The installation destination for all Oracle software and configuration-related files. For example, /app/oracle.
- The path to \$ORACLE_HOME.

Specify Hardware Cluster Installation Mode dialog box

The names of the cluster nodes and their associated host names. In the example Oracle installation, the nodes and their associated host names are galaxy and nebula.

14 When you come to the Select Installation Method dialog box, choose not to have a starter database created. Do not click **Create a starter database**. Symantec recommends you to create the database later.

- 15 From the OUI and for a basic installation, review and approve the various installation locations (Oracle Base and Oracle Home).
- 16 From the OUI, select the nodes in the hardware cluster where the OUI will install the software.

At this point in the installation process, the OUI runs a product specific prerequisite check. Any items that are flagged must be manually checked and configured.

- 17 From the OUI, specify the install locations for Oracle Base and software location.
- 18 When the OUI presents the Select Configuration Option screen, select the **Create a Database** option and the **Install Software only** option.

When the OUI presents the Privileged Operating System Groups screen, make any adjustments required for your configuration and select Next to proceed with the installation.

- 19 The OUI then presents a summary for the Oracle Database and begins the Oracle Database installation.
- 20 You will then be prompted to execute configuration script. The following script must be executed on your cluster nodes:

■ `/app/oracle/orahome/root.sh`

- 21 To execute the configuration script, perform the following steps:

Return to the earlier step in this section to run the root.sh.

Press RETURN after running root.sh on all nodes.

Verifying the Oracle RAC 11g CRS and database installation

To verify that the installations of the Oracle RAC 11g CRS and database have succeeded, issue the command described below from any node in the cluster.

The output should show processes running on all nodes, as in the following example:

```
# $CRS_HOME/bin/crs_stat
NAME=ora.galaxy.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy
```



```
NAME=ora.galaxy.gsd
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy

NAME=ora.galaxy.ons
TYPE=application
TARGET=ONLINE
STATE=ONLINE on galaxy

NAME=ora.nebula.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on nebula

NAME=ora.nebula.gsd
TYPE=application
TARGET=ONLINE
STATE=ONLINE on nebula

NAME=ora.nebula.ons
TYPE=application
TARGET=ONLINE
STATE=ONLINE on nebula
```

Completing Oracle RAC 11g post-installation tasks

After installing the Oracle software, complete these tasks:

- [Adding Oracle RAC 11g patches or patchsets](#)
- [Relinking the SF Oracle RAC libraries to Oracle RAC 11g](#)
- [Creating the Oracle RAC 11g database](#)
- [Adding UDP IP addresses to the Oracle init file](#)
- [Upgrading databases](#)
- [Configuring the Oracle RAC 11g service group in a VCS configuration](#)

Adding Oracle RAC 11g patches or patchsets

If required for your Oracle RAC 11g installation, apply any Oracle 11g patches or patchsets as needed.

Before installing any Oracle 11g patch or patchset software:

- Review the latest information on supported Oracle 11g patches and patchsets:
<http://entsupport.symantec.com/docs/280186>
- You must have installed Oracle 11g R1 software.
- Review the patch or patchset notes that accompany the patch or patchset for instructions on installing them and performing the post-installation operations.

Relinking the SF Oracle RAC libraries to Oracle RAC 11g

If you added or upgraded the Oracle patches, you must relink the SF Oracle RAC libraries to Oracle. Symantec recommends that you relink the SF Oracle RAC libraries only after completing all the required patch additions.

Relinking Oracle

- 1 Return to the SF Oracle RAC configuration program, navigate to the main menu, and select the **Relink Oracle** option

```
Main Menu > Install or relink Oracle > Choose task

*****
* Oracle Clusterware (CRS) must be installed before RDBMS server
* Relinking is required if you installed Oracle software or patches
  from outside SFRAC installer
*****

1) Install Oracle Clusterware (CRS)
2) Install Oracle RDBMS server
3) Relink Oracle

b) [Go to previous menu]

Select an option [1-3,b,q]  3
```

- 2 Review the information from the configuration program, and press RETURN to continue.

```
Main Menu > Install or relink Oracle > Relink Oracle

*****
This step will:
* Relink Oracle software using SFRAC static libraries - OR-
* Copy dynamic SFRAC libraries to Oracle directories
*****

Press [Return] to continue:
```

- 3 Enter Oracle UNIX user name. Select the default option (oracle).
 The Oracle UNIX user name was previously set up:
 See [“Creating Oracle user and groups”](#) on page 281.
- 4 Enter Oracle UNIX group name. Select the default option (oinstall).
 The Oracle UNIX group name was previously set up:
 See [“Creating Oracle user and groups”](#) on page 281.

- 5 Enter the location of the Oracle base directory.

In this example, enter /app/oracle.

- 6 Enter absolute path of CRS Home directory.

In this example, enter /app/crshome.

- 7 Enter absolute path of Database Home directory.

In this example, enter /app/oracle/orahome.

- 8 Enter the patch level.

For example:

```
Main Menu > Install or relink Oracle >
Install Oracle Clusterware (CRS) > Patch Level
*****
Specifies the patch level of Oracle CRS/Database Software
being installed or relinked. For example:
Enter 0.2 for 10.2.0.2. Enter 0.7 for 9.2.0.7. .
*****
Enter Oracle Patch level for 11.1: [b] (0.6)
```

9 The SF Oracle RAC configuration program then displays the configuration information for verification.

```
Veritas Storage Foundation for Oracle RAC 5.0MP3 Configuration Program
```

```
Oracle environment information verification
```

```
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 11.1
Oracle Patch Level: 0.6          Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
```

After reviewing and approving this information the Oracle relinking process is started.

The following is an example of output from the relinking process.

```
Verifying binaries in /app/oracle/orahome on nebula ..... ok
Verifying binaries in /app/oracle/orahome on galaxy ..... ok
Copying SFRAC libskgxn on nebula ..... ok
Copying SFRAC libskgxn on galaxy ..... ok
Copying SFRAC ODM library on nebula ..... ok
Copying SFRAC ODM library on galaxy ..... ok
Relinking Oracle on nebula ..... ok
Relinking Oracle on galaxy ..... ok
Oracle Relinking is now complete.
```

After the Oracle relinking process is finished, the SF Oracle RAC configuration screen appears. Enter q to quit the program.

```
Main Menu > Install or relink Oracle > Choose task
```

```
*****
* Oracle Clusterware (CRS) must be installed before RDBMS server
* Relinking is required if you installed Oracle software or patches
  from outside SFRAC installer
*****
```

- 1) Install Oracle Clusterware (CRS)
- 2) Install Oracle RDBMS server
- 3) Relink Oracle

```
b) [Go to previous menu]
Select an option [1-3,b,q]  q
```

Note: Copy the location of the installsfrac log files. If necessary, these log files can be used for troubleshooting purposes.

10 Proceed to create the Oracle 11g database:

See [“Creating the Oracle RAC 11g database”](#) on page 334.

Creating the Oracle RAC 11g database

Create the Oracle RAC 11g database on shared storage.

Use your own tools or scripts or review the guidelines on using the Oracle dbca (Database Creation Assistant) tool to create a database on shared raw VxVM volumes or shared VxFS file systems. Refer to the your Oracle documentation for instructions on how to install the Oracle RAC 11g database.

For information about creating a test database:

See [“About creating a test database”](#) on page 617.

Note: If you plan to configure global clusters, then set up the Oracle RAC 11g database only on the primary site. On the secondary site, the database will be replicated.

Adding UDP IP addresses to the Oracle init file

UDP IP addresses must be added to the Oracle init file as cluster_interconnects parameters.

Follow the procedure described below to add IP addresses to the Oracle init file.

To add IP addresses to the Oracle init file

- 1** Login to each system as oracle user.
- 2** Next, ensure that all instances database are shut down.

3 Use the `ifconfig -a` command to ensure that these UDP's private IPs are running and in UP state.

For example, on the cluster node galaxy:

```
[oracle@] /oracle> ifconfig -a
en0:
flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,
SIMPLEX,MULTICAST,GROUPRT,64BI
T,CHECKSUM_OFFLOAD(ACTIVE),PSEG,LARGESEND,CHAIN>
    inet 10.140.93.23 netmask 0xfffff800 broadcast 10.140.95.255
    inet 10.140.91.41 netmask 0xfffff800 broadcast 10.140.95.255
    tcp_sendspace 131072 tcp_recvspace 65536 rfc1323 0
en1:
flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,
SIMPLEX,MULTICAST,GROUPRT,64BI
T,CHECKSUM_OFFLOAD(ACTIVE),PSEG,LARGESEND,CHAIN>
    inet 192.168.1.1 netmask 0xfffff800 broadcast 192.168.15.255
    inet 192.168.1.2 netmask 0xfffff800 broadcast 192.168.15.255
en2:
flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,
SIMPLEX,MULTICAST,GROUPRT,64BI
T,CHECKSUM_OFFLOAD(ACTIVE),PSEG,LARGESEND,CHAIN>
    inet 192.168.2.1 netmask 0xfffff800 broadcast 192.168.31.255
lo0:
flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,
SIMPLEX,MULTICAST,GROUPRT,64BIT>
    inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
    inet6 ::1/0
    tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
[oracle@] /oracle>
```

4 As oracle user, enter the following commands for the cluster_interconnects:

```
$ sqlplus '/ as sysdba'

SQL> alter system set cluster_interconnects='192.168.1.2':'192.168.2.1'
scope=spfile sid='vrts1'

SQL> alter system set cluster_interconnects='192.168.1.4':'192.168.2.2'
scope=spfile sid='vrts2'

SQL> exit;
```

Refer to [Table 13-4](#) for information about the IP addresses used in the above command examples.

Upgrading databases

If you currently have Oracle databases running and want to migrate them to the latest Oracle patch level, refer to the README.html file downloaded with the patch.

For information about upgrading and migrating current databases to Oracle RAC 11g:

See [“About upgrading and migrating to Oracle RAC 11gR1 software”](#) on page 361.

Configuring the Oracle RAC 11g service group in a VCS configuration

After installing Oracle 11g and creating a database, proceed to modify the VCS configuration file. Review the sample VCS configurations and details on configuring service groups in an Oracle 11g environment.

See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.

Installing Oracle RAC 11g manually

This chapter includes the following topics:


- [About installing Oracle RAC 11g manually](#)
- [Installing Oracle RAC 11g manually](#)

About installing Oracle RAC 11g manually

Symantec recommends using the SF Oracle RAC installer and the `installsfraconfig` command for an Oracle RAC 11g binaries installation. However, some situations may require the manual installation of Oracle RAC 11g binaries as described in this chapter. These situations may arise when either a customized installation is required by the user or when there may be a DBA permissions issue.

[Figure 15-1](#) displays the steps required for a manual Oracle RAC 11g installation.

Figure 15-1 Oracle RAC 11g manual installation steps



Installing Oracle RAC 11g manually

The following procedures are required to manually install Oracle RAC 11g:

- [Configuring MultiPrivNIC for CRS/UDP cache fusion](#)
- [Setting up environment variables](#)
- [Linking Veritas Membership library](#)
- [Installing the Oracle Clusterware \(CRS\)](#)
- [Installing the Oracle RAC 11g database binaries](#)
- [Adding Oracle RAC 11g patches or patchsets](#)
- [Link the ODM library](#)
- [Creating the Oracle RAC 11g database](#)
- [Add CRS under VCS control](#)
- [Configure cluster interconnects](#)
- [Upgrading databases](#)
- [Configuring the Oracle RAC 11g service group in a VCS configuration](#)

Configuring MultiPrivNIC for CRS/UDP cache fusion

If you did not configure MultiPrivNIC for CRS/UDP cache fusion in the pre-installation procedures, proceed to manually configure MultiPrivNIC as described in this section:

See [“Manually configuring MultiPrivNIC for UDP cache fusion”](#) on page 299.

Setting up environment variables

Set up the required environment variables for CRS and Oracle binaries.

To set up the environment variables

- 1 Log in as oracle user on each node.
- 2 On each node, set the proper environment variables (the following examples are for ksh):

```
$ cat .profile
set -o vi

PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:\
/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]           # This is at Shell startup.  In normal
then echo "$MAILMSG"       # operation, the Shell checks
fi                          # periodically.

ORACLE_HOME=/app/oracle/orahome; export ORACLE_HOME
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data; export ORA_NLS33
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib; export CLASSPATH
DISPLAY=YOUR.DISPLAY.IP; export DISPLAY
CRS_HOME=/app/crshome; export CRS_HOME
ORA_CRS_HOME=/app/crshome; export ORA_CRS_HOME

LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$CRS_HOME/lib:\
$CRS_HOME/lib32:/usr/lib:/opt/VRTSdbms3/lib; export LIBPATH
PATH=$PATH:$ORACLE_HOME/bin:$CRS_HOME/bin:/opt/VRTS/bin:\
/opt/VRTSdbed/bin; export PATH

ORACLE_SID=PROD; export ORACLE_SID

umask 022

set -o vi
$
```

- 3 On the first node, set the DISPLAY variable.

- For Bourne Shell (sh or ksh), type:

```
$ DISPLAY=host:0.0 export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY host:0.0
```

Linking Veritas Membership library

Perform the following steps to link the Veritas membership library.

To link Veritas Membership library

- 1 Create required sub-directories if they not present.

For example, enter the following commands:

```
# mkdir -p /usr/sbin/cluster/utilities
# echo `lltstat -C` > /usr/sbin/cluster/utilities/cldomain
# mkdir /opt/ORCLcluster
# mkdir /opt/ORCLcluster/lib
# mkdir /opt/ORCLcluster/lib32
# mkdir /opt/ORCLcluster/bin
```

- 2 Archive skgxn library.

For example, enter the following commands:

```
# cd /opt/VRTSvcs/rac/lib
# rm libskgxn2.a
# cp libvcsmm.so shr_skgxn2.o
# ar -X32 -qv libskgxn2.a shr_skgxn2.o
# rm shr_skgxn2.o
# cd /opt/VRTSvcs/rac/lib64
# rm libskgxn2.a
# cp libvcsmm.so shr_skgxn2.o
# ar -X64 -qv libskgxn2.a shr_skgxn2.o
# rm shr_skgxn2.o
```

3 Copy the skgxn library and clsinfo to /opt/ORCLcluster directory.

For example, enter the following commands:

```
# cp /opt/VRTSvcs/rac/bin/clsinfo \
/opt/ORCLcluster/bin

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

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

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

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

Installing the Oracle Clusterware (CRS)

Use the Oracle RAC 11g runInstaller utility to install the CRS component before installing the Oracle RAC 11g binaries and the database.

The software is installed on each node in the location created earlier:

See [“Preparing \\$CRS_HOME and \\$ORACLE_HOME on each node”](#) on page 303.

To install the CRS, complete the following procedure.

Note: For additional information, refer to your Oracle RAC 11g documentation.

To install the CRS

- 1 Ensure that the Oracle Clusterware oracle user has the following capabilities: CAP_NUMA_ATTACH, CAP_BYPASS_RAC_VMM, and CAP_PROPAGATE.

To check existing capabilities, enter the following command as `root`; in this example, the Oracle Clusterware oracle user is `oracle`:

```
# /usr/bin/lsuser -a capabilities oracle
```

To add capabilities, enter a command similar to the following:

```
# /usr/bin/chuser \  
capabilities=CAP_NUMA_ATTACH CAP_BYPASS_RAC_VMM CAP_PROPAGATE oracle
```

- 2 The `ncargs` attribute controls the maximum command line length, including the environment variables on the AIX nodes. Change this attribute to a larger value before invoking the Oracle Installer.

The following is an example of the command to set this attribute:

```
# /usr/sbin/chdev -l sys0 -a ncargs=128
```

- 3 Log in as Oracle user and launch the Oracle RAC 11g `runInstaller` utility on the first node.

With the Oracle RAC 11g CRS disc in the CD-ROM drive, enter the following command as Oracle user:

```
$ cd /cdrom/cdrom0
```

```
$ ./runInstaller
```

- 4 As you run the installer, prepare the following information required for the installation and configuration of the CRS component:

- The installation destination, `$CRS_HOME`, on each node (for example, `/app/crshome`), and the path to `$CRS_HOME` for the Specify File Locations dialog box.
- The names of the nodes and their associated host names for the Cluster Configuration dialog box.
- The cluster name and the nodes to be managed by the Oracle Clusterware. Note that the Oracle program may already have this information and display it.

- Identify the planned use for each global interface that is displayed by Oracle: Public, Private, or Do Not Use. Note that the Oracle program may already have this information and display it.
 - The name of a file in the OCR directory or raw volume for the Oracle Cluster Registry dialog box.
For example: /ocrvote/ocr
 - The name of a file in the VOTE-disk directory or raw volume for the Voting Disk dialog box.
For example: /ocrvote/vote
- 5 When runInstaller prompts you to run the script
/\$ORACLE_BASE/oraInventory/orainstRoot.sh, verify that the script exists on each node before proceeding.
- If the script does not exist on each node, copy it from the first node to each of the other cluster nodes.
- 6 Run the /\$ORACLE_BASE/oraInventory/orainstRoot.sh script on each node.
- 7 When runInstaller prompts you to run the script, run root.sh on each node:
- ```
cd $CRS_HOME
```
- ```
# ./root.sh
```
- This starts the CRS daemons on the node where you enter the command.
- 8 Click OK to exit the installer.

Installing the Oracle RAC 11g database binaries

The following procedure describes how to install the Oracle RAC 11g database binaries.

To install the Oracle RAC 11g database binaries

- 1 On the first node, launch the Oracle RAC 11g runInstaller utility. With the Oracle RAC 11g Enterprise Edition disc in the CD-ROM drive, enter the following:
- ```
$ cd /cdrom/cdrom0
```
- ```
$ ./runInstaller
```
- 2 As you run the installer, prepare the following information for the installation and configuration of the Oracle RAC 11g binaries component:

- The destination of the installation (for example, \$ORACLE_HOME), and the /app/oracle/orahome path for the Specify File Locations dialog box. This location is either on shared storage or an identical location on each of the local cluster nodes.
 - The names of the nodes (for example, galaxy and nebula) for the Specify Hardware Cluster Installation Mode dialog box.
 - The desired installation edition for the Select Installation Type dialog box.
- 3 In the Select Database Configuration dialog box, you can choose not to create a starter database by selecting Do not create a starter database. Veritas recommends creating a database at a later point.
 - 4 In the Summary dialog box, click Install. When the installation ends, a Setup Privileges message appears. Do not click OK; instead, first run the root.sh script on each node. For example:

```
# cd /app/oracle/orahome

# ./root.sh
```
 - 5 Click OK on the Setup Privileges message.

Adding Oracle RAC 11g patches or patchsets

If required for your Oracle RAC 11g installation, apply any Oracle 11g patches or patchsets as needed.

Before installing any Oracle 11g patch or patchset software, perform the following tasks:

- Review the latest information on supported Oracle 11g patches and patchsets: <http://entsupport.symantec.com/docs/280186>
- You must have installed Oracle 11g R1 software.
- Review the patch or patchset notes that accompany the patch or patchset for instructions on installing them and performing the post-installation operations.

Link the ODM library

Replace Oracle's ODM library with Veritas' ODM library by issuing the following commands:

```
$ mv $ORACLE_HOME/lib/libodm11.so libodm11.so.oracle
$ ln -s /opt/VRTSodm/lib/libodm64.so libodm11.so
```

Add CRS under VCS control

Perform the following steps to add CRS under VCS control.

To add CRS under VCS control

- 1 Log in to one system as root.
- 2 Save your existing configuration to prevent any changes while you modify main.cf:

```
# haconf -dump makero
```

If your configuration is not writable, a warning appears: "Cluster not writable". You may safely ignore the warning.

- 3 Make sure that VCS is not running while you edit the main.cf by using the `hastop` command. This command stops the VCS engine on all systems and leaves the resources available.

```
# hastop -all -force
```

- 4 Using vi or another text editor, edit the main.cf file:

```
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  
)  
  
:  
:  
cssd requires ocrvote_mnt  
cssd requires multi_priv
```

- 5 Save and close the file.
- 6 Verify the syntax of the file `/etc/VRTSvcs/conf/config/main.cf` by typing the following command:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 7 Use the following command to start the VCS engine on the system which make the changes on main.cf:

```
# hstart
```

- 8 Enter the following command:

```
# hstatus
```

- 9 When "LOCAL_BUILD" is listed in the message column, start VCS on the other systems with the following command:

```
# hstart
```

- 10 Verify that the service group resources are brought online.

On one system, enter the following command:

```
# hagrps -state
```

Creating the Oracle RAC 11g database

Create the Oracle RAC 11g database on shared storage.

Use your own tools or scripts or review the guidelines on using the Oracle dbca (Database Creation Assistant) tool to create a database on shared raw VxVM volumes or shared VxFS file systems. Refer to your Oracle documentation for instructions on how to install the Oracle RAC 11g database.

For information about creating a repository database, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Note: If your configuration supports replicating global clusters using SF Oracle RAC, then the Oracle RAC 11g database should not be set up on a GCO secondary site. Follow the above steps only when setting up Oracle on a primary GCO site.

Configure cluster interconnects

Perform the following steps to configure cluster interconnects.

To configure cluster interconnects

- 1 On each database instance, log in as sysdba and enter the following queries.

For example, on galaxy for instance 1 – vrts1 :

```
alter system set cluster_interconnects='192.168.1.2:192.168.2.1' scope=spfile sid='vrts1';
```

For example, on nebula for instance 2 – vrts2 :

```
alter system set cluster_interconnects='192.168.1.4:192.168.2.2' scope=spfile sid='vrts2';
```

Refer to [Table 13-4](#) for information about the IP addresses used in the above command examples.

- 2 Restart the database.
- 3 You can verify the cluster_interconnects on each system using following sql query:

```
select * from v$configured_interconnects where IS_PUBLIC='NO';
```

Upgrading databases

If you currently have Oracle databases running and want to migrate them to the latest Oracle patch level, refer to the README.html file downloaded with the patch.

For information about upgrading and migrating current databases to Oracle RAC 11g:

See [“About upgrading and migrating to Oracle RAC 11gR1 software”](#) on page 361.

Configuring the Oracle RAC 11g service group in a VCS configuration

After installing Oracle 11g and creating a database, proceed to modify the VCS configuration file. Review the sample VCS configurations and details on configuring service groups in an Oracle 11g environment:

See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.

Configuring VCS service groups for Oracle RAC 11g

This chapter includes the following topics:


- [About VCS service group for Oracle RAC 11g dependencies](#)
- [Configuring VCS resources for Oracle](#)
- [Location of VCS log files](#)

About VCS service group for Oracle RAC 11g dependencies

A VCS service group is a collection of resources working together to provide application services to clients. A VCS service group typically includes multiple resources that are both hardware and software based. For example, a resource may be a physical component such as a disk or network interface card, or a software component such as Oracle 11g or a Web server, or a configuration component such as an IP address or mounted file system.

[Figure 16-1](#) displays the steps required to configure the VCS service groups for an Oracle RAC 11g installation. Each step is discussed in detail in this chapter.

Figure 16-1 VCS service group configuration steps



VCS service group dependencies with or without the VCS Oracle agent

VCS service group dependencies are based upon whether you use the VCS Oracle agent or not. The following figures illustrate the dependencies.

- In a configuration *with* the VCS Oracle agent, VCS controls the Oracle database. This is the preferred configuration. An online local firm dependency exists between the Oracle group and the CVM group.

Note: In a setup with multiple databases, the configuration with the VCS Oracle agent must be used.

For additional details on service group dependencies, refer to the *Veritas Cluster Server User's Guide*.

[Figure 16-2](#) displays a schematic of a configuration with the VCS Oracle agent, and the dependencies between the VCS service groups and resources. In this figure, the name inside the circular shape is the "resource type". The name outside of the circular shape is the "resource name".

- In a configuration *without* the VCS Oracle agent, CRS controls the database.

Note: The configuration without the VCS Oracle agent may be used only in single database setups.

Figure 16-3 displays a schematic of a configuration without the VCS Oracle agent, and the dependencies between the VCS service groups and resources. In this figure, the name inside the circular is the "resource type". The name outside of the circular shape is the "resource name".

Figure 16-2 Configuration with the VCS Oracle agent (Preferred configuration)




Table 16-1 defines the CVM service group resources displayed in Figure 16-2.

Figure 16-3 Configuration without the VCS Oracle Agent




Table 16-1 defines the CVM service group resources displayed in Figure 16-3.

Table 16-1 CVM service group resource definitions

Resource name	Resource type	Definitions
cssd	Application	This resource monitors the Oracle RAC 11g clusterware processes. The purpose of the cssd resource is to ensure that the dependency of cssd on the Oracle Cluster Registry (OCR) and VOTE resources and the MultiPrivNIC resource (optional) are satisfied. For additional information: See “CSSD agent” on page 638.
ora_priv	MultiPrivNIC	The MultiPrivNIC agent provides redundancy for IP addresses for Oracle's UDP/IP cache fusion capabilities and CRS heartbeat IP address. For additional information: See “MultiPrivNIC agent” on page 635.

Table 16-1 CVM service group resource definitions (*continued*)

Resource name	Resource type	Definitions
vxfsckd	CFSfsckd	The Veritas file system check daemon (vxfsckd) is responsible for ensuring file system consistency when a node fails that was a primary node for a shared file system. Cluster file system check daemon (CFSfsckd) is responsible for ensuring file system consistency in a cluster when a node fails that was a primary node for a shared file system.
cvm_clus	CVMCluster	The CVMCluster agent controls system membership on the cluster port associated with Veritas Volume manager (VxVM). For additional information: See “CVMCluster agent” on page 623.
cvm_vxconfigd	CVMVxconfigd	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 configuration information stored on disks. CVMVxconfigd must be present in the CVM service group. For additional information: See “CVMVxconfigd agent” on page 625.
ocrvote_voldg	CVMVolDg	The CVMVolDg agent represents and controls CVM disk groups and the CVM volumes within the disk groups. For additional information: See “CVMVolDg agent” on page 627.
ora_db1 and ora_db2	Oracle	Oracle database agent. Oracle application is the resource type. Note: This resource is only available for the configuration with the VCS Oracle agent.

Table 16-1 CVM service group resource definitions (continued)

Resource name	Resource type	Definitions
crs.ora.mnt	CFS_Mount	The Cluster Ready Service Oracle Mount agent (crs_ora_mnt) brings online, takes offline, and monitors an Oracle cluster file system mount point. The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point. Note: This resource is only available for the configuration with the VCS Oracle agent. For additional information: See “CFSMount agent” on page 629.
oradata_mnt	CFSMount	The oradata_mnt agent brings online, takes offline, and monitors the Oracle cluster file system mount point. The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point. For additional information: See “CFSMount agent” on page 629.
oradata_voldg	CVMVolDg	The oradata_voldg agent represents and controls Oracle disk groups and the Oracle volumes within the disk groups. The CVMVolDg agent represents and controls CVM disk groups and the CVM volumes within the disk groups. For additional information: See “CVMVolDg agent” on page 627.

Automatic database starting prevention

The CRS and Oracle agent may attempt to start the instance at the same time if the database mount is available. To prevent automatic database starting, you must change the Management policy for the database (automatic, manual) to MANUAL using the Oracle `SRVCTL` command. The command will change the `AUTO_START` attribute of CRS database and instance resources.

To prevent automatic database starting

- ◆ Depending on the status of the database:

- If the database is already registered and running:

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

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

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

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

Configuring VCS resources for Oracle

Oracle service groups can be configured using the following methods:

- By editing the VCS configuration file, `main.cf`, to define the Oracle service groups:

See [“Editing the main.cf file to configure the Oracle service groups”](#) on page 357.

- By configuring the Oracle service groups using the CLI:

See [“Configuring the Oracle service groups using the CLI”](#) on page 359.

Editing the main.cf file to configure the Oracle service groups

This section describes how to manually edit the `main.cf` file to configure Oracle service groups.

To edit the `main.cf` file to configure Oracle service groups, perform the following procedures:

- [To configure VCS service groups for Oracle RAC 11g](#)
- [To save and check the configuration](#)
- [To verify the state of newly added resources](#)

To configure VCS service groups for Oracle RAC 11g

- 1 Log in to one system as `root`.
- 2 As root user, save your existing configuration to prevent any changes while you modify `main.cf`:

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

If the configuration is not writable, a warning appears: "Cluster not writable." You may safely ignore the warning.

- 3 Make sure VCS is not running while you edit `main.cf` by using the `hastop` command to stop the VCS engine on all systems and leave the resources available:

```
# hastop -all -force
```

- 4 Make a backup copy of the `main.cf` file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 5 Using `vi` or another text editor, edit the `main.cf` file by creating and modifying the Oracle service groups.

Use the following `main.cf` samples as a guideline:

See [“Sample main.cf for Oracle RAC 11g with Oracle agent”](#) on page 525.

See [“Sample main.cf for Oracle RAC 11g without Oracle agent”](#) on page 522.

Modifying the VCS configuration

For additional information and instructions on modifying the VCS configuration by editing the `main.cf` file, refer to the *VERITAS Cluster Server User's Guide*.

Saving and checking the configuration

After you finish configuring the Oracle service groups by editing the `main.cf` file, verify the new configuration.

To save and check the configuration

- 1 Save and close the main.cf file.
- 2 Verify the syntax of the file `/etc/VRTSvcs/conf/config/main.cf`:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

- 3 Use the following command to start the VCS engine on one system:

```
# hstart
```

- 4 Enter the following command:

```
# hstatus
```

- 5 When "LOCAL_BUILD" is listed in the message column, start VCS on the other system with the following command:

```
# hstart
```

- 6 Verify that the service group resources are brought online. On one system, enter the following command:

```
# hagrps -display
```

To verify the state of newly added resources

- 1 Use `hagrps -state` command to check the status of groups.
- 2 Use `hares -state` command to check the status of resources.

A restart is now required to make sure that CRS and Oracle RAC database instances use Symantec libraries.

Configuring the Oracle service groups using the CLI

For detailed information and instructions on modifying the VCS configuration by using the CLI (command line interface), refer to the *Veritas Cluster Server User's Guide*.

Location of VCS log files

On all cluster nodes, access and review the following log files for any errors or status messages:

```
/var/VRTSvcS/log/engine_A.log
```

When large amounts of data are written, multiple log files may be required, such as engine_B.log, engine_C.log, and so on. The engine_A.log contains the most recent data.

Upgrading to Oracle RAC 11g and migrating the database

This chapter includes the following topics:

- [About upgrading and migrating to Oracle RAC 11gR1 software](#)
- [Overview of upgrade and migration tasks](#)
- [Upgrading from Oracle RAC 9iR2 to Oracle RAC 11gR1](#)
- [Upgrading from Oracle RAC 10g to Oracle RAC 11gR1](#)

About upgrading and migrating to Oracle RAC 11gR1 software

SF Oracle RAC 5.0 MP3 supports the following Oracle software upgrades:

- Oracle RAC 9iR2 to Oracle RAC 11g R1
- Oracle RAC 10gR1 to Oracle RAC 11g R1
- Oracle RAC 10gR2 to Oracle RAC 11g R1

The upgrade procedures for Oracle RAC 10gR1 and Oracle RAC 10gR2 are generally the same, except where specific steps are identified as applicable for a specific version of Oracle RAC.

Overview of upgrade and migration tasks

If you have not already installed and configured Veritas Storage Foundation for Oracle RAC 5.0 MP3, you will need to do so during the preparation for migration procedures.

The upgrade and migration procedures are specific for the OS, current version of Oracle, and target version of Oracle you are using. The high-level order of procedures is as follows:

- Upgrade OS if required
See [“Supported operating systems”](#) on page 39.
- Upgrade to Veritas Storage Foundation for Oracle RAC 5.0 MP3 if you have not already upgraded
See [“Supported upgrade options for SF Oracle RAC 5.0 MP3”](#) on page 145.
- Install target version of Oracle RAC, including patchsets if required
- Migrate the database
- Relink the Symantec libraries

Upgrading from Oracle RAC 9iR2 to Oracle RAC 11gR1

The upgrade and migration procedure assumes that the beginning configuration includes the following components, and that these components are up and running on the cluster nodes:

- Supported operating system
- SF Oracle RAC 5.0 MP3
- Oracle RAC 9i R2

The migration procedure consists of the following tasks:

- [Performing pre-upgrade tasks for Oracle RAC 11g](#)
- [Installing Oracle RAC 11gR1](#)
- [Migrating an existing Oracle RAC 9iR2 database to Oracle RAC 11g](#)
- [Performing post-upgrade tasks Oracle RAC 11gR1](#)

Performing pre-upgrade tasks for Oracle RAC 11g

The migration procedures assume that you have installed Oracle RAC 9i on your cluster nodes.

Before you upgrade and migrate the Oracle RAC 9i software, do the following:

- Review the configuration of VCS service groups.
See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.
See [“Configuring VCS resources for Oracle”](#) on page 357.
- Review the example main.cf file.
See [“Sample main.cf for Oracle RAC 11g with Oracle agent”](#) on page 525.
- Configure Oracle RAC 11gR1 prerequisites.
See [“About preparing to install Oracle RAC 11g”](#) on page 275.
- Install CRS for Oracle RAC 11g R1.
See [“Installing Oracle Clusterware \(CRS\)”](#) on page 321.
- Install Oracle RAC 11gR1 binaries.
See [“Installing Oracle RAC 11g database binaries”](#) on page 323.
- Add required patchsets
See [“Adding Oracle RAC 11g patches or patchsets”](#) on page 329.

Perform the following pre-upgrade tasks before migrating from Oracle RAC 9iR2 to Oracle RAC 11gR1.

To perform pre-upgrade tasks

- 1 Upgrade OS and install any patches, if required. For details:
See Oracle metalink documentation.
- 2 Take a hot or cold backup of the existing Oracle RAC 9iR2 database.
- 3 Take a backup of the existing Oracle Home and Central Inventory.
- 4 Freeze the VCS service groups that contain resources to monitor Oracle database instances and/or listener processes. As root user, enter:

```
# haconf -makerw

# hagrps -freeze oracle_group -persistent

# haconf -dump -makero
```
- 5 Shutdown Oracle database instances and listener processes on each node in the cluster.
- 6 Stop the Global Services Daemon, gsd. As "oracle" user, enter:

```
$ $ORACLE_HOME/bin/psctl stop
```

- 7 Rename the SRVM configuration Location file. As "oracle" user, enter:

```
$ cd /var/opt/oracle
```



```
$ mv srvConfig.loc srvConfig.loc_backup
```
- 8 Configure Private IP addresses for CRS communication.
See [“Configuring CRS and UDP private IP addresses for failover”](#) on page 291.
- 9 Create a home directory for CRS binaries and a new home directory for Oracle RAC 11gR1 database binaries.
See [“Identify storage for the Oracle RAC 11g components”](#) on page 278.
- 10 Create a disk group, volume and file system for OCR and vote disk.
See [“Identify storage for the Oracle RAC 11g components”](#) on page 278.
- 11 Add a CFSMount and CVMVolDg resource to VCS configuration for monitoring the shared diskgroup and filesystem.
See [“Configuring VCS resources for Oracle”](#) on page 357.

Installing Oracle RAC 11gR1

Install the Oracle RAC 11g binaries.

To install Oracle RAC 11g

- 1 Install Oracle RAC 11gR1 Clusterware Software.
See [“Installing Oracle Clusterware \(CRS\)”](#) on page 321.
- 2 Verify that the CRS daemons are running. As user "oracle", enter:

```
$ $CRS_HOME/bin/crs_stat -t
```
- 3 Add the CSSD resource for monitoring the CRS daemons to VCS configuration.
See [“Configuring VCS resources for Oracle”](#) on page 357.
- 4 Install Oracle RAC 11gR1 RDBMS software. For detailed instructions:
See [“Installing Oracle RAC 11g database binaries”](#) on page 323.
- 5 Install the latest patchset. For detailed instructions:
See [“Completing Oracle RAC 11g post-installation tasks”](#) on page 329.

Migrating an existing Oracle RAC 9iR2 database to Oracle RAC 11g

To migrate an existing database from Oracle RAC 9iR2 database to Oracle RAC 11gR1, upgrade the database to Oracle RAC 11gR1.

For details:, see Oracle metalink documentation.

Performing post-upgrade tasks Oracle RAC 11gR1

To perform post-upgrade tasks

- 1 Unfreeze the VCS service groups that were frozen earlier.

See [“Performing pre-upgrade tasks for Oracle RAC 11g”](#) on page 362.

As root user, enter:

```
# haconf -makerw

# hagrps -unfreeze oracle_group -persistent

# haconf -dump -makero
```

- 2 Relink Symantec libraries. For details:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 11g”](#) on page 330.

Upgrading from Oracle RAC 10g to Oracle RAC 11gR1

The upgrade and migration procedure assumes that the beginning configuration includes the following components, and that these components are up and running on the cluster nodes:

- Storage Foundation for Oracle RAC 5.0
- Oracle RAC 10g
- Oracle RAC 10g database

The migration procedure consists of the following tasks:

- [Preparing to upgrade to Oracle RAC 11gR1](#)
- [Installing Oracle RAC 11g R1](#)
- [Migrating the existing Oracle RAC 10g database to Oracle RAC 11gR1](#)
- [Performing post-upgrade tasks](#)

The upgrade procedures for Oracle RAC 10gR1 and Oracle RAC 10gR2 are generally the same, except where specific steps are identified as applicable for a specific version of Oracle RAC.

Preparing to upgrade to Oracle RAC 11gR1

Perform the following pre-upgrade tasks before migrating from Oracle RAC 10g to Oracle RAC 11gR1.

To perform pre-upgrade tasks

- 1 Upgrade the OS and install any patches, if required. For Oracle OS requirements:

See Oracle metalink documentation.

- 2 Take a hot or cold backup of the existing database.
- 3 Take a backup of the existing Oracle home and central inventory.
- 4 Shutdown the Oracle instance.

- If Oracle is under VCS control, freeze the `oracle` group:

```
# haconf -makerw

# hagrps -freeze oracle_group -persistent

# haconf -dump -makero
```

- Use oracle commands to shutdown oracle.

- 5 Shutdown CRS.

- If CRS is under VCS

```
# haconf -makerw

# hagrps -freeze cssd_group -persistent

# haconf -dump -makero
```

- Stop CRS

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

Installing Oracle RAC 11g R1

After completing the pre-upgrade tasks, perform the upgrade procedure.

Note: Review your Oracle installation manuals and the appropriate Oracle support websites before installing Oracle RAC 11g.

To upgrade from Oracle10g to Oracle 11g R1

- 1 Install the Oracle RAC 11g R1 CRS.
See [“Installing Oracle Clusterware \(CRS\)”](#) on page 321.
- 2 Make sure that Oracle 11g R1 CRS is running.
 - To list the version of CRS software installed

```
# $ORA_CRS_HOME/bin/crsctl query crs softwareversion
```
 - To list the CRS software operating version

```
# $ORA_CRS_HOME/bin/crsctl query crs activeversion
```
- 3 Install the 11gR1 RDBMS
See [“Installing Oracle RAC 11g database binaries”](#) on page 323.

Migrating the existing Oracle RAC 10g database to Oracle RAC 11gR1

Upgrade the database to Oracle RAC 11g R1.

For details:, see Oracle metalink documentation.

Performing post-upgrade tasks

To perform post-upgrade tasks

- 1 Unfreeze the VCS service groups that were frozen earlier.
See [“Performing pre-upgrade tasks for Oracle RAC 11g”](#) on page 362.
As root user, enter:

```
# haconf -makerw  
  
# hagrps -unfreeze oracle_group -persistent  
  
# hagrps -unfreeze cssd_group -persistent  
  
# haconf -dump -makero
```
- 2 Relink Symantec libraries. For details:
See [“Relinking the SF Oracle RAC libraries to Oracle RAC 11g”](#) on page 330.

Adding and removing SF Oracle RAC nodes

- [Adding a node to SF Oracle RAC clusters using Oracle RAC 10g](#)
- [Removing a node from SF Oracle RAC clusters using Oracle RAC 10g](#)
- [Adding a node to SF Oracle RAC clusters using Oracle RAC 11g](#)
- [Removing a node from SF Oracle RAC clusters using Oracle RAC 11g](#)

Adding a node to SF Oracle RAC clusters using Oracle RAC 10g

This chapter includes the following topics:

- [About adding a node to an Oracle RAC 10g cluster](#)
- [SF Oracle RAC add node procedures](#)
- [Preparing the new node for installing Oracle](#)
- [Adding the new system to the SFDB repository](#)

About adding a node to an Oracle RAC 10g cluster

SF Oracle RAC enables you to add and remove nodes to an Oracle RAC 10g cluster.

An Oracle RAC 10g cluster can have as many as eight nodes.

The example procedure in this chapter describes how to add a single node to a two-node cluster. The example procedure adds a node (saturn) to an existing cluster (rac_cluster101) that already has two nodes (galaxy and nebula).

Note: Prior to proceeding with the tasks described in this chapter for adding a node to your existing cluster, first perform the pre-installation operations described in See [“Pre-installation tasks”](#) on page 197.

SF Oracle RAC add node procedures

The following SF Oracle RAC procedures enable you to add nodes to an Oracle RAC 10g cluster:

- Modifying `/etc/pse.conf` to enable the dlpi driver
- Checking system requirements for the new node
- Physically adding a new system to the cluster
- Installing Veritas Storage Foundation for Oracle RAC on the new system
- Starting Volume Manager
- Verifying the existing security setup on the node
- Configuring LLT, GAB, VCSMM, and VXFEN drivers
- Adding a node in a VxSS group
- Preparing the new node for installing Oracle
- Adding the new system to the SFDB repository

Modifying `/etc/pse.conf` to enable the dlpi driver

The file `/etc/pse.conf` must be configured to enable the Streams DLPI driver. The following procedure describes checking the file and modifying it, if necessary, and loading the dlpi driver.

Checking and modifying `/etc/pse.conf`

- 1 Check to see if the Ethernet driver is configured in the `/etc/pse.conf` file:

```
# egrep 'ethernet driver' /etc/pse.conf
```
- 2 In the output, examine the line containing the "ethernetdriver" expression:

```
# d+ dlpi en /dev/dlpi/en # streams dlpi ethernet driver
```
- 3 If the comment symbol ("#") precedes the line, the DLPI driver is not configured. Using vi or another text editor, edit the file:

```
# vi /etc/pse.conf
```
- 4 Find the section in the file labeled "#PSE drivers" and look for the line shown in step 2.
- 5 Uncomment the line by removing the initial "#" symbol.

- 6 Save and close the file.
- 7 To load the dlpi driver, enter:

```
# strload -f /etc/dlpi.conf
```

Checking system requirements for new node

Before installing the SF Oracle RAC software on the new node, you must verify that the new node meets the installation requirements.

Verify that the new nodes joining the cluster meet all of the requirements for installing and using SF Oracle RAC.

- The new system must have the identical operating system and patch level as the existing systems.
- Use a text window of 80 columns minimum by 24 lines minimum; 80 columns by 24 lines is the recommended size for the optimum display of the installsfrac script.

Physically adding a new system to the cluster

The new system must have the identical operating system and patch level as the existing systems. When you physically add the new system to the cluster, make sure the node:

- Is connected to the same shared storage devices as the existing nodes.
- Has private network connections to two independent switches for the cluster.

Refer to the *Veritas Cluster Server Installation Guide*.

After installing SF Oracle RAC on the new system and starting VxVM, ensure that the new system can access the same shared storage devices (including coordinator disks). If the new system does not see the same disks as the existing nodes, the new system cannot join the cluster as a new node, as indicated by a CVM error on the console.

Installing Veritas Storage Foundation for Oracle RAC on the new system

Read the pre-installation instructions in this guide before proceeding.

See [“About preparing to install and configure SF Oracle RAC”](#) on page 57.

To install Veritas Storage Foundation for Oracle RAC without configuration

- 1 Log in as root on one of the systems for installation.
- 2 Enter the following command from the top-level directory of the mounted disc:

```
# ./installsffrac -installonly [-rsh]
```

The `-installonly` option is required to perform the installation without configuring the software. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to install the software simultaneously on several systems.

Note: You must install SF Oracle RAC using the `-installonly` option as described in this step, and later manually configure the new node joining the cluster. Install the Veritas Storage Foundation for Oracle RAC software as described previously in this guide, but run the product installation script instead of the generic installer script.

- 3 After the script performs initial checks, confirm to start the installation.
- 4 Review the output as the script checks system licensing and installs the licensing fileset.
- 5 Enter the license key as the installer prompts.

```
Enter a SFRAC license key for saturn: [?] XXXX-XXXX-XXXX-XXXX-XX
XXXX-XXXX-XXXX-XXXX-XX successfully registered on saturn
SFRAC license registered on saturn
```

- 6 Enter keys for additional product features such as VVR, if you want to set up a Global Cluster environment.

```
Do you want to enter another license key for saturn? [y,n,q,?]
(n)
```

- 7 Review the output as the script lists the filesets and patches to install and checks whether any of them are present on the node.
- 8 After the requirements checks are complete, press **Return** to start installing the packages. If you are installing multiple nodes, you have the option of installing them simultaneously. You will be prompted after the installation is complete.

9 After the script installs the packages and patches, note the location of the summary, log, and response files in the output.

10 Ignore the message advising you to run `installsfrac -configure`.

When adding a node to a cluster running Veritas Storage Foundation for Oracle RAC, you must manually configure the system using the following procedure.

See [“To start Volume Manager”](#) on page 375.

11 Before proceeding to the following procedure, you must restart the new node.

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

Starting Volume Manager

As you run the utility, answer `n` to prompts about licensing. You installed the appropriate license when you ran the `installsfrac` utility.

To start Volume Manager

1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
# vxinstall
```

2 Decline to set up a systemwide disk group for the system by entering `n`. 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.

Verifying the existing security setup on the node

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 secure mode, then proceed with configuring LLT and GAB.

See [“Configuring LLT, GAB, VCSMM, and VXFEN drivers”](#) on page 377.

Note: The procedures in this section use as an example "node east" to help describe what is required.

To verify the existing security setup on the node

- 1 If node east is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Otherwise, proceed to configuring the authentication broker on node east.

See [“Configuring the authentication broker on node east”](#) on page 376.

- 2 Find out the root broker to which the node east belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node east already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.

See [“Setting up VCS related security configuration”](#) on page 377.

- 4 If the node east belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node east.

- Kill /opt/VRTSat/bin/vxatd process.
- Remove the credential that RB2 has given to AB on node east.

```
# vssat deletcred --domain type:domainname \
--prplname prplname
```

Configuring the authentication broker on node east

Configure a new authentication broker (AB) on node east. This AB belongs to root broker RB1.

To configure the authentication broker on node east

- 1 Create a principal for node east on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \
--prplname prplname --password password \
--prpltype service
```

- 2 Ensure that there is no clock skew between the times on node east and RB1.
- 3 Copy the /opt/VRTSat/bin/root_hash file from RB1 to node east.

- 4 Configure AB on node east to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \  
rootbroker -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 VCS related security configuration

Perform the following steps to configure VCS related security settings.

Setting up VCS related security configuration

- 1 Start /opt/VRTSat/bin/vxatd process.
- 2 Create HA_SERVICES domain for VCS.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

- 3 Add VCS and webserver principal to AB on node east.

```
# 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
```

Configuring LLT, GAB, VCSMM, and VXFEN drivers

Configure the LLT, GAB, VCSMM, and VXFEN drivers.

To configure LLT, GAB, VCSMM, and VXFEN drivers

- 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 en1 /dev/dlpi/en:1 - ether - -

link en2 /dev/dlpi/en: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 system. 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 If you are adding the new node to a single node cluster, then fencing must be enabled and configured on the original node before proceeding to step 7.

Note: For I/O fencing information, refer to the earlier section on Setting up I/O fencing.

- 7** Set up the `/etc/vcsmmtab` and `/etc/vxfendg` files on the new system by copying them from one of the other existing nodes:

```
# rcp galaxy:/etc/vcsmmtab /etc
# rcp galaxy:/etc/vxfendg /etc
# rcp galaxy:/etc/vxfenmode /etc
```

- 8** Run the commands to start LLT and GAB on the new node:

```
# /etc/init.d/llt.rc start
# /etc/init.d/gab.rc start
```

- 9** On the new node, start the VXFEN, VCSMM, and LMX drivers. First rcp the `/etc/vxfenmode` file to enable fencing mode.

Use the commands in the order shown:

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

- 10** On the new node, start the GMS and ODM drivers. Use the commands in the order shown:

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

- 11** On the new node, verify that the GAB port memberships are a, b, d, and o. Run the command:

```
# /sbin/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
```

Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

Adding a node in a VxSS group

- 1 On the existing system, galaxy run the following command:

```
# haconf -dump -makero  
  
# hastop -all -force
```

- 2 Add the node saturn to the existing VxSS group.

Modify the main.cf to add new node to the group VxSS using vi or another text editor. The main.cf located at: /etc/VRTSvcs/conf/config.

For example:

```
group VxSS (  
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }  
    Parallel = 1  
    AutoStartList = { galaxy, nebula, saturn }  
    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  
//     }
```

- 3 Start VCS on galaxy, by running the following command:

```
# hstart
```

- 4 Start VCS on all the other nodes.

Run the `hstart` command on all the other nodes:

```
# hstart
```

- 5 Make sure the VxSS group is running.

Run the following command on all nodes.

```
# hagrps -state VxSS
```

- 6 Save the configuration by running the following command from any node.

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

VxSS is enabled on all nodes.

Preparing the new node for installing Oracle

Performing pre-installation operations involves both manual and automated tasks from the SF Oracle RAC configuration program. Before installing Oracle RAC 10g you must perform the following additional tasks:

- Setting the system parameters
- Creating Oracle user and groups
- Setting up Oracle user equivalence for RSH and RCP
- Setting Oracle Clusterware oracle user capabilities
- Configuring shell limits for Oracle user
- Configuring kernel parameter for processes per user
- Tuning the Virtual Memory Manager (VMM)
- Verifying the user "nobody" exists
- Configuring private IP addresses for CRS
- Identifying the public virtual IP addresses for use by Oracle
- Creating Oracle disk groups, volumes, and mount points
- Preparing `$CRS_HOME` and `$ORACLE_HOME` on new node

- Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node
- Configuring CVM
- Linking Veritas membership library

After performing the preparatory task for installing Oracle on a new node, the following additional procedures must be performed:

- Using the Oracle RAC 10g add node procedure
- Relinking the SF Oracle RAC libraries to Oracle RAC 10g

Setting the system parameters

To prepare the new node for installing Oracle, you must set the system parameters, including the OS kernel and networking parameters.

Refer to the latest Oracle documentation for information about the system parameters and procedures. Restart the nodes for the new values to take effect .

Creating Oracle user and groups

To prepare the new node for installing Oracle, you must create the oracle user and groups.

Refer to the following section for information:

See [“Creating Oracle user and groups”](#) on page 201.

Setting up Oracle user equivalence for RSH and RCP

To prepare the new node for installing Oracle, you must set up Oracle user equivalence for RSH and RCP.

Refer to the following section for information:

See [“Setting up Oracle user equivalence for RSH and RCP”](#) on page 206.

Setting Oracle Clusterware oracle user capabilities

To prepare the new node for installing Oracle, you must set up the Oracle clusterware oracle user capabilities.

Refer to the following section for information:

See [“Setting Oracle Clusterware oracle user capabilities”](#) on page 207.

Configuring shell limits for Oracle user

To prepare the new node for installing Oracle, you must configure shell limits for Oracle user.

Refer to the following section for information:

See [“Configuring shell limits for Oracle user”](#) on page 207.

Configuring kernel parameter for processes per user

To prepare the new node for installing Oracle, you must configure kernel parameter for processes per user.

Refer to the following section for information:

See [“Configuring kernel parameter for processes per user”](#) on page 208.

Tuning the Virtual Memory Manager (VMM)

To prepare the new node for installing Oracle, you must tune the Virtual Memory Manager (VMM).

Refer to the following section for information:

See [“Tuning the Virtual Memory Manager \(VMM\)”](#) on page 209.

Verifying the user "nobody" exists

To prepare the new node for installing Oracle, you must verify that the user "nobody" exists.

Refer to the following section for information:

See [“Verifying the user "nobody" exists”](#) on page 209.

Configuring private IP addresses for CRS

To prepare the new node for installing Oracle, you must configure private IP addresses for CRS.

The CRS daemon requires a private IP address on each node to enable communications and heartbeating. Determine a private NIC device for which LLT is configured.

Access and view the file /etc/llttab. For example, if a device is used as an LLT interconnect on one system, you can configure an available IP address for it.

Table 18-1 Required information to configure private IP addresses for CRS

Required information	Example
On galaxy:	192.168.12.1
On nebula:	192.168.12.2
On saturn:	192.168.12.7

To configure private IP addresses for CRS

- 1
- Make a backup copy of the main.cf file. For example:
- ```
cd /etc/VRTSvcs/conf/config
cp main.cf main.cf.2node
```
- 2
- Use the following commands to add private IP into ora\_priv resource on the active node:

```
haconf -makerw
hasys -add saturn
hagrps -modify cvm SystemList -add saturn 2
hagrps -modify cvm AutoStartList -add saturn
hares -modify ora_priv Device -add en1 0 -sys Saturn
hares -modify ora_priv Device -add en2 1 -sys saturn
hares -modify ora_priv Address "192.168.12.7" -sys saturn
haconf -dump -makero
```

## Identifying the public virtual IP addresses for use by Oracle

To prepare the new node for installing Oracle, you must identify the public virtual IP addresses for use by Oracle.

Refer to the following section for information:

See [“Identifying the public virtual IP addresses for use by Oracle”](#) on page 215.



## Creating Oracle disk groups, volumes, and mount points

To prepare the new node for installing Oracle, you must create the Oracle disk groups, volumes, and mount points.

Refer to the following section for information:

See [“Creating Oracle disk groups, volumes, and mount points”](#) on page 216.

## Preparing \$CRS\_HOME and \$ORACLE\_HOME on new node

To prepare the new node for installing Oracle, you must prepare \$CRS\_HOME and \$ORACLE\_HOME on the new node.

The local or cluster file system on shared storage should be the same as on the existing nodes.

### Local file system for \$CRS\_HOME and \$ORACLE\_HOME

Follow the directions for the following procedures to prepare \$CRS\_HOME and \$ORACLE\_HOME on new node:

- Create a file system on local storage for Oracle/CRS binaries (/app):  
See [“To create a file system on local storage for Oracle/CRS binaries \(/app\)”](#) on page 386.
- Prepare \$CRS\_HOME on the new node:  
See [“To prepare \\$CRS\\_HOME on the new node”](#) on page 387.
- Prepare \$ORACLE\_BASE on the new node:  
See [“To prepare \\$ORACLE\\_BASE on the new node”](#) on page 387.
- Prepare \$ORACLE\_HOME on the new node:  
See [“To prepare \\$ORACLE\\_HOME on the new node”](#) on page 387.

Identify the directories required for installing Oracle and CRS software.

**Table 18-2** Required information \$CRS\_HOME on each system

| Required information                                             | Example                                                      |
|------------------------------------------------------------------|--------------------------------------------------------------|
| Disk on each node for Oracle and CRS for Oracle and CRS binaries | Disk 1                                                       |
| Disk group on each local storage                                 | On saturn: orabindg_saturn                                   |
| Volume for each local disk group                                 | On saturn: orabinvol_saturn                                  |
| File system on each local volume                                 | vol_saturn<br>/dev/vx/rdisk/orabindg_saturn/orabinvol_saturn |

**Table 18-2** Required information \$CRS\_HOME on each system *(continued)*

| Required information                      | Example             |
|-------------------------------------------|---------------------|
| Mount point for each local file system    | /app                |
| Directory to install CRS                  | /app                |
| Directory to install Oracle (ORACLE_HOME) | /app/oracle/orahome |

**To create a file system on local storage for Oracle/CRS binaries (/app)**

- 1
- As root user, first create a VxVM local diskgroup, orabindg\_hostname:  
  
# vxdbg init orabindg\_saturn Disk\_1
- 2
- Create a volume, orabinvol\_hostname:  
  
# vxassist -g orabindg\_saturn make orabinvol\_saturn 12G
- 3
- Create directory, /app  
  
# mkdir /app
- 4
- Create a filesystem with this volume, orabinvol\_hostname  
  
# mkfs -V vxfs /dev/vx/dsk/orabindg\_saturn/orabinvol\_saturn
- 5
- Mount /app:  
  
# mount -V vxfs /dev/vx/dsk/orabindg\_saturn/orabinvol\_saturn /app
- 6
- Add an entry for this filesystem. For example, in the /etc/filesystems add:  
  
/app:  
dev = dev/vx/dsk/orabindg\_saturn/orabinvol\_saturn  
vfs = vxfs  
mount = true  
options = rw  
account = false

### To prepare \$CRS\_HOME on the new node

- 1 On each system, log in as "root".
- 2 Create the directory for CRS\_HOME:  
  
# **mkdir -p /app/crshome**
- 3 Change ownership and permissions.  
  
# **chown -R oracle:oinstall /app/crshome**  
  
# **chmod -R 744 /app/crshome**

### To prepare \$ORACLE\_BASE on the new node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE\_BASE:  
  
# **mkdir -p /app/oracle**
- 3 Change ownership and permissions.  
  
# **chown -R oracle:oinstall /app/oracle**  
  
# **chmod -R 744 /app/oracle**

### To prepare \$ORACLE\_HOME on the new node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE\_HOME:  
  
# **mkdir -p /app/oracle/orahome**
- 3 Change ownership and permissions:  
  
# **chown -R oracle:oinstall /app/oracle/orahome**  
  
# **chmod -R 744 /app/oracle/orahome**

## Cluster file system for \$CRS\_HOME and \$ORACLE\_HOME

Create the Oracle base mount point:

```
mkdir -p /app
```

## Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node

To prepare the new node for installing Oracle, you must prepare the OCR and VOTE-disk mount point and the Oracle database mount point on the new node.

Create the mount point for the CRS file system having the OCR and Vote Disk files.

For example:

```
mkdir -p /ocrvote
```

Create the mount point for the Oracle Database.

For example:

```
mkdir -p /oradata
```

## Configuring CVM

To prepare the new node for installing Oracle, you must configure CVM.

### Configuring CVM

- 1 To enable the existing cluster to recognize the new node, execute on one of existing nodes:

```
haconf -makerw

hares -modify cvm_clus CVMNodeId -add saturn 2

haconf -dump -makero

/etc/vx/bin/vxclustadm -m vcs -t gab reinit

/etc/vx/bin/vxclustadm nidmap
```

- 2 Repeat the following step in the remaining existing cluster:

```
/etc/vx/bin/vxclustadm -m vcs -t gab reinit

/etc/vx/bin/vxclustadm nidmap
```

- 3 Start CVM on the newly added node.

```
hastart
```

- 4 Verify that the CVM group has come online on the newly added node.

```
hastatus -sum
```

## Linking Veritas membership library

To prepare the new node for installing Oracle, you must link the Veritas membership library.

Perform the following steps to link the Veritas membership library.

### To link Veritas Membership library

- 1 Create required sub-directories if they not present.

For example, enter the following commands:

```
mkdir -p /usr/sbin/cluster/utilities
echo `lltstat -C` > /usr/sbin/cluster/utilities/cldomain
mkdir /opt/ORCLcluster
mkdir /opt/ORCLcluster/lib
mkdir /opt/ORCLcluster/lib32
mkdir /opt/ORCLcluster/bin
```

- 2 Archive skgxn library.

For example, enter the following commands:

```
cd /opt/VRTSvcs/rac/lib
rm libskgxn2.a
cp libvcsmm.so shr_skgxn2.o
ar -X32 -qv libskgxn2.a shr_skgxn2.o
rm shr_skgxn2.o
cd /opt/VRTSvcs/rac/lib64
rm libskgxn2.a
cp libvcsmm.so shr_skgxn2.o
ar -X64 -qv libskgxn2.a shr_skgxn2.o
rm shr_skgxn2.o
```

### 3 Copy the skgxn library and clsinfo to /opt/ORCLcluster directory.

For example, enter the following commands:

```
cp /opt/VRTSvcs/rac/bin/clsinfo \
/opt/ORCLcluster/bin

cp /opt/VRTSvcs/rac/lib/libskgxn2.a \
/opt/ORCLcluster/lib32

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

cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
/opt/ORCLcluster/lib

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

## Using the Oracle RAC 10g add node procedure

After performing the preparatory tasks for installing Oracle on a new node, follow the Oracle RAC 10g add node procedure.

For the Oracle RAC 10g procedure for adding a node, see: Metalink Article 270512.1, Adding a Node to a 10g RAC Cluster

In this procedure, Oracle copies the CRS\_HOME and ORACLE\_HOME from an existing node in the cluster .

## Relinking the SF Oracle RAC libraries to Oracle RAC 10g

After the CRS and Oracle binaries are installed on the new node, relink the Veritas library on \$ORACLE\_HOME.

For information about this procedure, refer to:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

## Adding the new system to the SFDB repository

Add the new system to the Storage Foundation database repository using the following procedure.

### Add the new system to the SFDB repository

- 1 Add the system using the following sequence of commands:

```
haconf -makerw

hagrpl -modify Sfua_Base SystemList -add saturn

hares -modify sfua_ip Device en3 -sys saturn

haconf -dump -makero
```

- 2 Copy the /etc/vx/vxldb/.odbc.ini file from an existing node to the new system using a remote file copy utility such as rcp, tcp, or scp.

For example, to use rcp:

```
rcp /etc/vx/vxldb/.odbc.ini saturn:/etc/vx/vxldb
```

- 3 Copy the /etc/vx/vxldb/admin.properties file from an existing node to the new system using a remote file copy utility such as rcp, tcp, or scp.

For example, to use rcp:

```
rcp /etc/vx/vxldb/admin.properties saturn:/etc/vx/vxldb
```

- 4 Make a repository mount point on the newly added node using the following commands:

```
grep sfua_db_mountpoint /etc/vx/vxldb/admin.properties sfua_db_mountpoint=/sfua_rep

mkdir /sfua_rep
```



# Removing a node from SF Oracle RAC clusters using Oracle RAC 10g

This chapter includes the following topics:

- [About removing a node from an Oracle RAC 10g cluster](#)
- [Using the Oracle remove node procedure](#)
- [Removing SF Oracle RAC using the `uninstallsrac` utility](#)


## About removing a node from an Oracle RAC 10g cluster

You can remove one or more nodes from an Oracle RAC 10g cluster by using the procedures described in this chapter.

To remove a node from a Veritas Storage Foundation for Oracle RAC cluster, you must follow the Oracle remove node procedure, as well as the procedure using the `uninstallsrac` utility and subsequent editing of the configuration files.

[Figure 19-1](#) displays the steps required to remove a node from a Veritas Storage Foundation for Oracle RAC cluster. Each step is discussed in this chapter.

**Figure 19-1** Removing a node from an Oracle RAC 10g cluster



## Using the Oracle remove node procedure

Use the Oracle remove node procedure, to remove the node from an Oracle RAC 10g cluster.

For the Oracle procedure for removing a node, refer to your Oracle documentation and the appropriate Oracle web sites.

Follow the instructions provided by Oracle.

## Removing SF Oracle RAC using the uninstallsfrac utility

Run the uninstallsfracscript from any node in the cluster, including a node from which you are uninstalling Veritas Storage Foundation for Oracle RAC. Before invoking the uninstallsfrac script, you must take all service groups offline and shut down VCS.

The example procedure involves removing SF Oracle RAC from the node saturn.

## To remove SF Oracle RAC using the uninstallsfrac utility

- 1 Offline all Oracle parallel groups.

```
hagrpf -offline oracle_group -sys saturn
```

- 2 Offline CRS.

```
hares -offline cssd_resource -sys saturn
```

If CRS is not under control.

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

- 3 Stop VCS on the node on which you want to remove SF Oracle RAC.

```
/opt/VRTSvcs/bin/hastop -local
```

- 4 As superuser, start the uninstallation process from any node from which you are uninstalling Veritas Storage Foundation for Oracle RAC.

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

```
./uninstallsfrac
```

- 5 Decline to uninstall SF Oracle RAC from all systems.

VCS configuration files exist on this system with the following information:

```
Cluster Name: rac_cluster101
Cluster ID Number: 101
Systems: galaxy nebula saturn
Service Groups: cvm oradb1_grp
```

Do you want to uninstall SFRAC from these systems? [y,n,q] (y) **n**

---

**Warning:** Be sure to answer "n". Otherwise, the utility begins the process of uninstalling SF Oracle RAC from all systems.

---

- 6 Specify the name of the system on which you are uninstalling SF Oracle RAC.

- 7 The installer prompts you to specify the name of the system from which you are uninstalling Veritas Storage Foundation for Oracle RAC:

Enter the system names separated by spaces on which to uninstall SFRAC:**saturn**

- 8 Review the output as the script checks for packages currently installed on your system. The utility also checks for dependencies between packages to determine the packages it can safely uninstall and in which order.
- 9 Confirm the uninstall when the uninstaller has completed checking.
- 10 Review the output as the script stops processes and drivers running on each node.
- 11 After the script uninstalls the packages, make a note of the location of the summary and log files that the uninstaller creates.

## Editing VCS configuration files on existing nodes

After running `uninstallsrac`, modify the configuration files on the existing remaining nodes to remove references to the deleted nodes. For an example `main.cf`:

See [“Sample main.cf for adding and removing an Oracle RAC 10g node”](#) on page 549.

The following procedures must be performed to edit the VCS configuration files on the existing remaining nodes:

- [Editing /etc/llthosts](#)
- [Editing /etc/gabtab](#)
- [Modifying the VCS configuration to remove the node](#)

### Editing /etc/llthosts

On each of the existing nodes, use `vi` or another editor to edit `/etc/llthosts`. Remove lines corresponding to the removed nodes; for example, if `saturn` is the node removed from the cluster, remove the line "2 saturn" from the file:

```
1 galaxy
2 nebula
3 saturn
```

Change to:

```
1 galaxy
2 nebula
```

## Editing `/etc/gabtab`

In the file `/etc/gabtab`, change the command contained in the file to reflect the number of systems after the node is removed:

```
/sbin/gabconfig -c -nN
```

where `N` is the number of nodes remaining.

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 using one of the following three methods:

- Directly editing `/etc/VRTSvcs/conf/config/main.cf`
- Using the VCS Cluster Manager
- Using the command line, as illustrated in the following example

---

**Note:** For details on configuring VCS, refer to *Veritas Cluster Server User's Guide*.

---

At this point in the process, all Oracle binaries have been removed from the system to be deleted. The instance has been removed from the database, that is, the thread disabled, and the `spfile<SID>.ora` edited by Oracle to remove any references to this instance. The next step is to remove all references in the `main.cf` to the deleted nodes.

As root user execute the following procedure on the CVM master node only.

Perform the following procedure to modify the CVM group in the `main.cf` file.

### Modify the VCS configuration to remove the node

- 1 To determine the CVM master node execute:

```
vxdctl -c mode
```

- 2 Make a backup copy of the `main.cf` file.

```
cd /etc/VRTSvcs/conf/config
```

```
cp main.cf main.cf.3node.bak
```

- 3 Use the following commands to reconfigure the VxSS group. Execute the following commands:

```
haconf -makerw
```

```
hagrps -modify VxSS SystemList -delete saturn
```

```
haconf -dump -makero
```

- 4 Use the following commands to reconfigure the CVM group. Execute:

```
haconf -makerw
```

```
hares -modify cvm_clus CVMNodeId -delete saturn
```

```
hagrps -modify cvm SystemList -delete saturn
```

```
hasys -delete saturn
```

```
haconf -dump -makero
```

Example of `main.cf` file:

See [“Sample main.cf for adding and removing an Oracle RAC 10g node”](#) on page 549.

- 5 Verify the syntax of `main.cf` file:

```
hacf -verify .
```

The `main.cf` file now should not contain entries for system saturn.

- 6 Copy the new version of the `main.cf` to each system in the cluster.

```
rcp (or scp) main.cf galaxy:/etc/VRTSvcs/conf/config
```

```
rcp (or scp) main.cf nebula:/etc/VRTSvcs/conf/config
```

# Adding a node to SF Oracle RAC clusters using Oracle RAC 11g

This chapter includes the following topics:

- [About adding a node to an Oracle RAC 11g cluster](#)
- [Using the SF Oracle RAC add node procedures](#)
- [Preparing the new node for installing Oracle](#)
- [Adding the new system to the SFDB repository](#)

## About adding a node to an Oracle RAC 11g cluster

SF Oracle RAC enables you to add and remove nodes to an Oracle RAC 11g cluster.

An SF Oracle RAC cluster can have as many as eight nodes.

The example procedure in this chapter describes how to add a node to a two-node cluster. The example procedure adds a node (saturn) to an existing cluster (rac\_cluster101) with two nodes (galaxy and nebula).

## Using the SF Oracle RAC add node procedures

The following SF Oracle RAC procedures enable you to add nodes to an Oracle RAC 11g cluster:

- Modifying /etc/pse.conf to enable the dlpi driver
- Checking system requirements for new node

- Physically adding a new system to the cluster
- Installing Veritas Storage Foundation for Oracle RAC on the new system
- Starting Volume Manager
- Verifying the existing security setup on the node
- Configuring LLT, GAB, VCSMM, and VXFEN drivers
- Adding a node in a VxSS group
- Preparing the new node for installing Oracle
- Adding the new system to the SFDB repository

## Modifying /etc/pse.conf to enable the dlpi driver

The file /etc/pse.conf must be configured to enable the Streams DLPI driver. The following procedure describes checking the file and modifying it, if necessary, and loading the dlpi driver.

### Checking and modifying /etc/pse.conf

- 1 Check to see if the Ethernet driver is configured in the /etc/pse.conf file:  

```
egrep 'ethernet driver' /etc/pse.conf
```
- 2 In the output, examine the line containing the "ethernet driver" expression:  

```
d+ dlpi en /dev/dlpi/en # streams dlpi ethernet driver
```
- 3 If the comment symbol("#") precedes the line, the DLPI driver is not configured. Using vi or another text editor, edit the file:  

```
vi /etc/pse.conf
```
- 4 Find the section in the file labeled "#PSE drivers" and look for the line shown in step 2.
- 5 Uncomment the line by removing the initial "#" symbol.
- 6 Save and close the file.
- 7 To load the dlpi driver, enter:  

```
strload -f /etc/dlpi.conf
```



## Checking system requirements for new node

Before installing the SF Oracle RAC software on the new node, you must verify that the new node meets the installation requirements.

Verify that the new nodes joining the cluster meet all of the requirements for installing and using SF Oracle RAC.

- The new system must have the identical operating system and patch level as the existing systems.
- Use a text window of 80 columns minimum by 24 lines minimum; 80 columns by 24 lines is the recommended size for the optimum display of the `installsf` script.

## Physically adding a new system to the cluster

The new system must have the identical operating system and patch level as the existing systems. When you physically add the new system to the cluster, make sure the node:

- Is connected to the same shared storage devices as the existing nodes.
- Has private network connections to two independent switches for the cluster.

Refer to the *Veritas Cluster Server Installation Guide*.

After installing SF Oracle RAC on the new system and starting VxVM, the new system can access the same shared storage devices (including coordinator disks). If the new system does not see the same disks as the existing nodes, the new system cannot join the cluster as a new node, as indicated by a CVM error on the console.

## Installing Veritas Storage Foundation for Oracle RAC on the new system

Read the pre-installation instructions in this guide before proceeding.

See [“About preparing to install and configure SF Oracle RAC”](#) on page 57.

### To install Veritas Storage Foundation for Oracle RAC without configuration

- 1 Log in as root on one of the systems for installation.
- 2 Enter the following command from the top-level directory of the mounted disc:

```
./installsffrac -installonly [-rsh]
```

The `-installonly` option is required to perform the installation without configuring the software. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to install the software simultaneously on several systems.

- 3 After the script performs initial checks, confirm to start the installation.
- 4 Review the output as the script checks system licensing and installs the licensing fileset.
- 5 Enter the license key as the installer prompts.

```
Enter a SFRAC license key for saturn: [?] XXXX-XXXX-XXXX-XXXX-XX
XXXX-XXXX-XXXX-XXXX-XX successfully registered on saturn
SFRAC license registered on saturn
```

- 6 Enter keys for additional product features such as VVR, if you want to set up a Global Cluster environment.

```
Do you want to enter another license key for saturn? [y,n,q,?]
(n)
```

- 7 Review the output as the script lists the filesets and patches to install and checks whether any of them are present on the node.
- 8 After the requirements checks are complete, press **Return** to start installing the packages. If you are installing multiple nodes, you have the option of installing them simultaneously. You will be prompted after the installation is complete.
- 9 After the script installs the packages and patches, note the location of the summary, log, and response files in the output.

- 10 Ignore the message advising you to run `installsfrac -configure`.

When adding a node to a cluster running Veritas Storage Foundation for Oracle RAC, you must manually configure the system using the following procedure.

See [“To start Volume Manager”](#) on page 375.

- 11 Before proceeding to the following procedure, you must restart the new node.

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

You must install SF Oracle RAC using the `-installonly` option and later manually configure the new node joining the cluster. Install the Veritas Storage Foundation for Oracle RAC software as described previously in this guide, but run the product installation script instead of the generic installer script.

## Starting Volume Manager

As you run the utility, answer `n` to prompts about licensing. You installed the appropriate license when you ran the `installsfrac` utility.

### To start Volume Manager

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
vxinstall
```

- 2 Decline to set up a systemwide disk group for the system by entering `n`. 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.

## Verifying the existing security setup on the node

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.

See [“Configuring LLT, GAB, VCSMM, and VXFEN drivers”](#) on page 405.

---

**Note:** The procedures in this section use as an example "node east" to help describe what is required.

---

### To verify the existing security setup on the node

- 1 If node east 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 east.  
See [“Configuring the authentication broker on node east”](#) on page 404.
- 2 Find out the root broker to which the node east belongs using the following command.

```
vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node east already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.  
See [“Setting up VCS related security configuration”](#) on page 405.
- 4 If the node east belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node east.

- Kill /opt/VRTSat/bin/vxatd process.
- Remove the credential that RB2 has given to AB on node east.

```
vssat deletcred --domain type:domainname \
--prplname prplname
```

## Configuring the authentication broker on node east

Configure a new authentication broker (AB) on node east. This AB belongs to root broker RB1.

### To configure the authentication broker on node east

- 1 Create a principal for node east on root broker RB1. Execute the following command on root broker RB1.

```
vssat addprpl --pdrtype root --domain domainname \
--prplname prplname --password password \
--prpltype service
```

- 2 Ensure that there is no clock skew between the times on node east and RB1.
- 3 Copy the /opt/VRTSat/bin/root\_hash file from RB1 to node east.

#### 4 Configure AB on node east to talk to RB1.

```
vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -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 VCS related security configuration

Perform the following steps to configure VCS related security settings.

#### Setting up VCS related security configuration

##### 1 Start /opt/VRTSat/bin/vxatd process.

##### 2 Create HA\_SERVICES domain for VCS.

```
vssat createpd --pdrtype ab --domain HA_SERVICES
```

##### 3 Add VCS and webserver principal to AB on node east.

```
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
```

## Configuring LLT, GAB, VCSMM, and VXFEN drivers

Configure LLT, GAB, VCSMM, and VXFEN drivers.

### To configure LLT, GAB, VCSMM, and VXFEN drivers

- 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 en1 /dev/dlpi/en:1 - ether - -

link en2 /dev/dlpi/en: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 system. 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 If you are adding the new node to a single node cluster, then fencing must be enabled and configured on the original node before proceeding to step 7.

---

**Note:** For I/O fencing information, refer to the earlier section on Setting up I/O fencing.

---

- 7** Set up the `/etc/vcsmmtab` and `/etc/vxfendg` files on the new system by copying them from one of the other existing nodes:

```
rcp galaxy:/etc/vcsmmtab /etc
rcp galaxy:/etc/vxfendg /etc
rcp galaxy: /etc/vxfenmode /etc
```

- 8** Start LLT and GAB on the new node:

```
/etc/init.d/llt.rc start
/etc/init.d/gab.rc start
```

- 9** On the new node, start the VXFEN, VCSMM, and LMX drivers. Remove the `/etc/vxfenmode` file to enable fencing. Use the commands in the order shown:

```
/etc/init.d/vxfen.rc stop
/etc/init.d/vxfen.rc start
/etc/init.d/vcsmm.rc start
/etc/init.d/lmx.rc start
```

- 10** On the new node, start the GMS and ODM drivers. Use the commands in the order shown:

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

- 11** On the new node, verify that the GAB port memberships are a, b, d, and o. Run the command:

```
/sbin/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
```

## Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

### Adding a node in a VxSS group

- 1 On the existing system, galaxy run the following commands:

```
haconf -dump -makero

hastop -all -force
```

- 2 Add the node saturn to the existing VxSS group.

Modify the main.cf to add new node to the group VxSS using vi or another text editor. The main.cf located at: /etc/VRTSvcs/conf/config.

For example:

```
group VxSS (
 SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
 Parallel = 1
 AutoStartList = { galaxy, nebula, saturn }
 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
// }
```



- 3 Start VCS on galaxy, by running the following command:

```
hstart
```

- 4 Start VCS on all the other nodes.

Run the `hstart` command on all the other nodes:

```
hstart
```

- 5 Make sure the VxSS group is running.

Run the following command on all nodes.

```
hagrps -state VxSS
```

- 6 Save the configuration by running the following command from any node.

```
haconf -dump -makero
```

VxSS is enabled on all nodes.

## Preparing the new node for installing Oracle

Performing pre-installation operations involves both manual and automated tasks from the SF Oracle RAC configuration program. Before installing Oracle RAC 11g you must perform the following additional tasks:

- Setting the system parameters
- Creating Oracle user and groups
- Setting up Oracle user equivalence for RSH and RCP
- Setting Oracle Clusterware oracle user capabilities
- Configuring shell limits for Oracle user
- Configuring kernel parameter for processes per user10g
- Tuning the Virtual Memory Manager (VMM)
- Verifying the user "nobody" exists
- Configuring private IP addresses for CRS
- Identifying the public virtual IP addresses for use by Oracle
- Creating Oracle disk groups, volumes, and mount points
- Preparing `$CRS_HOME` and `$ORACLE_HOME` on new node

- Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node
- Configuring CVM
- Linking Veritas membership library

After performing the above preparatory task for installing Oracle on a new node, the following additional procedures must be performed:

- Using the Oracle RAC 11g add node procedure
- Relinking the SF Oracle RAC libraries to Oracle

## Setting the system parameters

To prepare the new node for installing Oracle, you must set the system parameters, including the OS kernel and networking parameters.

Refer to the latest Oracle documentation for information about the system parameters and procedures. Restart the nodes for the new values to take effect.

## Creating Oracle user and groups

To prepare the new node for installing Oracle, you must create the oracle user and groups.

Refer to the following section for information:

See [“Creating Oracle user and groups”](#) on page 281.

## Setting up Oracle user equivalence for RSH and RCP

To prepare the new node for installing Oracle, you must set up Oracle user equivalence for RSH and RCP.

Refer to the following section for information:

See [“Setting up Oracle user equivalence for RSH and RCP”](#) on page 286.

## Setting up Oracle Clusterware oracle user capabilities

To prepare the new node for installing Oracle, you must set up the Oracle clusterware oracle user capabilities.

Refer to the following section for information:

See [“Setting Oracle Clusterware oracle user capabilities”](#) on page 288.

## Configuring shell limits for Oracle user

To prepare the new node for installing Oracle, you must configure shell limits for Oracle user.

Refer to the following section for information:

See [“Configuring shell limits ”](#) on page 288.

## Configuring kernel parameters for processes per user

To prepare the new node for installing Oracle, you must configure kernel parameter for processes per user.

Refer to the following section for information:

See [“Configuring the user process parameter”](#) on page 289.

## Tuning the Virtual Memory Manager (VMM)

To prepare the new node for installing Oracle, you must tune the Virtual Memory Manager (VMM).

Refer to the following section for information:

See [“Tuning the Virtual Memory Manager \(VMM\)”](#) on page 290.

## Verifying the user "nobody" exists

To prepare the new node for installing Oracle, you must verify that the user "nobody" exists.

Refer to the following section for information:

See [“Verifying the user "nobody" exists”](#) on page 209.

## Configuring private IP addresses for CRS

To prepare the new node for installing Oracle, you must configure private IP addresses for CRS.

The CRS daemon requires a private IP address on each node to enable communications and heartbeating. Determine a private NIC device for which LLT is configured.

Access and view the file /etc/llttab. For example, if a device is used as an LLT interconnect on one system, you can configure an available IP address for it.

**Table 20-1** Required information to configure private IP addresses for CRS

| Required information | Example                                   |
|----------------------|-------------------------------------------|
| On galaxy:           | 192.168.1.1<br>192.168.1.2<br>192.168.2.1 |
| On nebula:           | 192.168.1.3<br>192.168.1.4<br>192.168.2.2 |
| On saturn:           | 192.168.1.5<br>192.168.1.6<br>192.168.2.3 |

**To configure private IP addresses for CRS**

- 1
- Make a backup copy of the main.cf file. For example:  
  
# cd /etc/VRTSvcs/conf/config  
  
# cp main.cf main.cf.2node
- 2
- Use the following commands to add private IP into multi\_priv resource on the active node:  
  
# haconf -makerw  
  
# hasys -add saturn  
  
# hagrps -modify cvm SystemList -add saturn 2  
  
# hagrps -modify cvm AutoStartList -add Saturn  
  
# hares -modify multi\_priv Device -add en1 0 -sys saturn  
  
# hares -modify multi\_priv Device -add en2 1 -sys Saturn  
  
# hares -modify multi\_priv Address -add 192.168.1.5 0 -sys Saturn  
  
# hares -modify multi\_priv Address -add 192.168.1.6 0 -sys saturn  
  
# hares -modify multi\_priv Address -add 192.168.2.3 1 -sys saturn  
  
# haconf -dump -makero

## Identifying the public virtual IP addresses for use by Oracle

To prepare the new node for installing Oracle, you must identify the public virtual IP addresses for use by Oracle.

Refer to the following section for information:

See [“Identifying the public virtual IP addresses for use by Oracle”](#) on page 301.

## Creating Oracle disk groups, volumes, and mount points

To prepare the new node for installing Oracle, you must create the Oracle disk groups, volumes, and mount points.

Refer to the following section for information:

See [“Creating Oracle disk groups, volumes, and mount points”](#) on page 302.

## Preparing \$CRS\_HOME and \$ORACLE\_HOME on new node

To prepare the new node for installing Oracle, you must prepare \$CRS\_HOME and \$ORACLE\_HOME on the new node.

The local or cluster file system on shared storage should be the same as on the existing nodes.

### Local file system for \$CRS\_HOME and \$ORACLE\_HOME

Follow the directions for the following procedures to prepare \$CRS\_HOME and \$ORACLE\_HOME on new node:

- Create a file system on local storage for Oracle/CRS binaries (/app):  
 See [“To create a file system on local storage for Oracle CRS/binaries \(/app\)”](#) on page 414.
- Prepare \$CRS\_HOME on the new node:  
 See [“To prepare \\$CRS\\_HOME on the new node”](#) on page 415.
- Prepare \$ORACLE\_BASE on the new node:  
 See [“To prepare \\$ORACLE\\_BASE on the new node”](#) on page 415.
- Prepare \$ORACLE\_HOME on the new node:  
 See [“To prepare \\$ORACLE\\_HOME on the new node”](#) on page 416.

Identify the directories required for installing Oracle and CRS software.

**Table 20-2** Required information \$CRS\_HOME on each system

| Required information                                             | Example                                        |
|------------------------------------------------------------------|------------------------------------------------|
| Disk on each node for Oracle and CRS for Oracle and CRS binaries | Disk 1                                         |
| Disk group on each local storage                                 | On saturn: orabindg_saturn                     |
| Volume for each local disk group                                 | On saturn: orabinvol_saturn                    |
| File system on each local volume                                 | /dev/vx/rdisk/orabindg_saturn/orabinvol_saturn |
| Mount point for each local file system                           | /app                                           |
| Directory to install CRS                                         | /app                                           |
| Directory to install Oracle (ORACLE_HOME)                        | /app/oracle/orahome                            |

**To create a file system on local storage for Oracle CRS/binaries (/app)**

- 1** As root user, first create a VxVM local diskgroup, orabindg\_hostname:

```
vxvg init orabindg_saturn Disk_1
```

- 2** Create a volume, orabinvol\_hostname:

```
vxassist -g orabindg_saturn make orabinvol_saturn 12G
```

- 3** Create directory, /app

```
mkdir /app
```

- 4** Create a filesystem with this volume, orabinvol\_hostname

```
mkfs -V vxfs /dev/vx/rdisk/orabindg_saturn/orabinvol_saturn
```

**5 Mount /app:**

```
mount -V vxfs /dev/vx/dsk/orabindg_saturn/orabinvol_saturn /app
```

**6 Add an entry for this filesystem. For example, in the /etc/filesystems add:**

```
/app:
dev = dev/vx/dsk/orabindg_saturn/orabinvol_saturn
vfs = vxfs
mount = true
options = rw
account = false
```

**To prepare \$CRS\_HOME on the new node**

- 1** On each system, log in as "root".
- 2** Create the directory for CRS\_HOME:

```
mkdir -p /app/crshome
```

- 3** Change ownership and permissions:

```
chown -R oracle:oinstall /app/crshome
chmod -R 744 /app/crshome
```

**To prepare \$ORACLE\_BASE on the new node**

- 1** On each system, log in as "root".
- 2** Create the directory for ORACLE\_BASE:

```
mkdir -p /app/oracle
```

- 3** Change ownership and permissions:

```
chown -R oracle:oinstall /app/oracle
chmod -R 744 /app/oracle
```

#### To prepare \$ORACLE\_HOME on the new node

- 1 On each system, log in as "root".
- 2 Create the directory for ORACLE\_HOME:  
  

```
mkdir -p /app/oracle/orahome
```
- 3 Change ownership and permissions:  
  

```
chown -R oracle:oinstall /app/oracle/orahome
```

```
chmod -R 744 /app/oracle/orahome
```

## Preparing OCR and VOTE-disk mount point and Oracle database mount point on new node

To prepare the new node for installing Oracle, you must prepare the OCR and VOTE-disk mount point and the Oracle database mount point on the new node.

Create the mount point for the CRS file system having the OCR and Vote Disk files.

For example:

```
mkdir -p /ocrvote
```

For example:

```
mkdir -p /oradata
```

## Configuring CVM

To prepare the new node for installing Oracle, you must configure CVM.



## Configuring CVM

- 1 To enable the existing cluster to recognize the new node, execute the following commands on one of existing nodes:

```
haconf -makerw

hares -modify cvm_clus CVMNodeId -add saturn 2

haconf -dump -makero

/etc/vx/bin/vxclustadm -m vcs -t gab reinit

/etc/vx/bin/vxclustadm nidmap
```

- 2 Repeat the following step in the remaining existing cluster:

```
/etc/vx/bin/vxclustadm -m vcs -t gab reinit

/etc/vx/bin/vxclustadm nidmap
```

- 3 Start CVM on the newly added node.

```
hstart
```

- 4 Verify that the CVM group has come online on the newly added node.

```
hastatus -sum
```

## Linking Veritas membership library

To prepare the new node for installing Oracle, you must link the Veritas membership library.

Perform the following steps to link the Veritas membership library.

### To link Veritas Membership library

- 1 Create required sub-directories if they not present. For example, enter the following commands:

```
mkdir -p /usr/sbin/cluster/utilities
echo `lltstat -C` > /usr/sbin/cluster/utilities/cldomain
mkdir /opt/ORCLcluster
mkdir /opt/ORCLcluster/lib
mkdir /opt/ORCLcluster/lib32
mkdir /opt/ORCLcluster/bin
```

- 2 Archive skgxn library. For example, enter the following commands:

```
cd /opt/VRTSvcs/rac/lib
rm libskgxn2.a
cp libvcsmm.so shr_skgxn2.o
ar -X32 -qv libskgxn2.a shr_skgxn2.o
rm shr_skgxn2.o
cd /opt/VRTSvcs/rac/lib64
rm libskgxn2.a
cp libvcsmm.so shr_skgxn2.o
ar -X64 -qv libskgxn2.a shr_skgxn2.o
rm shr_skgxn2.o
```

- 3 Copy the skgxn library and clsinfo to /opt/ORCLcluster directory. For example, enter the following commands:

```
cp /opt/VRTSvcs/rac/bin/clsinfo \
/opt/ORCLcluster/bin

cp /opt/VRTSvcs/rac/lib/libskgxn2.a \
/opt/ORCLcluster/lib32

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

cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
/opt/ORCLcluster/lib

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

## Using the Oracle RAC 11g add node procedure

For the Oracle RAC 11g procedure for adding a node, refer to your Oracle documentation or the Oracle websites for information about this procedure.

In this procedure, Oracle copies the CRS\_HOME and ORACLE\_HOME from an existing node in the cluster .

## Relinking the SF Oracle RAC libraries to Oracle

After CRS and Oracle binaries is installed on the new node, relink the Veritas library on \$ORACLE\_HOME.

For information about this procedure:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 11g”](#) on page 330.

## Adding the new system to the SFDB repository

Add the new system to the Storage Foundation database repository using the following procedure.

### Add the new system to the SFDB repository

- 1 Add the system using the following sequence of commands:

```
haconf -makerw

hagrpl -modify Sfua_Base SystemList -add saturn

hares -modify sfua_ip Device en3 -sys saturn

haconf -dump -makero
```

- 2 Copy the /etc/vx/vxdbed/.odbc.ini file from an existing node to the new system using a remote file copy utility such as rcp, tcp, or scp. For example, to use rcp, enter:

```
rcp /etc/vx/vxdbed/.odbc.ini saturn:/etc/vx/vxdbed
```

- 3 Copy the /etc/vx/vxdbed/admin.properties file from an existing node to the new system using a remote file copy utility such as rcp, tcp, or scp.

For example, to use rcp:

```
rcp /etc/vx/vxdbed/admin.properties saturn:/etc/vx/vxdbed
```

- 4 Make a repository mount point on the newly added node using the following commands:

```
grep sfua_db_mountpoint /etc/vx/vxdbed/\
admin.properties sfua_db_mountpoint=/sfua_rep

mkdir /sfua_rep
```

# Removing a node from SF Oracle RAC clusters using Oracle RAC 11g

This chapter includes the following topics:

- [About removing a node from an Oracle RAC 11g cluster](#)
- [Using the Oracle remove node procedure](#)
- [Removing SF Oracle RAC using the `uninstallsrac` utility](#)


## About removing a node from an Oracle RAC 11g cluster

You can remove one or more nodes from a Oracle RAC 11g cluster by using the procedures described in this chapter.

To remove a node from a Veritas Storage Foundation for Oracle RAC cluster, you must follow the Oracle remove node procedure, as well as the procedure using the `uninstallsrac` utility and subsequent editing of the configuration files.

[Figure 21-1](#) displays the steps required to remove a node from a Veritas Storage Foundation for Oracle RAC cluster. Each step is discussed in this chapter.

**Figure 21-1** Removing a node from an Oracle RAC 11g cluster flowchart



## Using the Oracle remove node procedure

Use the Oracle remove node procedure, to remove the node from an Oracle RAC 11g cluster.

For the Oracle procedure for removing a node, refer to your Oracle documentation and the appropriate Oracle web sites.

Follow the instructions provided by Oracle.

## Removing SF Oracle RAC using the uninstallsfrac utility

Run the uninstallsfrac script from any node in the cluster, including a node from which you are uninstalling Veritas Storage Foundation for Oracle RAC.

---

**Warning:** Prior to starting the uninstallsfrac script, you must take all service groups offline and shut down VCS.

---

The example procedure below involves removing SF Oracle RAC from the node saturn.

**To remove SF Oracle RAC using the uninstallsfrac utility**

- 1** Offline all Oracle parallel groups.

```
hagr -offline oracle_group -sys nebula
```

- 2** Offline CRS using the following command:

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

If CRS is not under control, then run the following command:

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

- 3** Stop VCS on the node on which you want to remove SF Oracle RAC.

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

- 4** As superuser, start the uninstallation from any node from which you are uninstalling Veritas Storage Foundation for Oracle RAC.

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

```
./uninstallsfrac
```

- 5** Decline to uninstall SF Oracle RAC from all systems.

VCS configuration files exist on this system with the following information:

```
Cluster Name: rac_cluster101
Cluster ID Number: 101
Systems: galaxy nebula saturn
Service Groups: cvm oradb1_grp
```

```
Do you want to uninstall SFRAC from these systems? [y,n,q] (y) n
```

---

**Warning:** Be sure to answer "n". Otherwise, the utility begins the process of uninstalling SF Oracle RAC from all systems.

---

- 6** Specify the name of the system on which you are uninstalling SF Oracle RAC.

- 7 The installer prompts you to specify the name of the system from which you are uninstalling Veritas Storage Foundation for Oracle RAC:

Enter the system names separated by spaces on which to uninstall SFRAC:**saturn**

- 8 Review the output as the script checks for packages currently installed on your system. The utility also checks for dependencies between packages to determine the packages it can safely uninstall and in which order.
- 9 Confirm the uninstall when the uninstaller has completed checking.
- 10 Review the output as the script stops processes and drivers running on each node.
- 11 After the script uninstalls the packages, make a note of the location of the summary and log files that the uninstaller creates.

## Editing VCS configuration files on existing nodes

After running `uninstallsrac`, modify the configuration files on the existing remaining nodes to remove references to the deleted nodes. For an example `main.cf`:

See [“Sample main.cf for adding and removing an Oracle RAC 11g node”](#) on page 552.

The following procedures must be performed to edit the VCS configuration files on the existing remaining nodes:

- [Editing `/etc/llthosts`](#)
- [Editing `/etc/gabtab`](#)
- [Modifying the VCS configuration to remove the node](#)

### Editing `/etc/llthosts`

On each of the existing nodes, use `vi` or another editor to edit `/etc/llthosts`. Remove lines corresponding to the removed nodes; for example, if `saturn` is the node removed from the cluster, remove the line "2 saturn" from the file:

```
1 galaxy
2 nebula
3 saturn
```

Change to:



```
1 galaxy
2 nebula
```

## Editing `/etc/gabtab`

In the file `/etc/gabtab`, change the command contained in the file to reflect the number of systems after the node is removed:

```
/sbin/gabconfig -c -nN
```

where `N` is the number of nodes remaining.

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 using one of the following methods:

- Editing `/etc/VRTSvcs/conf/config/main.cf` directly
- Using the VCS Cluster Manager
- Using the command line, as illustrated in the following example

---

**Note:** For details on configuring VCS, refer to *Veritas Cluster Server User's Guide*.

---

At this point in the process, all Oracle binaries have been removed from the system to be deleted. The instance has been removed from the database, that is, the thread disabled, and the `spfile<SID>.ora` edited by Oracle to remove any references to this instance. The next step is to remove all references in the `main.cf` to the deleted nodes.

As root user execute the following on the CVM master node only.

Perform the following procedure to modify the CVM group in the `main.cf` file.

### Modify the VCS configuration to remove the node

- 1 To determine the CVM master node execute:

```
vxdctl -c mode
```

- 2 Make a backup copy of the `main.cf` file.

```
cd /etc/VRTSvcs/conf/config
```

```
cp main.cf main.cf.3node.bak
```

- 3 Use the following commands to reconfigure the VxSS group. Execute the following commands:

```
haconf -makerw
```

```
hagrpl -modify VxSS SystemList -delete saturn
```

```
haconf -dump -makero
```

- 4 Use the following commands to reconfigure the CVM group. Execute the following commands:

```
haconf -makerw
```

```
hares -modify cvm_clus CVMNodeId -delete saturn
```

```
hagrpl -modify cvm SystemList -delete saturn
```

```
hasys -delete saturn
```

```
haconf -dump -makero
```

Example of `main.cf` file:

See [“Sample main.cf for adding and removing an Oracle RAC 11g node”](#) on page 552.

- 5 Verify the syntax of `main.cf` file:

```
hacf -verify .
```

The `main.cf` file now should not contain entries for system saturn.

- 6 Copy the new version of the `main.cf` to each system in the cluster.

```
rcp (or scp) main.cf galaxy:/etc/VRTSvcS/conf/config
```

```
rcp (or scp) main.cf nebula:/etc/VRTSvcS/conf/config
```



## Configuring an SF Oracle RAC global cluster

- [Preparing for a replicated SF Oracle RAC global cluster](#)
- [Configuring a global cluster using VVR](#)



# Preparing for a replicated SF Oracle RAC global cluster

This chapter includes the following topics:

- [Replication in the SF Oracle RAC environment](#)
- [Requirements for SF Oracle RAC global clusters](#)
- [Global cluster configuration overview](#)
- [Configuring an SF Oracle RAC cluster at the primary site](#)
- [Configuring an SF Oracle RAC cluster at the secondary site](#)
- [Configuring replication on clusters at both sites](#)
- [Configuring VCS service groups for global clusters](#)
- [Testing a global cluster configuration](#)

## 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 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.

## Requirements for SF Oracle RAC global clusters

Review the requirements information to make sure your configuration is supported for SF Oracle RAC.

### Supported software and hardware for SF Oracle RAC

For supported hardware and software:

- See “[Installation requirements](#)” on page 34.
- 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 22-1** 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            |
| 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 |



**Table 22-1** Supported replication options for SF Oracle RAC global clusters  
(continued)

| Replication technology                            | Supported modes                                                                                                                                         | Supported software                                                                            |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| 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>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

**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*
- Technical Support TechNote for the latest updates or software issues for replication agents:  
<http://entsupport.symantec.com/docs/282004.htm>

## Global cluster configuration overview

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

For sample primary and secondary cluster configuration files:

See “[Sample main.cf for Oracle RAC 10g for hardware-based replication primary site](#)” on page 538.

See “[Sample main.cf for Oracle RAC 10g for hardware-based replication secondary site](#)” on page 543.

Some SF Oracle RAC HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration.

For requirements:

See “[Installation requirements](#)” on page 34.

For instructions on configuring AT in a global cluster:

See the *Veritas Cluster Server User's Guide*

## Configuring an SF Oracle RAC 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 [“Important preinstallation information”](#) on page 33.

See [“Overview of SF Oracle RAC installation and configuration tasks”](#) on page 46.

See [“Global cluster information”](#) on page 565.

See [“Sample main.cf files for Oracle replication configurations”](#) on page 529.

See [“SF Oracle RAC installation and configuration information”](#) on page 559.

If you have an existing SF Oracle RAC cluster, and you want to set up a global cluster, skip the steps below and proceed to obtain and install license keys for HA/DR.

See [“Obtaining and installing license keys for a global cluster”](#) on page 77.

If you are setting up an SF Oracle RAC global cluster from the beginning, follow the steps below.

### To set up the primary site

- 1 Install and configure servers and storage.
- 2 Install the software for managing your array.
- 3 Obtain and install the licenses for HA/DR (required) and VVR (optional).

See [“Obtaining and installing license keys for a global cluster”](#) on page 77.

To verify existing licenses for global clustering or VVR:

See [“To verify VVR and global clustering licenses are enabled on all nodes”](#) on page 446.

- 4 Install and configure SF Oracle RAC. Prepare for your installation by planning your configuration, reviewing the the requirements, and performing pre-installation tasks.

For requirements:

See [“Installation requirements”](#) on page 34.

For preparation:

See [“About preparing to install and configure SF Oracle RAC”](#) on page 57.

For installation:

See [“About installing SF Oracle RAC”](#) on page 81.

After verifying a successful SF Oracle RAC installation:

Verify the CVM group is online on all nodes in the primary cluster:

```
hagrps -state cvm
```

- 5 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
- Non-replicated storage for OCR and Vote disk
- Replicated storage for database files

## 6 Install and configure Oracle RAC

For Oracle RAC 9i:

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

For Oracle RAC 10g:

See [“About preparing to install Oracle RAC 10g”](#) on page 195.

See [“About installing Oracle RAC 10g”](#) on page 231.

For Oracle 11g RAC:

See [“About preparing to install Oracle RAC 11g”](#) on page 275.

See [“About installing Oracle RAC 11g”](#) on page 315.

---

**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
```

## 7 Identify the disks that will be replicated, create the required CVM disk group, volume, and file system. Create the database on that file system.

## 8 Configure the VCS service groups for the database.

To configure service groups for Oracle RAC 9i:

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

To configure service groups for Oracle RAC 10g RAC:

See [“About VCS service group for Oracle RAC 10g dependencies”](#) on page 253.

To configure service groups for Oracle RAC 11g RAC:

See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.

Verify that all VCS service groups are online.

For sample configurations:

See [“Sample main.cf files for Oracle replication configurations”](#) on page 529.

## Configuring an SF Oracle RAC cluster at the secondary site

You can set up a multi-node or single-node cluster on the secondary site. 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 Oracle group and user IDs for both the primary and secondary clusters.

### To set up the secondary site

- 1 Install and configure servers and storage.
- 2 Install the software for managing your array.
- 3 Obtain and install the licenses for HA/DR (required) and VVR (optional).

See [“Obtaining and installing license keys for a global cluster”](#) on page 77.

To verify existing licenses for global clustering or VVR:

See [“To verify VVR and global clustering licenses are enabled on all nodes”](#) on page 446.

- 4 Install and configure SF Oracle RAC. Prepare for your installation by planning your configuration, reviewing the requirements, and performing pre-installation tasks.

For requirements:

See [“Installation requirements”](#) on page 34.

For preparation:

See [“About preparing to install and configure SF Oracle RAC”](#) on page 57.

For installation:

See [“About installing SF Oracle RAC”](#) on page 81.

After verifying a successful SF Oracle RAC installation:

Verify the CVM group is online on all nodes in the primary cluster:

```
hagrps -state cvm
```

- 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 setting up I/O fencing”](#) on page 116.

- 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
- Non-replicated storage for OCR and Vote disk
- Replicated storage for database files

- 8 Install and configure Oracle RAC

For Oracle RAC 9i:

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

For Oracle RAC 10g:

See [“About preparing to install Oracle RAC 10g”](#) on page 195.

See [“About installing Oracle RAC 10g”](#) on page 231.

For Oracle RAC 11g:

See [“About preparing to install Oracle RAC 11g”](#) on page 275.

See [“About installing Oracle RAC 11g”](#) on page 315.

---

**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 Do not create the database, the disk group, or volumes. 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.
- 10 Copy the init file from \$ORACLE\_HOME/dbs at the primary to \$ORACLE\_HOME/dbs.
- 11 Create the following subdirectories as on the primary site:  
For Oracle RAC 10g

```
$ mkdir $ORACLE_HOME/admin
$ mkdir $ORACLE_HOME/admin/$DB_NAME
$ cd $ORACLE_HOME/admin/$DB_NAME
$ mkdir bdump cdump create hdump pfile udump
```

For Oracle RAC 11g there is no need to create the subdirectories since they are created automatically by Oracle RAC 11g when the database starts.



## 12 Configure the VCS service groups for the database.

To configure service groups for Oracle RAC 9i:

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

To configure service groups for Oracle RAC 10g:

See [“About VCS service group for Oracle RAC 10g dependencies”](#) on page 253.

To configure service groups for Oracle RAC 11g:

See [“About VCS service group for Oracle RAC 11g dependencies”](#) on page 351.

Verify that all VCS service groups are online.

For sample configurations:

See [“Sample main.cf files for Oracle replication configurations”](#) on page 529.

---

**Note:** The database service group will be offline at the secondary site.

---

# Configuring replication on clusters at both sites

You must configure replication for the database files at both the primary and secondary site. 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 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 [“Preparing clusters for replication using VVR”](#) on page 446.

For Hardware-based replication:

See your hardware documentation for details on setting up replication between the two sites.

- 3 Start replication between the sites.  
See your hardware documentation for the appropriate procedures.
- 4 Test the replication between sites.

**To test replication between sites**

- 1 Create some sample files at the primary site on the replicated storage.
- 2 When replication is complete, split or suspend the replication.
- 3 Import the required disk groups, mount the coresponding file system, and verify that the sample files exist at the secondary site.
- 4 Unmount the file system, deport the disk group, and establish replication.

## Configuring VCS service groups for global clusters

You have configured VCS service groups for 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.

If you are using VVR for replication, configure VCS service groups after installing Oracle on the secondary site, and before establishing replication.

If you did not configure the ClusterService group when you installed SF Oracle RAC on the secondary site, configure it when you configure global clustering.

Configure a global cluster by setting:

- Heartbeat
- Wide areacluster (wac)
- GCO IP (gcoip)
- remote cluster resources

For sample main.cf files for global cluster configurations:

See [“Sample main.cf files for Oracle replication configurations”](#) on page 529.

**To configure VCS service groups for global clusters**

- 1 Configure and enable global groups for databases and resources.
  - Configure VCS service groups at the secondary site and test local failover at that site (assuming more than one node in the cluster at the secondary site).
  - Configure the replication agent at both sites.

- Make the RAC service group a global service group, enabling failover across clusters.
- 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.
- 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*
- 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*.

## Testing a global cluster configuration

### To test a planned failover

- 1 Offline the VCS service group for the Oracle RAC database on the cluster at the primary site.
- 2 Online the VCS service group for the Oracle RAC database on the cluster at the secondary site.

**To test distaster recovery at the recovery site**

- 1** Plan downtime to test the disaster recovery configuration.
- 2** Simulate a disaster at the primary site.

For example:

Shut down the hosts and storage arrays at the primary. If you can not shut down the storage arrays, disconnect the replication link between the sites.

- 3** Use VCS to fail over the Oracle RAC database to the cluster at the secondary site.

**To test failback on the primary site**

- 1** Offline the VCS service group for the Oracle RAC database on the cluster at the secondary site.
- 2** Restart the nodes and the storage array at the primary site.
- 3** Reconnect the replication link if it was broken.
- 4** Resynchronize the data from the secondary to make sure the data at the primary site is current.
- 5** Bring the VCS service group online at the primary site.

# Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring global clustering using VVR](#)
- [Preparing clusters for replication using VVR](#)
- [Setting up replication using VVR](#)
- [Configuring VCS to replicate the database volume using VVR](#)

## About configuring global clustering using VVR

After setting up a secondary cluster running SF Oracle RAC, you can configure a global cluster environment. 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.
- Setting up replication on both clusters.
- Starting replication of the database.
- Configuring VCS on the primary site for replication.
- Configuring VCS on the secondary site for replication.

For sample primary and secondary cluster configuration files:

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(primary replication site\)”](#) on page 529.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(secondary replication site\)”](#) on page 534.

## Preparing clusters for replication using VVR

Before configuring clusters for global clustering, make sure both clusters have the following:

- SF Oracle RAC and Oracle RAC installed and configured
- License keys installed for VVR
- VVR and global clustering enabled

For global clustering and VVR licensing information:

See [“Obtaining and installing license keys for a global cluster”](#) on page 77.

**To verify VVR and global clustering licenses are enabled on all nodes**

- 1 On all nodes of the primary and secondary clusters, enter:

```
vxlicrep |grep -i global
Global Cluster Option#VERITAS Cluster Server = Enabled
Global Cluster Option = Enabled
```

- 2 On both primary and secondary clusters, enter:

```
vxlicrep |grep -i vvr
VVR#VERITAS Volume Manager = Enabled
VVR = Enabled
```

The example procedures in the following topics assume rac\_cluster101 as the local cluster with nodes galaxy and nebula, and rac\_cluster102 as the remote cluster on the secondary site with nodes mercury and jupiter.

**To view sample configuration files**

- 1 Navigate to the directory containing sample configuration files:

```
cd /etc/VRTSvcs/conf/sample_rac
```

- 2 List the SF Oracle RAC VVR/CVR sample main.cf examples:

```
ls *cvmvvr*
cvmvvr_primary_main.cf* cvmvvr_secondary_main.cf*
```

These main.cf examples are for VVR/CVR primary and secondary configurations and will help guide you through the following sections.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(primary replication site\)”](#) on page 529.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(secondary replication site\)”](#) on page 534.

Preparing clusters for replication in both sites requires:

- [Adding the VVR resource types to the VCS configuration](#)
- [Configuring global clustering using VVR](#)
- [Defining the remote cluster and heartbeat Cluster Objects](#)

## Adding the VVR resource types to the VCS configuration

After SF Oracle RAC is installed on two clusters and fencing is configured, add the VVR types to the VCS configuration.

To add VVR types to the VCS configuration on each cluster perform the tasks described in the procedure below. In the procedure below and as an example, there is a local cluster and a remote cluster with the following names:

```
LOCAL_CLUSTER_NAME: rac_cluster101 (galaxy,nebula)
REMOTE_CLUSTER_NAME: rac_cluster102 (mercury,jupiter)
```

The primary site is rac\_cluster101 and the secondary site is rac\_cluster102. A virtual IP address of 10.180.8.188 is associated with rac\_cluster101 and a second virtual IP address of 10.190.99.199 is associated with rac\_cluster102.

### Adding VVR Types to the VCS Configuration

- 1 On the first cluster, make sure that CVM is up.
- 2 Make sure you have installed VVR license keys with "VVR" enabled.  
See [“Obtaining SF Oracle RAC license keys”](#) on page 75.
- 3 On any node in the first cluster, use the following command to enable write access to the VCS configuration:

```
haconf -makerw
```

- 4 Run the script to add definitions for VVR types:

```
cd /etc/VRTSvcs/conf/sample_vvr
./addVVRTypes.sh
```

When you run the script "addVVRTypes.sh" for adding VVR resource definitions to the VCS configuration, the following warning messages are displayed:

```
VCS WARNING V-16-1-10421
Non static attribute does not exist 'Primary' (RVG)
VCS WARNING V-16-1-10421
Non static attribute does not exist 'SRL' (RVG)
VCS WARNING V-16-1-10421
Non static attribute does not exist 'RLinks' (RVG)
```

These warnings may be ignored.

- 5 Save the VCS configuration and change the access to read-only:

```
haconf -dump -makero
```

- 6 Perform the previous steps on the second cluster, `rac_cluster102`.

## Configuring global clustering using VVR

You can configure global clustering on SF Oracle RAC clusters by editing the configuration files:

See [“Modifying the configuration for VVR using the main.cf”](#) on page 449.

Before you configure global clustering, review the following requirements:

- Cluster names on the primary and secondary sites must be unique.
- Node and resource names must be unique within a cluster but not across clusters.
- Each cluster requires a virtual IP address associated with the cluster. The VCS installation and creation of the ClusterService group typically involves defining this IP address. If you did not configure the ClusterService group when you installed SF Oracle RAC, configure it when you configure global clustering.
- One WAN (Wide Area Network) heartbeat must travel between clusters, assuming each cluster has the means to monitor the health of the remote cluster. Configure the heartbeat resource manually.
- All Oracle user and group IDs must be the same on all nodes.



- The Oracle RAC database, which VVR replicates from the storage on the primary site to the secondary site, must be defined in a global group having the same name on each cluster. Each resource in the group may differ from cluster to cluster, but clients redirected to a remote cluster after a wide-area failover must see the same application as the one in the primary cluster.

See the *Veritas Cluster Server User's Guide* for complete details on global clustering.

## Modifying the configuration for VVR using the main.cf

### To modify the main.cf for VVR

- 1 Edit the main.cf file to specify the virtual IP address for the local cluster and define the ClusterService group for the local cluster.

The example global clustering configuration shows the rac\_cluster101 cluster on the primary site. The additions to the configuration appear in bold text.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "VVRTypes.cf"

cluster rac_cluster101 (
 UserNames = { admin = "cDRpdxPmHpzS." }
 ClusterAddress = "10.10.10.101"
 Administrators = { admin }
 CounterInterval = 5
 UseFence = SCSI3
)

system galaxy (
)

system nebula (
)

group ClusterService (
 SystemList = { galaxy = 0, nebula = 1 }
 AutoStartList = { galaxy, nebula }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)
```

```

Application wac (
 StartProgram = "/opt/VRTSvcs/bin/wacstart"
 StopProgram = "/opt/VRTSvcs/bin/wacstop"
 MonitorProcesses = {"/opt/VRTSvcs/bin/wac" }
 RestartLimit = 3
)

IP gcoip (

 Device = en0

 Address = "10.10.10.101"
 NetMask = "255.255.255.0"
)
NIC csgnic (
 Device = en0
)
gcoip requires csgnic
wac requires gcoip

.
.
.

group cvm (
.
.
.

```

- 2 Edit the main.cf file for the secondary cluster as illustrated in the previous step using configuration values that will be appropriate for the secondary cluster.

- 3 On all nodes enter:

```
hastop -local -force
```

Wait for port 'h' to close on all systems

- 4 On all nodes enter:

```
hstart
```

## Defining the remote cluster and heartbeat Cluster Objects

After configuring global clustering, add the remote cluster 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 through step 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
```

The final output should resemble the output displayed below, from `rac_cluster01` (primary):

```
hastatus -sum
.....
-- WAN HEARTBEAT STATE
-- Heartbeat To State

L Icmp rac_cluster102 ALIVE

-- REMOTE CLUSTER STATE
-- Cluster State

M rac_cluster102 RUNNING

-- REMOTE SYSTEM STATE
-- cluster:system State Frozen

N rac_cluster102:mercury RUNNING 0
N rac_cluster102:jupiter RUNNING 0
```

## 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:

```
.
.
remoteclass 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:

```
.
.
remoteclass 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.

## Setting up replication using VVR

Setting up replication with VVR in a global cluster environment involves the following tasks:

- Create a disk group on the storage on the primary site ( if you have not already done so).  
See [“Configuring an SF Oracle RAC cluster at the primary site”](#) on page 435.
- Creating the Storage Replicator Log (SRL) in the disk group for the database.  
See [“Creating the SRL volume on the primary site”](#) on page 454.
- Creating the Replicated Volume Group (RVG) on the primary site.  
See [“Setting up replication objects on the primary site”](#) on page 455.
- Setting up replication objects on the secondary site.  
See [“Configuring replication for the secondary site”](#) on page 456.

### 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.  
See [“Configuring an SF Oracle RAC cluster at the secondary site”](#) on page 438.
- 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 racl_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 replication objects 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.
- 3 Check that the CVM group is online.
- 4 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
vradmind: 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/vradmind
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`.

## Configuring replication for 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:

- Create a disk group on the storage with the same name as the equivalent disk group on the primary site (if you have not already done so).  
See [“Configuring an SF Oracle RAC cluster at the secondary site”](#) on page 438.
- Create volumes for the database and SRL on the secondary site.  
See [“Creating the data and SRL volumes on the secondary site”](#) on page 456.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.  
See [“Editing the /etc/vx/vras/.rdg files”](#) on page 457.
- 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 457.
- Create the replication objects on the secondary site.  
See [“Setting up disk group on secondary site for replication”](#) on page 458.

## 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 volume 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 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 vradmin 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 the network interface is en0:1
- The virtual IP address is 10.10.9.101
- The net mask is 255.255.255.0
- ```
# ifconfig en0 10.10.9.101 netmask 255.255.255.0 \
    broadcast 10.180.95.255 alias
# ifconfig en0 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 the network interface is en0:1
- 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_clus101_priv has IP address 10.10.9.101
- rac_clus102_priv has IP address 10.11.9.102

- 4 Use the ping command to verify the links are functional.

Setting up disk group on secondary site for replication

Create the replication objects on the secondary site from the master node on 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: `oradatadg`

- `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: `10.10.9.101` or `rac_clus101_priv`
- `sec_host` is the virtual IP address or resolvable virtual host name of the cluster on the secondary site. For example: `10.11.9.102` or `rac_clus102_priv`

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_clus101_priv  
rac_clus102_priv
```

- 2 On the secondary site, the 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_clus102_priv_rac1_rvg` and the "secondary" RLINK created is `rlk_rac_clus101_priv_rac1_rvg`.

- 3 Verify the list of RVGs in the RDS by executing the following command.

```
# vradmin -g oradg -l printrvg
```

For example:

```
Reeplicated Data Set: rac1_rvg
Primary:
HostName: 10.180.88.187 <localhost>
RvgName: rac1_rvg
DgName: oradatadg
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_10.11.9.102_ 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_10.10.9.101_ rac1_rvg, detached=on,
synchronous=off
```

Note: Once the replication is started the value off 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_clus102_priv
```

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_clus102_priv
```

- 2 To start replication after full synchronization, enter the following command:

```
# vradmin -g oradatadg -c rac1_ckpt startrep rac1_rvg  
rac_clus102_priv
```

Verifying replication status

Verify that replication is properly functioning.

To verify replication status

- 1 Use the `vxprint` command on the primary site:

```
vxprint -g diskgroup -l rlink_name
```

- 2 Review the `flags` output for the status. The output may appear as `connected` and `consistent`. For example:

```
# vxprint -g oradatadg -l rlk_10.182.13.221_oradatadg
Rlink: rlk_10.182.13.221_oradatadg
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.

The following sample `main.cf` files illustrate the VCS configuration changes after setting up an existing Oracle RAC database for replication:

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(primary replication site\)”](#) on page 529.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(secondary replication site\)”](#) on page 534.

To view the sample `main.cf` files on your system, enter the following commands:

```
# pwd

/etc/VRTSvcs/conf/sample_rac
```

```
# ls cvmvvr*

cvmvvr_primary_main.cf*    cvmvvr_secondary_main.cf*
```

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:

See [“Configuration examples before and after modification”](#) on page 465.

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 627.

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 CFSSMount 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 23-1](#)

- Configuration after modification:

- [Figure 23-2](#)


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.

For Oracle RAC 9i, the CVM service group may be different on your configuration, depending on the type of IP you set up for the Listener.

Dependencies


Configuration before modification for replication:

Figure 23-1 Illustration of dependencies for AIX: configuration before modification for replication (Oracle RAC 10g)



Configuration after modification for replication:

Figure 23-2 Illustration of dependencies for AIX: configuration 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 RVGShared service 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.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(primary replication site\)”](#) on page 529.

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 = en0  
  
    Address = "10.10.9.101"  
    NetMask = "255.255.255.0"  
)  
  
NIC nic (  
    Device = en0  
  
    NetworkType = ether  
)  
  
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.

The following is an example of an Oracle RAC database service group configured for replication:

```
group oradb1_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"
    Pfile @galaxy = "/oracle/orahome/dbs/initvrts1.ora"
    Pfile @nebula = "/oracle/orahome/dbs/initvrts2.ora"
    ShutDownOpt = SRVCTLSTOP
    MonScript = "./bin/Oracle/SqlTest.pl"
)

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.

```
# hastop -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.

See [“Sample main.cf for Oracle RAC 10g and CVM/VVR \(secondary replication site\)”](#) on page 534.

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 = en0

    Address = "10.11.9.102"
    NetMask = "255.255.255.0"
)

NIC nic (

Device = en0

    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, oradb1_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 oradb1_grp (
    SystemList = { mercury = 0, jupiter = 1 }
    ClusterList = { rac_cluster102 = 0, rac_cluster101 = 1 }
    Parallel = 1
    OnlineRegryInterval = 300
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { mercury, jupiter }
)

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 @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
    MonScript = "./bin/Oracle/SqlTest.pl"
)

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
```

Starting VCS on all nodes in both clusters

After modifying the `main.cf` files for replication on the primary and secondary clusters, start VCS on both clusters.

To start VCS on all nodes in both clusters

- 1 From the primary site, use the following command to start the VCS engine on one node:

```
# hstart
```

- 2 Next, enter the following command:

```
# hastatus
```

- 3 When `LOCAL_BUILD` or `RUNNING` is listed in the message column, use the following command to start VCS on the other node:

```
# hstart
```

- 4 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.

On the primary site, enter the following commands:

```
# hagrps -online rlogowner -sys galaxy
```

```
# hagrps -online oradb1_grp -sys galaxy
```

```
VCS WARNING V-16-1-50817 Please use hagrps -online -force to online a group for the first time
```

```
# hagrps -online -force oradb1_grp -sys galaxy
```

On the secondary site, enter the following command:

```
# hagrps -online rlogowner -sys mercury
```

- 5 On the secondary site, start VCS from one node using the following command:

```
# hstart
```

- 6 Next, enter the following command:

```
# hastatus
```

- 7 When LOCAL_BUILD or RUNNING is listed in the message column, use the following command to start VCS on the other node:

```
# hstart
```

- 8 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.

Uninstalling SF Oracle RAC

- [Uninstalling SF Oracle RAC from an Oracle RAC 10g cluster](#)
- [Uninstalling SF Oracle RAC from an Oracle RAC 11g cluster](#)

Uninstalling SF Oracle RAC from an Oracle RAC 10g cluster

This chapter includes the following topics:


- [About uninstalling SF Oracle RAC from an Oracle RAC 10g cluster](#)
- [Preparing to uninstall SF Oracle RAC from an Oracle RAC 10g cluster](#)
- [Removing SF Oracle RAC from an Oracle RAC 10g cluster](#)

About uninstalling SF Oracle RAC from an Oracle RAC 10g cluster

You can uninstall the SF Oracle RAC software from the nodes in a cluster. When prompted by the installer, you can select the nodes in a cluster from which you want to uninstall the SF Oracle RAC.

[Figure 24-1](#) displays at a high-level, the steps that are required to uninstall SF Oracle RAC from a cluster having Oracle RAC 10g instances.

Figure 24-1 SF Oracle RAC with Oracle RAC 10g uninstallation flowchart



Preparing to uninstall SF Oracle RAC from an Oracle RAC 10g cluster

Perform the following steps before you uninstall SF Oracle RAC from a cluster having Oracle RAC 10g instances:

- [Stopping Oracle instances](#)

- Backing up the Oracle database
- Moving or removing the Oracle database (optional)
- Uninstalling Oracle RAC 10g (optional)
- Removing repository database
- Stopping applications that use CFS (outside of VCS control)
- Unmounting CFS file systems (outside of VCS control)
- Stopping applications that use VxFS (outside of VCS control)
- Unmounting VxFS file systems (outside of VCS control)
- Stopping VCS

Stopping Oracle instances

You must stop the Oracle instances and CRS on the cluster nodes where you want to uninstall SF Oracle RAC. Before stopping the Oracle instances, you must stop the service groups that are dependent on the Oracle service groups.

The following procedure provides the steps to stop the instances on a two-node cluster. In this procedure, the nodes are galaxy and nebula. Depending on the VCS configuration, the procedure to stop Oracle instances may vary.

If the Oracle RAC 10g instances are under VCS control, you must take the Oracle resources offline.

To stop Oracle instances

- 1 Log in as superuser.
- 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:

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

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

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

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

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

For example:

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

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

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

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

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

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, refer to your Oracle RAC 10g documentation.

After uninstalling SF Oracle RAC, you can not access the Oracle database as Veritas Volume Manager and Veritas File System are uninstalled from the cluster.

Moving or removing the Oracle database (optional)

You can move or remove the Oracle database after safely relocating the data as necessary. For instructions on moving or removing the Oracle database, refer to your Oracle RAC 10g documentation.

Uninstalling Oracle RAC 10g (optional)

You can either uninstall Oracle RAC 10g or choose to run Oracle RAC 10g after uninstalling SF Oracle RAC.

If you choose to uninstall Oracle RAC 10g, use the Oracle runInstaller utility.

- If Oracle RAC 10g is installed locally on each node, run the utility on each node.

- If Oracle RAC 10g is installed on a cluster file system, run the utility once.

To uninstall Oracle RAC 10g

- 1 Log in as oracle user.
- 2 Set the DISPLAY variable. Depending on the shell you use, execute the following command:

```
Bourne Shell (sh or ksh)  $ DISPLAY=host:0.0;export DISPLAY
```

```
C Shell (csh or tcsh)    $ setenv DISPLAY host:0.0
```

- 3 Run the Oracle RAC 10g runInstaller utility:

```
$ /cdrom/Disk1/runInstaller
```

- 4 Select the option to uninstall the Oracle RAC 10g software.
Refer to your Oracle RAC 10g documentation for details on this utility.
- 5 If necessary, uninstall Oracle RAC 10g from the other nodes using this procedure.

Removing repository database

To remove the Storage Foundation for Oracle repository database, use the following procedure. You must have superuser (root) privileges to remove the repository database.

To remove repository database

- 1 Run the following commands to remove the repository configuration from the VCS configuration and deport the repository disk group.

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

- 2 Import the repository disk group using the command:

```
# vxvg import <name_of_disk_group>
```

- 3 Mount the repository volume using the command:

```
# /opt/VRTSdbcom/config/sfua_rep_mount start
```
- 4 Drop the repository database using the command:

```
# /opt/VRTS/bin/sfua_db_config -o dropdb
```

Stopping applications that use CFS (outside of VCS control)

All Oracle users must stop all applications using the CFS mounts not under VCS control.

To verify that no processes use the CFS mount point

- 1 Enter the following command:

```
# fuser -c mount_point
```
- 2 Stop any processes using a CFS mount point with the mechanism provided by the application.

Unmounting CFS file systems (outside of VCS control)

All Oracle users must unmount any CFS file systems not under VCS control on all nodes.

To unmount CFS file systems not under VCS control

- 1 Log in as root.
- 2 Remove the existing unused modules in the kernel memory and in the library memory.

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

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

```
# umount mount_point
```

Stopping applications that use VxFS (outside of VCS control)

All Oracle users must stop all applications using the VxFS mounts not under VCS control.

To verify that no processes use the VxFS mount point

- 1 Enter the following command:

```
# fuser -c mount_point
```

- 2 Stop any processes using a VxFS mount point with the mechanism provided by the application.

Unmounting VxFS file systems (outside of VCS control)

All Oracle users must unmount any VxFS file systems not under VCS control on all nodes.

Note: To avoid issues on reboot, you must remove all entries of VxFS from /etc/filesystem folder.

To unmount VxFS file systems not under VCS control

- 1 Log in as root.
- 2 Remove the existing unused modules in the kernel memory and in the library memory.

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

- 3 Determine the file systems to unmount by checking the output of the mount file. For example, type:

```
# mount | grep vxfs | grep cluster
```

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

```
# umount mount_point
```

Stopping VCS

Stop VCS to take the service groups on all nodes offline:

To stop VCS

- 1 Log in as superuser on one cluster node.
- 2 Stop VCS on all nodes.

```
# /opt/VRTSvcs/bin/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.

```
# /sbin/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
```

Removing SF Oracle RAC from an Oracle RAC 10g cluster

You can remove the SF Oracle RAC filesets from all nodes in the SF Oracle RAC cluster using the `uninstallsfrac` program. Note that the `uninstallsfrac` program removes all SF Oracle RAC filesets regardless of the Oracle version used.

If you need to uninstall SF Oracle RAC after an incomplete installation, or if the `uninstallsfrac` program is not available in `/opt/VRTS/install`, you may need to use the `uninstallsfrac` program on the SF 5.0 MP3 for Oracle RAC disc.

Removing SF Oracle RAC filesets involve the following tasks:

- [Removing the SF Oracle RAC filesets](#)
- [Removing other configuration files \(optional\)](#)

Removing the SF Oracle RAC filesets

The installer provides you an option to remove Veritas Volume Manager and Veritas Volume Replicator filesets. Note that `uninstallsfrac` program can remove these filesets only if both the boot disk is not under VxVM control and there are no open volumes.

To remove the SF Oracle RAC filesets

- 1 Do one of the following to begin uninstalling:
 - If you can execute commands as superuser on the remote nodes in the cluster using SSH or RSH, run `uninstallsfrac` program on one node to uninstall SF Oracle RAC on all nodes in the cluster.
 - If you cannot execute commands as superuser on remote nodes in the cluster using SSH or RSH, you must run `uninstallsfrac` program on each node in the cluster.

- 2 Navigate to the folder containing the `uninstallsfrac` program.

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

- 3 Start the `uninstallsfrac` program.

```
# ./uninstallsfrac [-rsh]
```

The program specifies the directory where the logs are created and begins with a copyright message.

- 4 Indicate whether or not you want to remove VxVM and VVR filesets from all nodes in the cluster. Enter "y" only if the root disk is outside of VM control.
- 5 If you invoke the uninstallation program from a remote system in the same subnet, enter the system names where you want the uninstallation to take place.

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

- 6 If you invoke the uninstallation program from a node in the SF Oracle RAC cluster, review the cluster information and confirm to uninstall SF Oracle RAC.

VCS configuration files exist on this system with the following information:

```
Cluster Name: rac_cluster101  
Cluster ID Number: 101  
Systems: galaxy nebula  
Service Groups: cvm
```

Do you want to uninstall SFRAC from these systems? [y,n,q] (y)

- 7 Review the output as the uninstallation program checks the operating system on each system, verifies system-to-system communication, and verifies the system licensing.
- 8 Review the output as the uninstallation program checks for Storage Foundation for Oracle RAC filesets currently installed on the nodes.

This process involves identifying system uninstall requirements and dependencies between filesets to determine the safety and order of uninstalling filesets.
- 9 Confirm to uninstall SF Oracle RAC:


```
All SFRAC processes that are currently running will be stopped.  
Are you sure you want to uninstall SFRAC? [y,n,q] (y)
```
- 10 Make sure that the user VxFS file systems that are currently mounted on each node are unmounted before uninstalling SF Oracle RAC.


```
All user VxFS filesystems that are currently mounted on galaxy  
should be unmounted before uninstalling.
```


See [“Unmounting CFS file systems \(outside of VCS control\)”](#) on page 486.
- 11 Make sure that each node in the cluster does not have any attached volumes.
- 12 Review the output as the uninstallation program stops agents and performs verifications on each node to proceed with the uninstallation.
- 13 Review the output as the uninstallation program stops the SF Oracle RAC processes.
- 14 Review the output as the uninstallation program uninstalls the SF Oracle RAC filesets.
- 15 If necessary, review the summary and log files of uninstallation activities.

Removing other configuration files (optional)

You can remove Veritas configuration files and filesets that are left after running the `uninstallsfrac` utility.

To remove residual Veritas configuration files (optional)

- 1 List all VRTS filesets that can be removed.

```
lsllpp -l |grep -i vrts
```

- 2 Run `installp -u` command to remove left over VRTS filesets.

- 3 The following commands moves the residual Veritas configuration files to the `vrts.53` directories:

```
# cd /var
# mkdir vrts.53
# mv *VRTS* vrts.53
# mv vx vrts.53
# cd /var/opt
# mkdir vrts.53
# mv *VRTS* vrts.53
# cd /opt
# mkdir vrts.53
# mv *VRTS* vrts.53
# cd /etc
# mkdir vrts.53
# mv vx *llt* *fen* *gab* *vcs* vcsmmtab vrts.53
```

You can remove the `vrts.53` directories at a later time.

If you reinstall without removing the Veritas configuration files and filesets, the reinstall fails.

Rebooting the nodes

After uninstalling SF Oracle RAC, reboot each node:

```
# /usr/sbin/shutdown -Fr
```


Uninstalling SF Oracle RAC from an Oracle RAC 11g cluster

This chapter includes the following topics:


- [About uninstalling SF Oracle RAC from an Oracle RAC 11g cluster](#)
- [Preparing to uninstall SF Oracle RAC from an Oracle RAC 11g cluster](#)
- [Removing SF Oracle RAC from an Oracle RAC 11g cluster](#)

About uninstalling SF Oracle RAC from an Oracle RAC 11g cluster

You can uninstall the SF Oracle RAC software from the nodes in a cluster. When prompted by the installer, you can select the nodes in a cluster from which you want to uninstall the SF Oracle RAC.

[Figure 25-1](#) displays at a high-level , the steps that are required to uninstall SF Oracle RAC from a cluster having Oracle RAC 11g instances.

Figure 25-1 SF Oracle RAC with Oracle RAC 11g uninstallation flowchart



Preparing to uninstall SF Oracle RAC from an Oracle RAC 11g cluster

Perform the following steps before you uninstall SF Oracle RAC from a cluster having Oracle RAC 11g instances:

- [Stopping Oracle instances](#)

- Backing up the Oracle database
- Moving or removing the Oracle database (optional)
- Uninstalling Oracle RAC 11g (optional)
- Removing repository database
- Stopping the applications that uses CFS (outside of VCS control)
- Unmounting CFS file systems (outside of VCS control)
- Stopping the applications that use VxFS (outside of VCS control)
- Unmounting VxFS file systems (outside of VCS control)
- Stopping VCS

Stopping Oracle instances

If the Oracle RAC 11g instances are under VCS control, you must take the Oracle resources offline.

You must stop the Oracle instances and CRS on the cluster nodes where you want to uninstall SF Oracle RAC. Before stopping the Oracle instances, you must stop the Oracle instances that are dependent on the Oracle service groups.

The following procedure provides the steps to stop the Oracle instances on a two-node cluster. In this procedure, the nodes are galaxy and nebula. Depending on the VCS configuration, the procedure to stop Oracle instances may vary.

To stop the Oracle instances

- 1 Log in as superuser.
- 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:

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

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

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

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

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

For example:

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

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

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

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

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

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, refer to your Oracle RAC 11g documentation.

After uninstalling SF Oracle RAC, you can not access the Oracle database as Veritas Volume Manager and Veritas File System are uninstalled from the cluster.

Moving or removing the Oracle database (optional)

You can move or remove the Oracle database after safely relocating the data as necessary. For instructions on moving or removing the Oracle database, refer to your Oracle RAC 11g documentation.

Uninstalling Oracle RAC 11g (optional)

You can either uninstall Oracle RAC 11g or choose to run Oracle RAC 11g after uninstalling SF Oracle RAC.

To uninstall Oracle RAC 11g

- 1 Log in as oracle user.
- 2 Set the DISPLAY variable. Depending on the shell you use, execute the following command:

```
Bourne Shell (sh or ksh)  $ DISPLAY=host:0.0;export DISPLAY
```

```
C Shell (csh or tcsh)    $ setenv DISPLAY host:0.0
```

- 3 Run the Oracle RAC 11g RAC runInstaller utility:

\$ /cdrom/Disk1/runInstaller
- 4 Select the option to uninstall the Oracle RAC 11g software.
Refer to your Oracle RAC 11g documentation for details on this utility.
- 5 If necessary, uninstall Oracle RAC 11g from the other nodes using this procedure.

Removing repository database

To remove the Storage Foundation for Oracle repository database, use the following procedure. You must have superuser (root) privileges to remove the repository database.

To remove repository database

- 1 Run the following commands to remove the repository configuration from the VCS configuration and deport the repository disk group.

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

- 2 Import the repository disk group using the command:

```
# vxvg import <name_of_disk_group>
```

- 3 Mount the repository volume using the command:

```
# /opt/VRTSdbcom/config/sfua_rep_mount start
```
- 4 Drop the repository database using the command:

```
# /opt/VRTS/bin/sfua_db_config -o dropdb
```

Stopping the applications that uses CFS (outside of VCS control)

All Oracle users must stop all applications using the CFS mounts not under VCS control.

To verify that no processes use the CFS mount point

- 1 Enter the following command:

```
# fuser -c mount_point
```
- 2 Stop any processes that uses a CFS mount point with the mechanism provided by the application.

Unmounting CFS file systems (outside of VCS control)

All Oracle users must unmount any CFS file systems not under VCS control on all nodes.

To unmount CFS file systems not under VCS control

- 1 Log in as root.
- 2 Remove the existing unused modules in the kernel memory and in the library memory.

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

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

```
# umount mount_point
```

Stopping the applications that use VxFS (outside of VCS control)

All Oracle users must stop all applications using the VxFS mounts not under VCS control.

To verify that no processes use the VxFS mount point

- 1 Enter the following command:

```
# fuser -c mount_point
```

- 2 Stop any processes that use a CFS mount point with the mechanism provided by the application.

Unmounting VxFS file systems (outside of VCS control)

All Oracle users must unmount any VxFS file systems not under VCS control on all nodes.

Note: To avoid issues on reboot, you must remove all entries of VxFS from /etc/filesystem folder.

To unmount CFS file systems not under VCS control

- 1 Log in as root.
- 2 Remove the existing unused modules in the kernel memory and in the library memory.

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

- 3 Determine the file systems to unmount by checking the output of the mount file. For example, type:

```
# mount | grep vxfs | grep cluster
```

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

```
# umount mount_point
```

Stopping VCS

Stop VCS to take the service groups on all nodes offline:

To stop VCS

- 1 Log in as superuser on one cluster node.
- 2 Stop VCS on all nodes.

```
# /opt/VRTSvcs/bin/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.

```
# /sbin/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
```

Removing SF Oracle RAC from an Oracle RAC 11g cluster

You can remove the SF Oracle RAC filesets from all nodes in the SF Oracle RAC cluster using the `uninstallsfrac` program. Note that the `uninstallsfrac` program removes all SF Oracle RAC filesets regardless of the Oracle version used.

If you need to uninstall SF Oracle RAC after an incomplete installation, or if the `uninstallsfrac` program is not available in `/opt/VRTS/install`, you may need to use the `uninstallsfrac` program on the SF 5.0 MP3 for Oracle RAC disc.

Removing SF Oracle RAC filesets involve the following tasks:

- [Removing the SF Oracle RAC filesets](#)
- [Removing other configuration files \(optional\)](#)

Removing the SF Oracle RAC filesets

The installer provides you an option to remove Veritas Volume Manager and Veritas Volume Replicator filesets. Note that `uninstallsfrac` program can remove these filesets only if both the boot disk is not under VxVM control and there are no open volumes.

To remove the SF Oracle RAC filesets

- 1 Do one of the following to begin uninstalling:
 - If you can execute commands as superuser on the remote nodes in the cluster with SSH or RSH, run `uninstallsfrac` program on one node to uninstall SF Oracle RAC on all nodes in the cluster.
 - If you cannot execute commands as superuser on remote nodes in the cluster with SSH or RSH, you must run `uninstallsfrac` program on each node in the cluster.

- 2 Navigate to the folder containing the `uninstallsfrac` program.

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

- 3 Start the `uninstallsfrac` program.

```
# ./uninstallsfrac [-rsh]
```

The program specifies the directory where the logs are created and begins with a copyright message.

- 4 Indicate whether or not you want to remove VxVM and VVR filesets from all nodes in the cluster. Enter "y" only if the root disk is outside of VM control.
- 5 If you invoke the uninstallation program from a remote system in the same subnet, enter the system names where you want the uninstallation to take place.

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

- 6 If you invoke the uninstallation program from a node in the SF Oracle RAC cluster, review the cluster information and confirm to uninstall SF Oracle RAC.

VCS configuration files exist on this system with the following information:

```
Cluster Name: rac_cluster101  
Cluster ID Number: 101  
Systems: galaxy nebula  
Service Groups: cvm
```

Do you want to uninstall SFRAC from these systems? [y,n,q] (y)

7 Review the output as the uninstallation program checks the operating system on each system, verifies system-to-system communication, and verifies the system licensing.

8 Review the output as the uninstallation program checks for Storage Foundation for Oracle RAC filesets currently installed on the nodes.

This process involves identifying system uninstall requirements and dependencies between filesets to determine the safety and order of uninstalling filesets.

9 Confirm to uninstall SF Oracle RAC:

```
All SFRAC processes that are currently running will be stopped.  
Are you sure you want to uninstall SFRAC? [y,n,q] (y)
```

10 Make sure that the user VxFS file systems that are currently mounted on each node are unmounted before uninstalling SF Oracle RAC.

```
All user VxFS filesystems that are currently mounted on galaxy  
should be unmounted before uninstalling.
```

See [“Unmounting CFS file systems \(outside of VCS control\)”](#) on page 498.

11 Make sure that each node in the cluster does not have any attached volumes.

12 Review the output as the uninstallation program stops agents and performs verifications on each node to proceed with the uninstallation.

13 Review the output as the uninstallation program stops the SF Oracle RAC processes.

14 Review the output as the uninstallation program uninstalls the SF Oracle RAC filesets.

15 If necessary, review the summary and log files of uninstallation activities.

Removing other configuration files (optional)

You can remove Veritas configuration files and filesets that are left after running the `uninstallsfrac` utility.

To remove residual Veritas configuration files (optional)

- 1 List all VRTS filesets that can be removed.

```
lsllpp -l |grep -i vrts
```

- 2 Run `installp -u` command to remove left over VRTS filesets.

- 3 The following commands moves the residual Veritas configuration files to the `vrts.43` directories:

```
# cd /var
# mkdir vrts.43
# mv *VRTS* vrts.43
# mv vx vrts.43
# cd /var/opt
# mkdir vrts.43
# mv *VRTS* vrts.43
# cd /opt
# mkdir vrts.43
# mv *VRTS* vrts.43
# cd /etc
# mkdir vrts.43
# mv vx *llt* *fen* *gab* *vcs* vcsmmtab vrts.43
```

You can remove the `vrts.43` directories at a later time.

If you reinstall without removing the Veritas configuration files and filesets, the reinstall fails.

Rebooting the nodes

After uninstalling SF Oracle RAC, reboot each node:

```
# /usr/sbin/shutdown -Fr
```


Reference

- [Sample configuration files](#)
- [SF Oracle RAC installation and configuration sample values](#)
- [Response Files](#)
- [Automatic Storage Management](#)
- [Creating a test database](#)
- [High availability agent information](#)

Sample configuration files

This appendix includes the following topics:

- [About sample main.cf files](#)
- [Sample main.cf files for Oracle configurations](#)
- [Sample main.cf files for Oracle replication configurations](#)
- [Sample main.cf files for adding and removing nodes](#)

About sample main.cf files

You can examine the VCS configuration file, main.cf, to verify the SF Oracle RAC installation and configuration.

The main.cf file is located in the folder /etc/VRTSvcs/conf/config.

For the following configuration samples, please note the following:

- All configurations shown here assume that Oracle and CRS binaries are installed on local disks and that they are managed by the operating system. These file systems must be specified in the file /etc/filesystems.
- All sample configurations assume that Oracle binaries are installed on local disks and that they are managed by the operating system. These file systems must be specified in the file /etc/fstab.
- The "cluster" definition in all of the configurations should specify UseFence=SCSI3.

Sample main.cf files for Oracle configurations

Sample main.cf file examples are provided for the following Oracle configurations:

- [Sample main.cf for Oracle RAC 10g without Oracle agent](#)

- [Sample main.cf for Oracle RAC 10g with Oracle agent](#)
- [Sample main.cf for Oracle RAC 10g with Oracle agent and CMC](#)
- [Sample main.cf for Oracle RAC 10g without the Oracle agent and with VxSS](#)
- [Sample main.cf for Oracle RAC 11g without Oracle agent](#)
- [Sample main.cf for Oracle RAC 11g with Oracle agent](#)

Note: After an SF Oracle RAC installation, several sample main.cf file types can be viewed in the following directory: `/etc/VRTSvcs/conf/sample_rac`

Sample main.cf for Oracle RAC 10g without Oracle agent

The following are the configuration details for this Oracle RAC 10g RAC sample main.cf:

- Configuration file name: `10g_simple_main.cf`
- Path to file: `/etc/VRTSvcs/conf/sample_rac/10g_simple_main.cf`
- Use for single Oracle RAC 10g database only
- Has only one parallel service group: `cvm`
- `cvm` group includes PrivNIC and Application resource for CSSD

```
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 (
)
```

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

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
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
```

```
        CVMActivation = sw
    )

    CFSfsckd vxfsckd (

    )

    CVMCluster cvm_clus (
        CVMClustName = rac_cluster101
        CVMNodeId = { galaxy = 0, nebula = 1 }
        CVMTransport = gab
        CVMTimeout = 200
    )

    CVMVxconfigd cvm_vxconfigd (
        Critical = 0
        CVMVxconfigdArgs = { syslog }
    )

    PrivNIC ora_priv (
        Critical = 0
        Device@galaxy = { en1 = 0, en2 = 1 }
        Device@nebula = { en1 = 0, en2 = 1 }
        Address@galaxy = "192.168.12.1"
        Address@nebula = "192.168.12.2"
        NetMask = "255.255.240.0"
    )

    cssd requires ocrvote_mnt
    cssd requires oradata_mnt
    cssd requires ora_priv
    ocrvote_mnt requires ocrvote_voldg
    oradata_mnt requires oradata_voldg
    oradata_voldg requires cvm_clus
    ocrvote_voldg requires cvm_clus
    ocrvote_mnt requires vxfsckd
    oradata_mnt requires vxfsckd
    vxfsckd requires cvm_clus
    cvm_clus requires cvm_vxconfigd
```

Sample main.cf for Oracle RAC 10g with Oracle agent

The following are the configuration details for this Oracle RAC 10g sample main.cf:

- Configuration file name: 10g_main.cf
- Path to file: /etc/VRTSvcs/conf/sample_rac/10g_main.cf
- For multiple Oracle databases
- Has two parallel service groups: cvm and oradb1_grp
- oradb1_grp depends on cvm
- oradb1_grp has Oracle and oradata mount resource

Note: Set the database startup mode to manual in CRS.

```
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 (
)

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Oracle oral (
    Critical = 0
```

```
Sid @galaxy = vrts1
Sid @nebula = vrts2
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

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

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 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1}
    Device@nebula = { en1 = 0, en2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)
```

```
cssd requires ocrvote_mnt
cssd requires ora_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```

Sample main.cf for Oracle RAC 10g with Oracle agent and CMC

The following are the configuration details for this Oracle RAC 10g main.cf:

- Named: 10g_main.cf
- More general purpose, can have multiple Oracle databases
- Has three service groups: cvm, oradb1_grp and oradb2_grp
- oradb1_grp depends on cvm
- oradb1_grp has Oracle and oradata mount resource
- oradb2_grp depends on cvm
- oradb2_grp has Oracle and oradata mount resource

Note: Set the database startup mode to manual in CRS.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)

system nebula (
)
```

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

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 }
    CVMTransport = gab
    CVMTimeout = 200
)
```

```
CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1}
    Device@nebula = { en1 = 0, en2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)

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

group CMC (
    SystemList = { galaxy, nebula }
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

ClusterConnectorConfig CMC_ClusterConfig (
    MSAddress = "mgmtserver1.symantecexample.com"
    MSPort = 14145
    ClusterId = "1145613636"
    ClusterType = "vcs"
    ClusterPort = 14141
    VCSLoggingLevel = "TAG_A"
    Logging = "/opt/VRTScmccc/conf/cc_logging.properties"
    ClusterConnectorVersion = "5.0.1000.0"
)
```

```
Process CMC_ClusterConnector (
    PathName = "/bin/sh"
    Arguments = "/opt/VRTScmccc/bin/cluster_connector.sh"
)

CMC_ClusterConnector requires CMC_ClusterConfig

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata1_mnt (
    Critical = 0
    MountPoint = "/oradata1"
    BlockDevice = "/dev/vx/dsk/oradata1_dg/oradatavol1"
)

CVMVolDg oradata1_voldg (
    CVMDiskGroup = oradata1_dg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
oral requires oradata1_mnt
oradata1_mnt requires oradata1_voldg

group oradb2_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
```

```
Parallel = 1
AutoStartList = { galaxy, nebula }
)

Oracle ora2 (
    Critical = 0
    Sid @galaxy = hr1
    Sid @nebula = hr2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata2_mnt (
    Critical = 0
    MountPoint = "/oradata2"
    BlockDevice = "/dev/vx/dsk/oradata2_dg/oradatavol"
)

CVMVolDg oradata2_voldg (
    CVMDiskGroup = oradata2_dg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
ora2 requires oradata2_mnt
oradata2_mnt requires oradata2_voldg
```

Sample main.cf for Oracle RAC 10g without the Oracle agent and with VxSS

The following is a sample main.cf for Oracle RAC 10g without the Oracle agent and with the VxSS group.

```
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 (
)

group VxSS (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
    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
// }

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
```

```
AutoStartList = { galaxy, nebula }
)

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
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)
```



```
CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1}
    Device@nebula = { en1 = 0, en2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)

cssd requires ocrvote_mnt
cssd requires oradata_mnt
cssd requires ora_priv
ocrvote_mnt requires ocrvote_voldg
oradata_mnt requires oradata_voldg
oradata_voldg requires cvm_clus
ocrvote_voldg requires cvm_clus
ocrvote_mnt requires vxfsckd
oradata_mnt requires vxfsckd
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

// resource dependency tree
//
// group cvm
// {
```

```
// Application cssd
// {
//     CFSSMount ocrvote_mnt
//     {
//         CVMVolDg ocrvote_voldg
//         {
//             CFSSMount oradata_mnt
//             {
//                 CMVVolDg oradata_voldg
//             }
//         }
//     }
//     CFSfsckd vxfsckd
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
// }
//     PrivNIC ora_priv
// }
// }
```

Sample main.cf for Oracle RAC 11g without Oracle agent

The following are the configuration details for this Oracle RAC 11g sample main.cf:

- Configuration file name: 11g_simple_main.cf
- To be used for single 11g+ Oracle database only
- Has only one service group: cvm
- cvm group includes PrivNIC and Application resource for CSSD

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo}
    Administrators = { admin }
```

```
        UseFence = SCSI3
    )

system galaxy (
)

system nebula (
)

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

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
)

CFSMount oradata_mnt (
    Critical = 0
```

```

MountPoint = "/oradata"
BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

CFSfsckd vxfsckd (

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

MultiPrivNIC multi_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1 }
    Device@nebula = { en1 = 0, en2 = 1 }
    Address@galaxy = { "192.168.1.1" = 0, "192.168.1.2" = 0, "192.168.2.1" = 1 }
    Address@nebula = { "192.168.1.3" = 0, "192.168.1.4" = 0, "192.168.2.2" = 1 }
    NetMask = "255.255.240.0"
)

cssd requires ocrvote_mnt
cssd requires oradata_mnt
cssd requires multi_priv

ocrvote_mnt requires ocrvote_voldg

```

```
oradata_mnt requires oradata_voldg

oradata_voldg requires cvm_clus
ocrvote_voldg requires cvm_clus
ocrvote_mnt requires vxfsckd
oradata_mnt requires vxfsckd
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```

Sample main.cf for Oracle RAC 11g with Oracle agent

The following are the configuration details for this Oracle RAC 11g sample main.cf:

- Configuration file name: 11g_main.cf
- More general purpose, can have multiple Oracle databases
- Has three service groups: cvm, oradb1_grp and oradb2_grp
- oradb1_grp depends on cvm
- oradb1_grp has Oracle and oradata mount resource
- oradb2_grp depends on cvm
- oradb2_grp has Oracle and oradata mount resource

Note: Set the database startup mode to manual in CRS .

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
```

```

system nebula (
)

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

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
)

CFSSMount 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 }
CVMTransport = gab
CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

MultiPrivNIC multi_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1 }
    Device@nebula = { en1 = 0, en2 = 1 }
    Address@galaxy = { "192.168.1.1" = 0, "192.168.1.2" = 0, "192.168.2.1" = 1 }
    Address@nebula = { "192.168.1.3" = 0, "192.168.1.4" = 0, "192.168.2.2" = 1 }
    NetMask = "255.255.240.0"
)

cssd requires ocrvote_mnt
cssd requires multi_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
```

```
        Owner = oracle
        Home = "/app/oracle/orahome"
        StartUpOpt = "SRVCTLSTART"
        ShutDownOpt = "SRVCTLSTOP"
    )

CFSMount oradata1_mnt (
    Critical = 0
    MountPoint = "/oradata1"
    BlockDevice = "/dev/vx/dsk/oradata1_dg/oradata1vol"
)

CVMVolDg oradata1_voldg (
    CVMDiskGroup = oradata1_dg
    CVMVolume = { oradata1vol }
    CVMActivation = sw
)

requires group cvm online local firm
ora1 requires oradata1_mnt
oradata1_mnt requires oradata1_voldg

group oradb2_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Oracle ora2 (
    Critical = 0
    Sid @galaxy = hr1
    Sid @nebula = hr2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata2_mnt (
    Critical = 0
    MountPoint = "/oradata2"
```



```
BlockDevice = "/dev/vx/dsk/oradata2_dg/oradatavol"
)

CVMVolDg oradata2_voldg (
    CVMDiskGroup = oradata2_dg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
ora2 requires oradata2_mnt
oradata2_mnt requires oradata2_voldg
```

Sample main.cf files for Oracle replication configurations

This section displays the sample main.cf files for Oracle replication configurations.

Sample main.cf file examples are provided for the following Oracle replication configurations:

- [Sample main.cf for Oracle RAC 10g and CVM/VVR \(primary replication site\)](#)
- [Sample main.cf for Oracle RAC 10g and CVM/VVR \(secondary replication site\)](#)
- [Sample main.cf for Oracle RAC 10g for hardware-based replication primary site](#)
- [Sample main.cf for Oracle RAC 10g for hardware-based replication secondary site](#)

Note: After an SF Oracle RAC installation, several sample main.cf file types can be viewed in the following directory: `/etc/VRTSvcS/conf/sample_rac`

Sample main.cf for Oracle RAC 10g and CVM/VVR (primary replication site)

The following are the configuration details for this Oracle RAC 10g sample main.cf:

- Named: `cvmvvr_primary_main.cf`
- More general purpose, can have multiple Oracle databases
- Path to file: `/etc/VRTSvcS/conf/sample_rac/cvmvvr_primary_main.cf`

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "VVRTypes.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    ClusterAddress = "10.10.10.101"
    Administrators = { admin }
    UseFence = SCSI3
)

remoteclass rac_cluster102 (
    ClusterAddress = "10.11.10.102"
)

heartbeat Icmp (
    ClusterList = { rac_cluster102 }
    Arguments @rac_cluster102 = { "10.11.10.102" }
)

system galaxy (
)

system nebula (
)

group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)

IP gcoip (
    Device = en0
    Address = "10.10.10.101"
```

```
        NetMask = "255.255.240.0"
    )

    NIC csgnic (
        NetworkHosts = { "10.10.12.2", "10.10.12.3" }
        Device = en0
    )

    gcoip requires csgnic
    wac requires gcoip

group RVGgroup (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

CVMVolDg racdata_voldg (
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

requires group cvm online local firm
racdata_rvg requires racdata_voldg

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/ops/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/ops/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/ops/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/ops/bin/cssd-monitor"
    OnlineRetryLimit = 20
)
```

```

CFSfsckd vxfsckd (
)

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
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_priv (
    Critical = 0
    Device = { en1 = 0, en2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)

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

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    ClusterList = { rac_cluster101 = 0, rac_cluster102 = 1 }
    OnlineRetryInterval = 300
    ClusterFailOverPolicy = Manual

```

```
AutoStartList = { galaxy, nebula }
Authority = 1
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)

RVGSharedPri ora_vvr_sharedpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

requires group RVGgroup online local firm
oral requires oradata_mnt
oradata_mnt requires ora_vvr_sharedpri

group rlogowner (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 2
)

IP logowner_ip (
    Device = en0
    Address = "10.10.9.101"
    NetMask = "255.255.240.0"
)

NIC nic (
    Device = en0
    NetworkHosts = { "10.10.12.2", "10.10.12.3" }
)
```

```
RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

Sample main.cf for Oracle RAC 10g and CVM/VVR (secondary replication site)

The following are the configuration details for this Oracle RAC 10g sample main.cf:

- Named: cvmvvr_secondary_main.cf
- More general purpose, can have multiple Oracle databases
- Path to file: /etc/VRTSvcs/conf/sample_rac/cvmvvr_secondary_main.cf

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "VVRTypes.cf"

cluster rac_cluster102 (
    UserNames = { admin = bopHo }
    ClusterAddress = "10.11.10.102"
    Administrators = { admin }
    UseFence = SCSI3
)

remoteclass rac_cluster101 (
    ClusterAddress = "10.10.10.101"
)

heartbeat Icmp (
    ClusterList = { rac_cluster101 }
    Arguments @rac_cluster101 = { "10.10.10.101" }
)

system mercury (
)
```

```
system jupiter (
)

group ClusterService (
  SystemList = { mercury = 0, jupiter = 1 }
  AutoStartList = { mercury, jupiter }
  OnlineRetryLimit = 3
  OnlineRetryInterval = 120
)

Application wac (
  StartProgram = "/opt/VRTSvcs/bin/wacstart"
  StopProgram = "/opt/VRTSvcs/bin/wacstop"
  MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
  RestartLimit = 3
)

IP gcoip (
  Device = en0
  Address = "10.11.10.102"
  NetMask = "255.255.240.0"
)

NIC csgnic (
  Device = en0
  NetworkHosts = { "10.10.12.2", "10.10.12.3" }
)

gcoip requires csgnic
wac requires gcoip

group RVGgroup (
  SystemList = { mercury = 0, jupiter = 1 }
  Parallel = 1
  AutoStartList = { mercury, jupiter }
)

CVMVolDg racdata_voldg (
  CVMDiskGroup = oradatadg
  CVMActivation = sw
)

RVGShared racdata_rvg (
  RVG = rac1_rvg
  DiskGroup = oradatadg
)
```

```

requires group cvm online local firm
racdata_rvg requires racdata_voldg

group cvm (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { mercury, jupiter }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/ops/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/ops/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/ops/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/ops/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster102
    CVMNodeId = { mercury = 0, jupiter = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

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
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0

```



```
CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_privnic (
    Critical = 0
    Device = { en1 = 0, en2 = 1 }
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)

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

group oradb1_grp (
    SystemList = { mercury = 0, jupiter = 1 }
    Parallel = 1
    ClusterList = { rac_cluster101 = 0, rac_cluster102 = 1 }
    OnlineRetryInterval = 300
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { mercury, jupiter }
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

Oracle oral (
    Critical = 0
    Sid @mercury = vrts1
    Sid @jupiter = vrts2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)
```

```
RVGSharedPri ora_vvr_sharedpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

requires group RVGgroup online local firm
oral requires oradata_mnt
oradata_mnt requires ora_vvr_sharedpri

group rlogowner (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoStartList = { mercury, jupiter }
    OnlineRetryLimit = 2
)

IP logowner_ip (
    Device = en0
    Address = "10.11.9.102"
    NetMask = "255.255.240.0"
)

NIC nic (
    Device = en0
    NetworkHosts = { "10.10.12.2", "10.10.12.3" }
)

RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

Sample main.cf for Oracle RAC 10g for hardware-based replication primary site

The following are the configuration details for this Oracle RAC 10g sample main.cf:

- Configuration file name: `srdf_primary_main.cf`
- Path to file: `/etc/VRTSvcs/conf/sample_rac/srdf_primary_main.cf`
- More general purpose, can have multiple Oracle databases

- While this sample is based on SRDF, this template is valid for other supported hardware-based replication options.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SRDFTypes.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    ClusterAddress = "10.11.10.101"
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
)

remoteclass rac_cluster102 (
    ClusterAddress = "10.11.10.102"
)

heartbeat Icmp (
    ClusterList = { rac_cluster102 }
    Arguments @rac_cluster101 = { "10.11.10.102" }
)

system galaxy (
)

system nebula (
)

group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)
```

```

IP gcoip (
    Device = en0
    Address = "10.11.10.101"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = en0
    NetworkHosts = { "10.10.12.2", "10.10.12.3" }
)

gcoip requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//     Application wac
//     {
//         IP gcoip
//         {
//             NIC csgnic
//         }
//     }
//     }
//     }

group Async_ReplicatedDataGroup (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStart = 0
    Parallel = 1
    ClusterList = { rac_cluster101 = 1, rac_cluster102 = 0 }
    AutoStartList = { galaxy, nebula }
    ClusterFailOverPolicy = Auto
)

CFSMount oradataa_cfs (
    MountPoint = "/oradataa"
    BlockDevice = "/dev/vx/dsk/oradataa_dg_galaxy_nebula/oradataa_vol"
)

CVMVolDg oradataa_vol (
    Critical = 0
    CVMDiskGroup = oradataa_dg_galaxy_nebula
    CVMVolume = { oradataa_vol }
)

```

```
CVMActivation = sw
)

Oracle Oracle2 (
    Critical = 0
    Sid @galaxy = RACG1
    Sid @nebula = RACG2
    Owner = oracle
    Home = "/oracle/10g"
    Pfile @galaxy = "/oracle/10g/dbs/initRACG1.ora"
    Pfile @nebula = "/oracle/10g/dbs/initRACG2.ora"
)

SRDF SRDFASync_Group_1 (
    GrpName = RAC_Async
    Mode = Asynchronous
)

requires group cvm online local firm
Oracle2 requires oradataa_cfs
oradataa_cfs requires oradataa_vol
oradataa_vol requires SRDFASync_Group_1

// resource dependency tree
//
//      group Async_ReplicatedDataGroup
//      {
//      Oracle Oracle2
//      {
//          CFSSMount oradataa_cfs
//          {
//              CVMVolDg oradataa_vol
//              {
//                  SRDF SRDFASync_Group_1
//              }
//          }
//      }
//      }
//      }

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
```

```

AutoStartList = { galaxy, nebula }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline stopcrs"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
)

CFSSMount ocrvote_cfs (
    Critical = 0
    MountPoint @galaxy = "/ocrvote"
    MountPoint @nebula = "/ocrvote"
    BlockDevice @galaxy = "/dev/vx/dsk/ocrvote_dg_galaxy_nebula/ocrvote_vol"
    BlockDevice @nebula = "/dev/vx/dsk/ocrvote_dg_galaxy_nebula/ocrvote_vol"
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVolDg ocrvote_vol (
    Critical = 0
    CVMDiskGroup @galaxy = ocrvote_dg_galaxy_nebula
    CVMDiskGroup @nebula = ocrvote_dg_galaxy_nebula
    CVMVolume @galaxy = { ocrvote_vol }
    CVMVolume @nebula = { ocrvote_vol }
    CVMActivation @galaxy = sw
    CVMActivation @nebula = sw
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

```

```

PrivNIC ora_privnic (
    Critical = 0
    Device @galaxy = { en1 = 0, en2 = 1}
    Device @nebula = { en1 = 0, en2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)

cssd requires ocrvote_cf
cssd requires ora_priv
cvm_clus requires cvm_vxconfigd
ocrvote_cfs requires ocrvote_vol
ocrvote_vol requires vxfsckd
vxfsckd requires cvm_clus

// resource dependency tree
//
//      group cvm
//      {
//      Application cssd
//      {
//      CFSMount ocrvote_cfs
//      {
//      CVMVolDg ocrvote_vol
//      {
//      CFSfsckd vxfsckd
//      {
//      CVMCluster cvm_clus
//      {
//      CVMVxconfigd cvm_vxconfigd
//      }
//      }
//      }
//      }
//      }
//      }
//      }
//      PrivNIC ora_priv
//      }
//      }

```

Sample main.cf for Oracle RAC 10g for hardware-based replication secondary site

The following are the configuration details for this Oracle RAC 10g sample main.cf:

- Configuration file name: srdf_secondary_main.cf
- Path to file: /etc/VRTSvcs/conf/sample_rac/srdf_secondary_main.cf
- More general purpose, can have multiple Oracle databases

This is a placeholder file for now, it needs editing.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SRDFTypes.cf"

cluster rac_cluster102 (
    UserNames = { admin = bopHo }
    ClusterAddress = "10.11.10.102"
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
)

remotecluster rac_cluster101 (
    ClusterAddress = "10.11.10.101"
)

heartbeat Icmp (
    ClusterList = { rac_cluster101 }
    Arguments @rac_cluster101 = { "10.11.10.101" }
)

system mercury (
)

system jupiter (
)

group ClusterService (
    SystemList = { mercury = 0, jupiter = 0 }
    AutoStartList = { mercury, jupiter }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
```



```
MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
RestartLimit = 3
)

IP gcoip (
    Device = en0
    Address = "10.11.10.102"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = en0
    NetworkHosts = { "10.10.12.2", "10.10.12.3" }
)

gcoip requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//     Application wac
//     {
//         IP gcoip
//         {
//             NIC csgnic
//         }
//     }
// }

group Async_ReplicatedDataGroup (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoStart = 0
    Parallel = 1
    ClusterList = { rac_cluster102 = 1, rac_cluster101 = 0 }
    AutoStartList = { mercury, jupiter }
    ClusterFailOverPolicy = Auto
)

CFSMount oradataa_cfs (
    MountPoint = "/oradataa"
    BlockDevice = "/dev/vx/dsk/oradataa_dg_mercury_jupiter/oradataa_vol"
)
```

```

        CVMVolDg oradataa_vol (
            Critical = 0
            CVMDiskGroup = oradataa_dg_mercury_jupiter
            CVMVolume = { oradataa_vol }
            CVMActivation = sw
        )

Oracle Oracle2 (
    Critical = 0
    Sid @mercury = RACG1
    Sid @jupiter = RACG2
    Owner = oracle
    Home = "/oracle/10g"
    Pfile @mercury = "/oracle/10g/dbs/initRACG1.ora"
    Pfile @jupiter = "/oracle/10g/dbs/initRACG2.ora"
)

SRDF SRDFASync_Group_1 (
    GrpName = RAC_Async
    Mode = Asynchronous
)

requires group cvm online local firm
Oracle2 requires oradataa_cfs
oradataa_cfs requires oradataa_vol
oradataa_vol requires SRDFASync_Group_1

// resource dependency tree
//
//     group Async_ReplicatedDataGroup
//     {
//         Oracle Oracle2
//         {
//             CFMount oradataa_cfs
//             {
//                 CVMVolDg oradataa_vol
//                 {
//                     SRDF SRDFASync_Group_1
//                 }
//             }
//         }
//     }
// }

```

```
group cvm (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { mercury, jupiter }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline stopcrs"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
)

CFSMount ocrvote_cfs (
    Critical = 0
    MountPoint @mercury = "/ocrvote"
    MountPoint @jupiter = "/ocrvote"
    BlockDevice @mercury = "/dev/vx/dsk/ocrvote_dg_mercury_jupiter/ocrvote_vol"
    BlockDevice @jupiter = "/dev/vx/dsk/ocrvote_dg_mercury_jupiter/ocrvote_vol"
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster102
    CVMNodeId = { mercury = 0, jupiter = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVolDg ocrvote_vol (
    Critical = 0
    CVMDiskGroup @mercury = ocrvote_dg_mercury_jupiter
    CVMDiskGroup @jupiter = ocrvote_dg_mercury_jupiter
    CVMVolume @mercury = { ocrvote_vol }
    CVMVolume @jupiter = { ocrvote_vol }
    CVMActivation @mercury = sw
    CVMActivation @jupiter = sw
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
```

```

    CVMVxconfigdArgs = { syslog }
)

PrivNIC ora_privnic (
    Critical = 0
    Device @mercury = { en1 = 0, en2 = 1}
    Device @jupiter = { en1 = 0, en2 = 1}
    Address@mercury = "192.168.12.3"
    Address@jupiter = "192.168.12.4"
    NetMask = "255.255.240.0"
)

cssd requires ocrvote_cf
cssd requires ora_priv
cvm_clus requires cvm_vxconfigd
ocrvote_cfs requires ocrvote_vol
ocrvote_vol requires vxfsckd
vxfsckd requires cvm_clus

// resource dependency tree
//
//     group cvm
//     {
//     Application cssd
//     {
//         CFSMount ocrvote_cfs
//         {
//             CVMVolDg ocrvote_vol
//             {
//                 CFSfsckd vxfsckd
//                 {
//                     CVMCluster cvm_clus
//                     {
//                         CVMVxconfigd cvm_vxconfigd
//                     }
//                 }
//             }
//         }
//     }
//     PrivNIC ora_priv
//     }
//     }

```

Sample main.cf files for adding and removing nodes

This section displays sample main.cf files for adding and removing nodes.

The following sample main.cf file examples are provided for adding and removing nodes in a configuration:

- [Sample main.cf for adding and removing an Oracle RAC 10g node](#)
- [Sample main.cf for adding and removing an Oracle RAC 11g node](#)

Note: After an SF Oracle RAC installation, several sample main.cf file types can be viewed in the following directory: `/etc/VRTSvcs/conf/sample_rac`

Sample main.cf for adding and removing an Oracle RAC 10g node

Changes to the Oracle RAC 10g sample main.cf for adding a node are highlighted in bold for the saturn node.

```
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 (
)

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
```

```
AutoFailOver = 0
Parallel = 1
AutoStartList = { galaxy, nebula, saturn }
)
```

Note: In the above oradb1_grp section of this main.cf, the saturn node 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"
)
```

Note: In the above Oracle oral section of this main.cf, the saturn node has been added.

```
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
```

```
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 above group cvm and CVMCluster cvm_clu section of this main.cf, the saturn node has been added.

```
PrivNIC ora_priv (  
    Critical = 0  
    Device@galaxy = { en1 = 0, en2 = 1}  
    Device@nebula = { en1 = 0, en2 = 1}  
    Device@saturn = { en1 = 0, en2 = 1}  
    Address@galaxy = "192.168.12.1"  
    Address@nebula = "192.168.12.2"  
    Address@saturn = "192.168.12.3"  
    NetMask = "255.255.240.0"  
)
```

Note: In the above group PrivNIC section of this main.cf, the saturn node has been added.

```
cssd requires ocrvote_mnt  
cssd requires ora_priv  
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
vxfsckd requires cvm_clus  
cvm_clus requires cvm_vxconfigd
```

For a sample main.cf for removing an Oracle RAC 10g node, use the above sample main.cf and remove the highlighted saturn node changes.

Sample main.cf for adding and removing an Oracle RAC 11g node

Changes to the sample main.cf for adding a node are highlighted in bold for the saturn node.


```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Sample configuration files
    Sample main.cf files for Oracle configurations
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)

system nebula (
)
system saturn (
)

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 (  
)
```

Note: In the above group cvm section of this main.cf, the saturn node has been added.

```
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 above CVMCluster cvm_clus section of this main.cf, the saturn node has been added.

```
MultiPrivNIC multi_priv (  
    Critical = 0  
    Device@galaxy = { en1 = 0, en2 = 1}  
    Device@nebula = { en1` = 0, en2 = 1}
```

```

Device@saturn = { en1 = 0, en2 = 1}
Address@galaxy = { "192.168.1.1" = 0, "192.168.1.2" = 0, "192.168.2.1" = 1 }
Address@nebula = { "192.168.1.3" = 0, "192.168.1.4" = 0, "192.168.2.2" = 1 }
Address@saturn = { "192.168.1.5" = 0, "192.168.1.6" = 0, "192.168.2.3" = 1 }
NetMask = "255.255.240.0"
)

```

Note: In the above MultiPrivNIC section of this main.cf, the saturn node has been added.

```

cssd requires ocrvote_mnt
cssd requires multi_priv
cssd requires oradata_mnt
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata1_mnt (
    Critical = 0
    MountPoint = "/oradata1"
    BlockDevice = "/dev/vx/dsk/oradata1_dg/oradatavol1"
)

```

```

    )

CVMVolDg oradata1_voldg (
    CVMDiskGroup = oradata1_dg
    CVMVolume = { oradata1_vol }
    CVMActivation = sw
)

requires group cvm online local firm
ora1 requires oradata1_mnt
oradata1_mnt requires oradata1_vold

```

Note: In the above group oradb1_grp and Oracle ora1 section of this main.cf, the saturn node has been added.

```

group oradb2_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn =2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

Oracle ora2 (
    Critical = 0
    Sid @galaxy = hr1
    Sid @nebula = hr2
    Sid @saturn = hr3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata2_mnt (
    Critical = 0
    MountPoint = "/oradata2"
    BlockDevice = "/dev/vx/dsk/oradata2_dg/oradata1_vol"
)

CVMVolDg oradata2_voldg (

```

```
    CVMDiskGroup = oradata2_dg
    CVMVolume = { oradata2_vol }
    CVMActivation = sw
)

requires group cvm online local firm
ora2 requires oradata2_mnt
oradata2_mnt requires oradata2_voldg
```

Note: In the above group oradb2_grp and Oracle ora2 section of this main.cf, the saturn node has been added.

For a sample main.cf for removing an Oracle RAC 11g node, use the above sample main.cf and remove the highlighted saturn node changes.

SF Oracle RAC installation and configuration sample values

This appendix includes the following topics:

- [SF Oracle RAC installation and configuration information](#)
- [Installation and configuration worksheets](#)

SF Oracle RAC installation and configuration information

The SF Oracle RAC installation and configuration program prompts you for information about SF Oracle RAC. This installation and configuration program provides default values for some information which you can choose to use.

To expedite your installation and configuration process:

- Review the following information before you begin.
- Assemble the required information for your installation and configuration.
- Use the sample values in the worksheets as examples of the information required for a SF Oracle RAC installation and configuration.
- Have the required information for your configuration at hand during the installation process.

[Table B-1](#) below lists both the tasks and the information required to perform the tasks for a SF Oracle RAC installation and configuration.

Table B-1 SF Oracle RAC installation and configuration information

Task	Information
Install SF Oracle RAC packages	See “ SF Oracle RAC filesets information ” on page 561.
Configure the Veritas Cluster server component	See “ Veritas Cluster Server component information ” on page 561. This configuration step is required.
Configure Cluster Volume Manager	See “ Cluster Volume Manager information ” on page 562. This configuration step is required.
Configure SF Oracle RAC clusters in secure mode	See “ SF Oracle RAC secure mode cluster information ” on page 563. This configuration step is optional.
Add the SF Oracle RAC users	See “ SF Oracle RAC add user information ” on page 563. This step is optional.
Configure the Cluster Management Console cluster connector	This configuration step is optional.
Configure the Cluster Management console	This configuration step is optional.
Configure SMTP email notification	See “ SMTP email notification information ” on page 564. This configuration step is optional.
Configure SNMP trap notification	See “ SNMP trap notification information ” on page 565. This configuration step is optional.
Configure global clusters	See “ Global cluster information ” on page 565. This configuration step is optional.
Configure Veritas Volume Replicator	See “ Veritas Volume Replicator information ” on page 566. This configuration step is optional.
Configure I/O fencing	See “ I/O fencing information ” on page 566. This configuration step is optional.

SF Oracle RAC filesets information

Table B-2 displays the information required to install SF Oracle RAC filesets.

Table B-2 SF Oracle RAC filesets information

Information	Examples
System names on which to install SF Oracle RAC	Example: galaxy, nebula
License keys	License keys could be one of the following types: - Valid license keys for each system in the cluster - Valid site license key - Valid demo license key If you want to configure Veritas Volume Replicator to enable disaster recovery, you must enter appropriate license keys. See “Obtaining SF Oracle RAC license keys” on page 75. Note: There are SF Oracle RAC and SF Oracle RAC Disaster Recovery and High Availability options for license keys.
Do you want to install required SF Oracle RAC filesets or all SF Oracle RAC filesets?	Install only the required filesets if you do not want to configure any optional components or features. Default option is to install all filesets.

Veritas Cluster Server component information

Table B-3 displays the information required to configure the Veritas Cluster Server component.

Table B-3 Veritas Cluster Server component information

Information	Example
Name of the cluster	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 (_). Example: rac_cluster101
Unique ID number for the cluster	Number in the range of 0-65535. Within the site containing the cluster, each cluster must have a unique ID. 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. The interface names associated with each NIC for each network link must be the same on all nodes. For example: ■ en1 ■ en2 Do not use the network interface card that is used for the public network, which is typically en0 and en1.

Cluster Volume Manager information

[Table B-4](#) displays the information required to configure the Cluster Volume Manager.

Table B-4 Cluster Volume Manager information

Information	Examples
Do you want to configure enclosure-based naming scheme?	Enter Yes to configure an enclosure-based naming scheme. Enter No to not configure an enclosure-based naming scheme. Note: Dynamic Multipathing (DMP) is a prerequisite for enclosure-based naming schemes

SF Oracle RAC secure mode cluster information

Table B-5 displays the information required to configure SF Oracle RAC clusters in secure mode.

Note: Configuring SF Oracle RAC in secure mode is optional.

Table B-5 SF Oracle RAC secure mode cluster information

Information	Examples
Which mode do you want to choose to configure Authentication Service?	The installer provides you the following three modes to configure Authentication Service in the SF Oracle RAC clusters: ■ Automatic mode ■ Semiautomatic mode using encrypted files ■ Semiautomatic mode without using encrypted files Default option is automatic mode. See “Preparing to configure the clusters in secure mode” on page 58.
Host name of the Symantec Product Authentication Service Root Broker System	Example: venus

SF Oracle RAC add user information

Table B-6 displays the information required to add VCS users. If you are going to configure SF Oracle RAC cluster in secure mode, you will need to add VCS users.

Note: Adding VCS users is optional.

Table B-6 SF Oracle RAC add user information	
Information	Examples
User name	Example: smith
User password	Enter the password at the prompt.
User privilege	Users have three levels of privileges: ■ A=Administrator ■ O=Operator ■ G=Guest Example: A

SMTP email notification information

[Table B-7](#) displays the information required to configure SMTP email notification.

Note: Configuring SMTP email notification is optional.

Table B-7 SMTP email notification information	
Information	Example
Domain-based address of the SMTP server	The SMTP server sends notification email about the events within the cluster. Example: smtp.symantecexample.com
Email address of each SMTP recipient to be notified	Example: john@symantecexample.com
Minimum severity of events for SMTP email notification	Events have four levels of severity: ■ I=Information ■ W=Warning ■ E=Error ■ S=SevereError Example: E

SNMP trap notification information

[Table B-8](#) displays the information required to configure SNMP trap notification.

Note: Configuring SNMP trap notification is optional.

Table B-8 SNMP trap notification information

Information	Examples
Port number for the SNMP trap daemon	Enter a port number. Default port number is 162.
Machine name for each SNMP console	Example: saturn
Minimum severity of events for SNMP trap notification	Events have four levels of severity: ■ I=Information ■ W=Warning ■ E=Error ■ S=SevereError Example: E

Global cluster information

[Table B-9](#) displays the information required to configure global clusters.

Note: Global clusters are an optional feature that requires a license to enable. If you are not using a global cluster, the information needed for the global cluster configuration is not required.

Table B-9 Global cluster information

Information	Example
Name of the public NIC	You must specify appropriate values for the NIC when you are prompted. Example: en0 and en1
Virtual IP address of the NIC	You must specify appropriate values for the virtual IP address when you are prompted. Example: 10.10.12.1

Table B-9 Global cluster information (*continued*)

Information	Example
Netmask for the virtual IP address	You must specify appropriate values for the netmask when you are prompted. Example: 255 . 255 . 255 . 0

Veritas Volume Replicator information

[Table B-10](#) displays the information required to configure Veritas Volume Replicator (VVR).

Note: Configuring Veritas Volume Replicator is optional. Veritas Volume Replicator is an advanced option and users will be prompted for this configuration information only if they have entered a VVR license.

Table B-10 Veritas Volume Replicator information

Information	Examples
Frequency of VVR statistics collection.	Enter a number to represent the frequency in seconds. Default is 10 seconds.
Number of days to preserve the collected statistics.	Enter a number to represent the number of days. Default is 3 days.

For additional Veritas Volume Replicator information, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

I/O fencing information

[Table B-11](#) displays the information required to configure I/O fencing.

Note: Configuring I/O fencing is optional.

Table B-11 I/O fencing information

Information	Examples
The name of three disks that form the coordinator disk group.	The following are examples of disk names: ■ rhdisk75 ■ rhdisk76 ■ rhdisk77
The names for each disk in the coordinator disk group (if using DMP).	The following are examples: ■ /dev/vx/rdmp/rhdisk75 ■ dev/vx/rdmp/rhdisk76 ■ dev/vx/rdmp/rhdisk77

Installation and configuration worksheets

Symantec recommends using the worksheets below to record values for your systems prior to beginning your SF Oracle RAC installation and configuration.

SF Oracle RAC worksheet

[Table B-12](#) contains sample values used when installing and configuring SF Oracle RAC. Enter the SF Oracle RAC values for your systems in the table below.

Table B-12 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	
Host names for added/removed node	saturn	
RAC License key	xxxxxxxxxx	
Primary cluster name	rac_cluster101	
Primary cluster ID number	101	
Private network links	en1, en2	

Table B-12 SF Oracle RAC worksheet (*continued*)

Installation information	Sample value	Assigned value
Cluster Manager NIC (Primary NIC)	en0	
Cluster Manager IP	10.10.12.1, 10.10.12.2	
Netmask for the virtual IP address	255.255.240.0	
Domain-based address of SMTP server	smtp.symantecexample.com	
Email address of SMTP notification recipients	admin@symantecexample.com	
SNMP trap daemon port number the console	162	
System name for the SNMP console	system2	
Minimum severity level of events for SMTP notification	i	
CVM cluster reconfiguration timeout	200 (default)	
CVM enclosure-based naming	yes	
Vxfen disks	disk01, disk02, disk03	
Vxfen diskgroup	vxencoordg	

Oracle worksheet

Table B-13 contains sample values used when installing and configuring Oracle. Enter the Oracle values for your systems in the table below.

Table B-13 Oracle worksheet

Installation information	Sample value	Assigned value
Oracle user ID	100	
Oracle user	oracle	

Table B-13 Oracle worksheet (*continued*)

Installation information	Sample value	Assigned value
Oracle group ID	1000	
Oracle install group	oinstall	
DBA group	dba	
CRS installer path	/cdrom/Oracle10g/ clusterware/Disk1	
database installer path	DB_Disk	
Private IP address	192.168.12.1-192.168.12.2	
Private hostnames (set in /etc/hosts)	galaxy_priv, nebula_priv	
VCS resource name for PrivNIC	ora_priv	
Virtual IP addresses	10.10.11.1-10.10.11.2	
Virtual hostnames (set in /etc/hosts)	galaxy-vip, nebula-vip	
export display	10.20.12.150:0.0 or fully qualified name (with domain)	
Disks for binary mounts	HDS0_10 for first system, HDS0_11 for second system, ...	
Disk groups for binary mounts	orabindg_galaxy, orabindg_nebula	
Volumes for binary mounts	orabinvol_galaxy, orabinvol_nebula	
Volume options for binary mounts	(none)	
Mount point for binaries	/app (Non-shared)	
CRS Home directory	/app/crshome	
ORACLE Base directory	/app/oracle	
ORACLE Home directory	/app/oracle/orahome	

Table B-13 Oracle worksheet (*continued*)

Installation information	Sample value	Assigned value
Disk for OCR and Votedisk (recommend mirroring)	HDS0_20 IBM0_20 EMC0_20	
Disk group for OCR and Votedisk	ocrvotedg (Shared)	
Volumes for OCR and Votedisk	ocrvotevol. If separate, ocrvol & votevol	
Volume options OCR and Votedisk	nmirror=3	
Mount point for Votedisk and OCR	/ocrvote (Shared)	
OCR file	/ocrvote/ocr	
Votedisk file	/ocrvote/vote	
VCS resource names for Votedisk and OCR	ocrvote_voldg, ocrvote_mnt	
Mount point for archive logs	/oradata	
Disk for Oracle Data	HDS0_30 IBM0_30 EMC0_30 HDS1_30	
Disk group for Oracle Database	oradatadg (Shared)	
Volume for Oracle Database	oradatavol	
Volume options 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	vrts.veritas.com (database name.domain)	
SID prefix	ora1	

Replicated cluster using VVR worksheet

Table B-14 contains sample values used when installing and configuring CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the table below.

Table B-14 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	
Secondary Cluster Logowner IP	10.11.9.102	
RVG Name	rac1_rvg	
Global Database Name	vrts	
Database Resource Name	ora1	
Database Group Name (depends on cvm, includes resources oracle agent etc.)	oradb1_grp	
Srl Volume Name 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 IPs for Secondary Cluster	192.168.12.3 - 192.168.12.4	

Replicated cluster using SRDF worksheet

Table B-14 contains sample values used when installing and configuring CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the table below.

Table B-15 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	
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 (depends on cvm, includes resources oracle agent etc.)	oradb1_grp	
Srl Volume Name Name	rac1_srl	
Resolvable Virtual Hostname of the cluster on Primary Site (for VVR)	rac_clus101_priv	

Table B-15 Replicated cluster using SRDF worksheet (*continued*)

Installation information	Sample value	Assigned value
Resolvable Virtual Hostname of the cluster on Secondary Site (for VVR)	rac_clus102_priv	
Private IPs for Secondary Cluster	192.168.12.3 - 192.168.12.4	

Response Files

This appendix includes the following topics:

- [About Response files](#)
- [Response file installation procedures](#)
- [Response file syntax](#)
- [Install-only response file example](#)
- [Configure response file example](#)
- [Response file variable definitions](#)

About Response files

Response files can be used to standardize and automate installations on multiple cluster systems. By creating a response file and specifying this file when you start, you can automate the SF Oracle RAC installation. This automated installation method is useful if multiple installations on similarly configured systems are required.

At the end of the SF Oracle RAC installation, the installation utility creates the following informational files:

- A log file containing executed system commands and output.
- A summary file containing the output of the installation scripts.
- Response files to be used with the `-responsefile` option of the installer.

The installation utility also indicates where the above files are stored. The response files are saved and stored at:

```
/opt/VRTS/install/logs/  
installfrac-<installernumber>/installfrac-<installernumber>.response
```

In addition to the SF Oracle RAC response files, Oracle and Veritas Cluster Server (VCS) response files are also available.

Note: Oracle provides its own response files and procedures for their use in performing an Oracle installation. Refer to the appropriate Oracle documentation for information on these response files and procedures.

Veritas Cluster Server response files

Veritas Cluster Server (VCS) has its own response files and procedures for their use in a VCS installation. If the user is only installing VCS, refer to the appropriate VCS documentation for information on these response files and procedures. For an SF Oracle RAC installation that includes VCS, follow the procedures outlined in this appendix.

Response file installation procedures

Running the `installsfrac` program with the `-responsefile` option is useful not only for installing and configuring SF Oracle RAC within a secure environment, but also for conducting installations to other clusters as well. Typically, you can use the response files generated during the installation of SF Oracle RAC for a re-installation on that cluster, or for installation on other clusters.

When using response files for installation to other clusters, you copy the file to a system in another cluster and manually edit the file to contain the appropriate values.

This section describes the following two procedures:

- Re-installation procedure using the response files
- Installation procedure using the response files on other clusters

Re-installation using the response files

When you perform an interactive installation of SF Oracle RAC using the installer, response files are automatically generated in the following directory:

```
/opt/VRTS/install/logs/installsfrac-<installernumber>/ \
installsfrac-<installernumber>.response file
```

The response file name is generated as `installsfrac-<installernumber>.response`, where the number is random. You can use the response files for future installations on the same machine.

Perform the following steps for a response file re-installation.

Response file re-installation

- 1 Follow the directions in the SF Oracle RAC installation and configuration chapters to perform a single installation on a single cluster (nodes and their coordination points) using the desired settings.

After performing a single installation on a single cluster, you will have also created the following two response file types:

- An install-only response file
- A configure response file

- 2 Locate the response files in the /opt/VRTS/install/logs directory. Identify the install-only response file and the configure response file using the procedures described in the following step.

Note: If the directory contains many response files generated from several earlier installations, then determine which response files from which installation you wish to use. After making this initial determination, then proceed to identify the install-only response file and configure response file using the procedures described in the following step.

- 3 The install-only response file is created during the SF Oracle RAC installation process, and the configure response file is created after an installation reboot and during the subsequent configuration process.

To determine if the response file is an install-only response file type, use the `grep` command on the response file to search for this string:

```
$CPI::CFG{OPT}{INSTALLONLY}=1;
```

If this string appears in the `grep` command output, then this response file is an install-only response file type.

To determine if the response file is a configure response file type, use the `grep` command on the response file to search for this string:

```
$CPI::CFG{OPT}{CONFIGURE}=1;
```

If this string appears in the `grep` command output, then this response file is a configure response file type.

Note: Both the install-only response file and the configure response file are used for this procedure.

- 4 Use the generated install-only response file to automatically run the cluster installation based on the settings in that generated response file.

To start the installsfrac script, enter the following command on that cluster:

```
# installsfrac -responsefile /opt/VRTS/install/logs/ \
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Note: The installsfrac-<installernumber>.response is the install-only response file in this step.

By default, the installsfrac program uses SSH for remote communication. However, to use RSH, specify the -rsh option with the installsfrac program by using the following command:

```
# installsfrac -rsh -responsefile /opt/VRTS/install/logs/ \
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Additionally, you must edit the response file to include "\$CPI::CFG{OPT}{RSH}=1". Using vi or another text editor, edit the response file and add this line. For additional information about response file syntax:

See [“Response file syntax”](#) on page 583.

- 5 After the system reboots, use the configure response file to automatically run the cluster configuration based on the settings in the configure response file. Enter the following command on that cluster:

```
# installsfrac -responsefile /opt/VRTS/install/logs/ \
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Note: The installsfrac-<installernumber>.response is the configure response file in this step.

By default, the installsfrac program uses SSH for remote communication. However, to use RSH, specify the -rsh option with the installsfrac program by using the following command:

```
# installsfrac -rsh -responsefile /opt/VRTS/install/logs/ \
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Installation using the response file on other clusters

When you perform an interactive installation of SF Oracle RAC using the installer, response files are automatically generated in the following directory:

```
/opt/VRTS/install/logs/installsfrac-<installernumber>/ \
installsfrac-<installernumber>.response file
```

The response file name is generated as `installsfrac-<installernumber>.response`, where the `installernumber` is random. You can use this response file for future installations on other clusters.

Perform the following steps for a response file installation on other clusters:

Response file installation on other clusters

- 1 Follow the directions in the SF Oracle RAC installation and configuration chapters to perform a single installation on the first cluster using the desired settings.

After performing a single installation on a single cluster, you will have the following two response file types:

- An install-only response file
- A configure response file

The install-only response file is created during the SF Oracle RAC installation process, and the configure response file is created after an installation reboot and during the subsequent configuration process.

Note: Both the install-only response file and the configure response file are used for this procedure.

- 2 After the installation on the first cluster is complete, review the log and summary files to ensure an error-free installation.

The installer stores these files within the `/opt/VRTS/install/logs` directory.

3 Proceed to review the response files located at:

```
/opt/VRTS/install/logs/  
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Determine which response file is the install-only response file and which response file is the configure response file within this directory.

To determine if the response file is an install-only response file type, use the `grep` command on the response file to search for this string:

```
$CPI::CFG{OPT}{INSTALLONLY}=1;
```

If this string appears in the `grep` command output, then this response file is an install-only response file type.

To determine if the response file is a configure response file type, use the `grep` command on the response file to search for this string:

```
$CPI::CFG{OPT}{CONFIGURE}=1;
```

If this string appears in the `grep` command output, then this response file is a configure response file type

Note: If the directory contains many response files generated from several earlier installations, then determine which response files from which installation you wish to use. After making this initial determination, then proceed to identify the install-only response file and configure response file using the procedures described in this step.

- 4 Using a method of your choice (for example, by using NFS, ftp, or a CD-ROM), place a copy of both response file types (install-only response file and configure response file) in a directory such as `/tmp` on the next cluster to install SF Oracle RAC.
- 5 On that next cluster, edit the copied response files using a text editor (vi).
In particular, modify the hostnames in the response file to hostnames used in the new cluster.

Note: The response files themselves contain information about the variables within them. This information is helpful in understanding the variables listed in the response file that may require editing.

- 6 For demo or site licenses, the license key in the response file need not be changed. When license keys are “node-locked” to specific cluster nodes, you must edit the license key.
- 7 Change the permissions on the file to 700.
- 8 Use the edited install-only response file to automatically run the next cluster installation, based on the settings in the edited install-only response file.

To start the installsfrac script, enter the following command on that cluster:

```
# installsfrac -responsefile /opt/VRTS/install/logs/\
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Note: The installsfrac-<installernumber>.response is the edited install-only response file in this step.

By default, the installsfrac program uses SSH for remote communication. However, to use RSH, specify the -rsh option with the installsfrac program by using the following command:

```
# installsfrac -rsh -responsefile /opt/VRTS/install/logs/\
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Additionally, you must edit the response file to include "\$CPI::CFG{OPT}{RSH}=1". Using vi or another text editor, edit the response file and add this line. For additional information about response file syntax:

See “[Response file syntax](#)” on page 583.

- 9 After the system reboots, use the edited configure response file to automatically run the next cluster configuration based on the settings in the edited configure response file.

Enter the following command on that cluster:

```
# installsfrac -responsefile /opt/VRTS/install/logs/\
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

Note: The installsfrac-<installernumber>.response is the edited configure response file in this step.

- 10 By default, the `installsfrac` program uses SSH for remote communication. However, to use RSH, specify the `-rsh` option with the `installsfrac` program by using the following command:

```
# installsfrac -rsh -responsefile /opt/VRTS/install/logs/\
installsfrac-<installernumber>/installsfrac-<installernumber>.response
```

- 11 Repeat steps 2 through 10 for each new cluster installation.

Response file syntax

The Perl statement syntax 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"];
```

Install-only response file example

The following is an example of an install-only response file with some of its variable definitions:

```
#
# installsfrac configuration values:
#
$CPI::CFG{KEYS}{nebula}=[ qw(XXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX) ];
$CPI::CFG{KEYS}{galaxy}=[ qw(XXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX) ];
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{INSTALLONLY}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SYSTEMS}=[ qw(nebula galaxy) ];
$CPI::CFG{UPI}="SFRAC";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";

1;
```

```
# # $CPI::CFG{DONOTINSTALL} is an optional one dimensional list variable.
# This variable defines a list of optional filesets not to be installed on
# the systems.
#
# $CPI::CFG{DONOTREMOVE} is an optional one dimensional list variable.
# This variable defines a list of filesets not to be removed from the
# systems during uninstall.
#
# $CPI::CFG{KEYS}{<SYSTEM>} is an optional two dimensional list variable.
# This variable defines a list of license keys to be registered on the
# systems during install.....
```

Note: A user could modify the hostnames in the above example, (galaxy and nebula) in this response file to other hostnames. The user could then use this modified response file for installation on another system.

Note: The variables in the above response file are defined in [Table C-1](#) below. The VCS variables in the above response file are defined in Veritas Cluster Server Installation Guide.

Configure response file example

The following is response file example of a configure response file with some of its variable definitions:

```
#
# installsfrac configuration values:
#
$CPI::CFG{OBC_AGENTFQHN}{nebula}="nebula.veritas.com";
$CPI::CFG{OBC_AGENTFQHN}{galaxy}="galaxy.veritas.com";
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{CONFIGURE}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SFRAC_CONFIG_SFRAC_OPT}=1;
$CPI::CFG{SFRAC_MAIN_MENU_OPT}=2;
$CPI::CFG{SYSTEMS}=[ qw(nebula galaxy) ];
$CPI::CFG{UPI}="SFRAC";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=19798;
```



```
$CPI::CFG{VCS_CLUSTERNAME}="sfrac_97_98";
$CPI::CFG{VCS_LLTLINK1}{nebula}="en2";
$CPI::CFG{VCS_LLTLINK1}{galaxy}="en2";
$CPI::CFG{VCS_LLTLINK2}{nebula}="en3";
$CPI::CFG{VCS_LLTLINK2}{galaxy}="en3";
$CPI::CFG{VCS_USERENPW}=[ qw(hKLdKFkHLgLLjTlfKI) ];
$CPI::CFG{VCS_USERNAME}=[ qw(admin) ];
$CPI::CFG{VCS_USERPRIV}=[ qw(Administrators) ];

1;
# # $CPI::CFG{DONOTINSTALL} is an optional one dimensional list variable.
# This variable defines a list of optional filesets not to be installed
# on the systems.
#
# $CPI::CFG{DONOTREMOVE} is an optional one dimensional list variable.
# This variable defines a list of filesets not to be removed from
# the systems during uninstall.
#
# $CPI::CFG{KEYS}{<SYSTEM>} is an optional two dimensional list variable.
# This variable defines a list of license keys to be registered
# on the systems during install. ....
```

Note: A user could modify the hostnames in the above example, (galaxy and nebula) in this response file to other hostnames. The user could then use this modified response file for installation on another system.

Note: The variables in the above response file are defined in [Table C-1](#) below. The VCS variables in the above response file are defined in Veritas Cluster Server Installation Guide.

Response file variable definitions

[Table C-1](#) lists the variables used in the response files and their definitions.

Note that while some variables are labeled as required and others as optional, some of the optional variables, if used, make it necessary to define other optional variables.

Table C-1 Response file variables

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{DONOTINSTALL}	List	Optional	This variable defines a list of optional packages not to be installed on the systems.
\$CPI::CFG{DONOTREMOVE}	List	Optional	This variable defines a list of packages not to be removed from the systems during uninstall.
\$CPI::CFG{KEYS}{<SYSTEM>}	List	Optional	This variable defines a list of license keys to be registered on the systems during install.
\$CPI::CFG{OPT_LOGPATH}	Scalar	Optional	This variable defines the location where log files are copied to following an install. The default location is /opt/VRTS/install/logs.
\$CPI::CFG{OPT}{CONFIGURE}	Scalar	Optional	This variable performs configuration and startup of a product that has previously been installed using the -installonly option.
\$CPI::CFG{OPT}{INSTALL}	Scalar	Optional	This variable designates a standard installation including licensing, install, configuration and startup.
\$CPI::CFG{OPT}{INSTALLONLY}	Scalar	Optional	This variable installs packages but does not configure or start the product.
\$CPI::CFG{OPT}{KEYFILE}	Scalar	Optional	This variable defines the location of an ssh keyfile used to communicate with all remote systems.
\$CPI::CFG{OPT}{LICENSE}	Scalar	Optional	This variable licenses the product only.

Table C-1 Response file variables (*continued*)

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{OPT}{NOLIC}	Scalar	Optional	This variable installs the product without requiring entry of a license key.
\$CPI::CFG{OPT}{PATCHPATH}	Scalar	Optional	This variable 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.
\$CPI::CFG{OPT}{PKGPATH}	Scalar	Optional	This variable 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.
\$CPI::CFG{OPT}{RSH}	Scalar	Optional	This variable uses rsh instead of ssh as the communication method between systems.
\$CPI::CFG{OPT}{TMPPATH}	Scalar	Optional	This variable defines the location where a working directory is created to store temporary files and packages needed during the install. The default location is /var/tmp.
\$CPI::CFG{OPT}{UNINSTALL}	Scalar	Optional	This variable uninstalls a product.
\$CPI::CFG{STOPFAIL_ALLOW}	Scalar	Optional	This variable decides whether or not installer will proceed when meet # failures of stopping process or unloading drivers.
\$CPI::CFG{SYSTEMS}	List	Required	This variable defines the list of systems on which the product is to be installed, uninstalled, or configured.

Table C-1 Response file variables (*continued*)

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{SYSTEMSCFG}	List	Optional	This variable defines the list of systems to be recognized for configuration when a secure environment prevents all systems from being installed at once.
\$CPI::CFG{UPI}	Scalar	Required	This variable is an abbreviation that defines the product to be installed, uninstalled, or configured.
\$CPI::CFG{AT_ROOTDOMAIN}	List	Optional	This variable defines the name of the system where the root broker is installed.
\$CPI::CFG{SFRAC_CONFIG_SFRAC_OPT}	Scalar	Required	This variable defines the option that needs to be chosen after selecting to configure SFRAC.
\$CPI::CFG{SFRAC_MAIN_MENU_OPT}	Scalar	Required	This variable defines the option that needs to be chosen in the main menu of SFRAC.
\$CPI::CFG{VCS_CLUSTERID}	Scalar	Required	This variable must be an integer between 0 and 65535 which uniquely identifies the cluster.
\$CPI::CFG{VCS_CLUSTERNAME}	Scalar	Required	This variable defines the name of the cluster.
\$CPI::CFG{VCS_CSGNETMASK}	Variable	Optional	This variable defines the Netmask of the virtual IP address to be used by the Cluster Manager (Web Console).
\$CPI::CFG{VCS_CSGNIC}	Scalar	Optional	This variable defines the virtual IP address to be used by the Cluster Manager (Web Console).

Table C-1 Response file variables (*continued*)

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{VCS_GCONETMASK}	Scalar	Optional	This variable defines the Netmask of the virtual IP address to be used by the Global Cluster Option).
\$CPI:: CFG{VCS_GCONIC}{<SYSTEM>}	Scalar	Optional	This variable defines the NIC for the Virtual IP used for the Global Cluster Option. 'ALL' can be entered as a system value if the same NIC is used on all systems.
\$CPI::CFG{VCS_GCOVIP}	Scalar	Optional	This variable defines the virtual IP address to be used by the Global Cluster Option.
\$CPI:: CFG{VCS_LLTLINK#}{<SYSTEM>}	Variable	Required	This variable defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured.
\$CPI:: CFG{VCS_LLTLINKLOWPRI}{<SYSTEM>}	Variable	Optional	This variable defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication.
\$CPI::CFG{VCS_SMTPRECP}	Variable	Optional	This variable defines a list of full email addresses example: user@yourcompany.com) of SMTP recipients.
\$CPI::CFG{VCS_SMTPRSEV}	Variable	Optional	This variable defines the minimum severity level of messages (Information, Warning , Error, SevereError) that listed SMTP recipients are to receive.

Table C-1 Response file variables (*continued*)

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{VCS_SMTPSERVER}	Variable	Optional	This variable defines the domain-based hostname (example: smtp.yourcompany.com) of the SMTP server to be used for web notification.
\$CPI::CFG{VCS_SNMPCONS}	Variable	Optional	This variable defines a list of SNMP console system names.
\$CPI::CFG{VCS_SNMPSEV}	Variable	Optional	This variable defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive.
\$CPI::CFG{VCS_SNMPPORT}	Variable	Optional	This variable defines the SNMP trap daemon port (default=162).
\$CPI::CFG{VCS_USERENPW}	Variable	Optional	This variable defines an encrypted password for each VCS user.
\$CPI::CFG{VCS_USERNAME}	Variable	Optional	This variable defines a list of VCS usernames.
\$CPI::CFG{VCS_USERPRIV}	Variable	Optional	This variable defines each user's VCS privileges.
\$CPI::CFG{VM_DEFAULTDG}<SYSTEM>	Variable	Required	This variable defines the name of the default disk group.
\$CPI::CFG{VM_NO_OPEN_VOLS}	Variable	Required	This variable indicates that the user should not be asked if there are any open volumes when vxconfigd is not enabled. Such prompts are asked during uninstallations. (1: affirms there are no open volumes on the system)
\$CPI::CFG{VOL_MAX_NMPPOOL_SZ}	Variable	Optional	This variable defines the vol_max_nmpool_sz tunable.

Table C-1 Response file variables (*continued*)

Variable	List or Scalar	Optional or Required	Description
\$CPI::CFG{VOL_MAX_RDBACK_SZ}	Variable	Optional	This variable defines the vol_max_rdback_sz tunable.
\$CPI::CFG{VOL_MIN_LOWMEM_SZ}	Variable	Optional	This variable defines the vol_min_lowmem_sz tunable
\$CPI::CFG{VOL_RVIO_MAXPOOL_SZ}	Variable	Optional	This variable defines the vol_rvio_maxpool_sz tunable.
\$CPI::CFG{VRAS_STATS_FREQUENCY}	Variable	Optional	This variable defines the frequency for VVR stats collection.
\$CPI::CFG{VVR_DATA_PORTS}	Variable	Optional	This variable defines the port used by VVR for replication.
\$CPI::CFG{VVR_HB_PORT}	Variable	Optional	This variable defines the port used for VVR heartbeating.
\$CPI:: CFG{VVR_VRADMIND_PORT}	Variable	Optional	This variable defines the port used for vradmind heartbeating.
\$CPI:: CFG{VVR_VXRSYNCD_PORT}	Variable	Optional	This variable defines the port used by vxrsyncd daemon.

Automatic Storage Management

This appendix includes the following topics:

- [About ASM](#)
- [Veritas CVM and ASM integration](#)
- [Supported SF Oracle RAC and ASM configurations](#)
- [CVM and ASM configuration procedure](#)
- [ASM instance startup and shutdown commands](#)
- [Adding additional storage to an ASM disk group with CVM](#)
- [Veritas ASM agents](#)
- [Veritas CVM and Oracle service group with ASM](#)
- [Veritas CVM and ASM main.cf file](#)

About ASM

ASM provides an integrated cluster file system and volume management capabilities for Oracle database files. ASM can be configured with the Veritas Volume Manager (VM) or Cluster Volume Manager (CVM) to take advantage of the volume management and DMP multipathing features in the Veritas Storage Foundation for Oracle RAC.

Veritas CVM and ASM integration

ASM can be configured with Veritas Cluster Volume Manager (CVM) to take advantage of the volume management and DMP multipathing features in Veritas Storage Foundation for Oracle RAC. CVM provides uniform device names across all the nodes in the cluster. When a new array is added to a configuration, the upper level logical volumes remain the same to the user.

Dynamic MultiPathing

A key element of SF Oracle RAC is Dynamic MultiPathing (DMP). DMP provides multiple parallel I/O paths between servers and storage which improve both the availability and scalability of servers and storage.

Customers who wish to run ASM with Veritas Storage Foundation can benefit significantly from the DMP capabilities built into SF Oracle RAC. DMP provides maximum availability and performance in large heterogeneous SANs. As such, DMP allows users to confidently standardize on SF Oracle RAC to take full advantage of the agility that results from no longer being locked into a specific storage hardware vendor.

ASM on CVM rules

The following rules must be obeyed when running ASM on CVM:

- Do not create any file systems on VxVM volumes which will be used with ASM. Doing so may cause data corruption and/or result in headers getting overwritten.
- Do not include volumes with file systems or raw volumes in the same diskgroup being assigned to ASM. Veritas raw volumes for ASM should be used exclusively by ASM.
- Do not enable ODM when implementing ASM. ODM is only appropriate for data files residing on the Veritas File system (VxFS and/or CFS). If the user chooses to use Veritas Cluster File System for any other purpose, such as archive logs, ODM is not applicable.
- Do not stripe or concatenate the volumes using Veritas Volume Manager. ASM implements an additional layer of striping above hardware RAID configuration.
- It is recommended that the user create 2 Oracle ASM diskgroups (DGs).
 - Diskgroup1: data files, one set of redo logs, and one set of control files
 - Diskgroup2: Flash Recovery Area, archive logs, and a second set of redo logs and control files

Note: Please refer to Oracle's ASM best practices paper for more details.

- A VxVM volume should consist of a single disk/LUN as presented to the operating system.
- Do not change ownership or permissions on the DMP meta nodes, as this could prevent access to storage devices under DMP. For example:

```
/dev/vx/dmp/devicename
```

and

```
/dev/vx/rdmp/devicename
```


Note: A DMP metanode represents a metadvice, a Veritas Volume Manager abstraction that corresponds to a disk or LUN and all the I/O paths on which it can be accessed.

- Do not attempt to have ASM directly address the DMP meta nodes. All access to DMP meta nodes should be through the appropriate Veritas Volume Manager (clustered or unclustered).
- DMP combined with other multipathing software on the same system is neither allowed nor supported.
- Coordinator disks for I/O fencing cannot be configured under ASM.
- Do not attempt VxVM/CVM virtualization operations, i.e. volume resize, concatenation with LUNs that are under ASM control.
- Volumes presented to a particular ASM diskgroup should be of the same geometry, speed, and type.
- Veritas Storage Foundation commands `vxresize` or `vxassist` are not supported in this configuration.

Supported SF Oracle RAC and ASM configurations

[Figure D-1](#) below shows ASM and Veritas Storage Foundation working together in a clustered (Oracle RAC) configuration. In this figure, clustered versions of the Veritas Volume Manager (VxCVM) and Veritas File System (VxCFS) are shown in support of Oracle RAC.

Figure D-1 Supported Oracle RAC configuration



In this Oracle RAC configuration, binaries, clusterware files, and non-Oracle files can be stored in a Veritas File System (VxCFS). Database files are stored and managed by ASM, which can then leverage Storage Foundation Dynamic Multi-pathing (DMP) through Veritas Volume Manager (VxCVM).

Note: For any given host or RAC cluster only one multipathing solution should be active and managing the underlying devices.

The Oracle RAC Voting Disk and the Oracle Cluster Registry (OCR) (shown at the bottom of the figure) are not managed by ASM in Oracle, but can be stored in Veritas Volume Manager and/or File System (VxCVM or VxCFS).

Requirements for the supported SF Oracle RAC and ASM configurations

Symantec supports ASM configurations running with the Storage Foundation under the following conditions:

- Support is provided for ASM on any Storage Foundation products that are supported with Oracle 10.2.0.3 and later. Check with Symantec customer support for up to date listings of supported configurations.
- ASM support is provided for configurations that are running over raw logical volumes (VxVM and CVM).

ASM monitoring utilities

ASM offers monitoring utilities, as well as alert log files for troubleshooting. The ASM command line utility (ASMCMD) is useful to administer disk groups and monitor disk space usage.

For any ASM errors, check the corresponding log files for the Oracle instance.

Refer to your Oracle documentation for information about the ASM monitoring utilities and log files.


CVM and ASM configuration procedure

The CVM and ASM configuration procedure consists of the following tasks:

- Configure the CVM disk group and volume for the ASM disk group
- Install ASM in an ASM home directory
- Use the DBCA or manual procedure to create an ASM DG and instance
- Verify ASM
- Relink the SF Oracle RAC libraries to Oracle
- Unlink ODM

[Figure D-2](#) displays a flowchart of the above required tasks.

Figure D-2 CVM and ASM configuration procedure



Warning: Before beginning this process, please review the requirements in See [“ASM on CVM rules”](#) on page 594.

Before beginning this process, please review the SF Oracle RAC main.cf file for an Oracle with ASM configuration:

See [“Veritas CVM and ASM main.cf file”](#) on page 609.

Configure the CVM disk group and volume for the ASM disk group

As root, configure a CVM disk group using the following commands. The example below illustrates how to create the CVM disk group. This procedure needs to be performed from one node in the case of RAC, but that node must be the CVM master node.

To determine which node is the CVM master node, enter the following command:

```
# vxdctl -c mode
```

Note: This command will only display the CVM master node, if the VxVM cluster feature is enabled.

Configure the CVM disk group and volume for the ASM disk group

- 1 Enter the following commands for the first disk:

```
galaxy@root# vxdiskunsetup -C HDS0_31
```

```
galaxy@root# vxdisksetup -i HDS0_31 format=cddisk
```

- 2 Enter the following commands to initialize and configure the disk group:

```
galaxy@root# vxdg -s init ora_asm_dg HDS0_31
```

```
galaxy@root# vxassist -g ora_asm_dg make ora_asm_vol 2000M HDS0_31
```

```
galaxy@root# vxedit -g ora_asm_dg \  
set group=dba user=oracle mode=660 ora_asm_vol
```

As Oracle user configure ASM using either DBCA or EM. If you do not use either GUI, you will need to manually create the ASM instances and disk groups.

Install Automatic Storage Management in an ASM home directory

Installing ASM in its own home directory enables you to keep the ASM home separate from the database home directory (ORACLE_HOME). With ASM, Oracle recommends that Oracle binaries for the database should be installed in one ORACLE_HOME, and ASM binaries should be installed in a separate, different ORACLE_HOME. By using separate home directories, you can upgrade and patch ASM and the Oracle Database software independently, and you can deinstall Oracle Database software without affecting the ASM instance.

During the installation process, you are asked to configure ASM. You can configure ASM by creating disk groups that become the default location for files created in the database.

When you create a disk group, indicate whether the disk group is a normal redundancy disk group (2-way mirroring for most files by default) or a high-redundancy disk group (3-way mirroring), or an external redundancy disk group (no mirroring by ASM).

The default disk group type is normal redundancy.

To install ASM in an ASM home directory

- 1 Insert the Oracle CD.
- 2 Set the DISPLAY variable, mount the CD drive, and run the Oracle installer as Oracle User.

```
# /mnt/cdrom/runInstaller
```

- 3 Read the Welcome screen and click **Next**.

When you start the Oracle Universal Installer (OUI), the Select a Product to Install window appears.

- 4 Select the Enterprise Edition for RAC, and then click **Next**.
- 5 In the Specify Home Details window, specify a name for the ASM home directory, for example, OraASM11g_home. Select a directory that is a subdirectory of your Oracle base directory, for example, /app/oracle/ASM_HOME.

Click Browse to change the directory in which ASM will be installed.

After you have specified the ASM home directory, click **Next**. The Specify Hardware Cluster Installation Mode window appears.

- 6 Click Select All to select all nodes for installation and then click **Next**.

If your Oracle Clusterware installation was successful, then the Specify Hardware Cluster Mode window lists the nodes that you identified for your cluster, such as galaxy and nebula.

After you click Next, the Product Specific Prerequisites Checks window appears.

Follow the OUI wizard instructions for the rest of this procedure.

- 7 Apply the appropriate Oracle RAC 10g or Oracle RAC 11g patch or patchset software.

For Oracle RAC 10g:

See [“Adding Oracle RAC 10g patchsets”](#) on page 244.

For Oracle RAC 11g:

See [“Adding Oracle RAC 11g patches or patchsets”](#) on page 329.

Use the Oracle DBCA procedure to create ASM disk group and instances

Before performing the Oracle DBCA procedure below, a CVM volume must be created.

See “Configure the CVM disk group and volume for the ASM disk group” on page 598.

Perform the following steps for the Oracle DBCA (Database Creation Assistant) procedure. This procedure can be performed on an Oracle system configuration using either spfiles or pfiles.

Oracle DBCA procedure

- 1 From the Oracle DBCA interface, select the "Configure ASM Instance" option.
- 2 Create a default ASM disk group (Eg: DATADG).
- 3 Choose the appropriate redundancy.
- 4 Change the Discovery String to “/dev/vx/rdisk/ora_asm_dg/ora_asm*”
This should list all the VxVM volumes created and available for ASM.
- 5 Select the disk.

At the completion of this step, the ASM instance should be started and the disk group should be mounted on all nodes of the cluster.

To ensure that the disk group is mounted after a node restart, update the ASM init.ora parameter, asm_diskstring, with the disk group name. If ASM is using a spfile, then this update is done automatically by DBCA.

Verify the disk group status using the following SQL command:

```
SQL>select name, state from v$asm_diskgroup;
```

When using DBCA to create a database, select ASM for storage.

After the ASM disk group is created, all the processes associated with ASM instance start and the init.ora is created for each of the RAC ASM instances. The default ASM instance name is +ASMn where n is the instance number depending on the number of RAC instances.

Verify your ASM installation

Verify that all the database services for ASM are up and running.

To verify that ASM is operational following the installation

- 1 Change directories to the bin directory in the CRS home directory:

```
cd /app/crshome/bin
```

- 2 Run the following command as the oracle user, where galaxy is the name of the node you want to check:

```
./srvctl status asm -n galaxy
```

```
ASM instance +ASM1 is running on node galaxy
```

The example output shows that there is one ASM instance running on the local node.

- 3 Repeat the command shown in the step above, substituting nebula for galaxy to verify the successful installation on the other node in your cluster.

Relinking the SF Oracle RAC libraries to Oracle

Relink the Veritas library on \$ORACLE_HOME.

For information about this procedure for Oracle RAC 10g, refer to:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 10g”](#) on page 249.

Note: SF Oracle RAC only supports Oracle 10.2.0.3 and later on ASM.

For information about this procedure for Oracle RAC 11g refer to:

See [“Relinking the SF Oracle RAC libraries to Oracle RAC 11g”](#) on page 330.

Unlinking the SF Oracle RAC ODM library

ODM is only appropriate for data files residing on the Veritas File system. For this reason, do not enable ODM when implementing ASM.

Note: Symantec recommends that in an ASM environment, that the ODM library also be unlinked.

Unlinking the SF Oracle RAC ODM library

1 Log in as the user oracle.

2 Shutdown the instance if it is up:

```
$ srvctl stop asm -n galaxy
```

```
$ srvctl stop asm -n nebula
```

3 From another window, login to each node as root, and issue the following command:

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

4 On each node, enter the following command as oracle user:

```
$ cd $ORACLE_HOME/lib
```

5 On each node, enter the following command as oracle user.

For Oracle RAC 11g:

```
$ mv $ORACLE_HOME/lib/libodm11.so libodm11.so.veritas
```

For Oracle RAC 10g:

```
$ mv $ORACLE_HOME/lib/libodm10.so libodm10.so.veritas
```

6 On each node, enter the following command as oracle user.

For Oracle RAC 11g:

```
$ ln -s $ORACLE_HOME/lib/libodmd11.so libodm11.so
```

For Oracle RAC 10g:

```
$ ln -s $ORACLE_HOME/lib/libodmd10.so libodm10.so
```

- 7 On each node, enter the following commands as oracle user.

For Oracle RAC 11g:

```
$ cd /app/oracle/ASM_HOME/lib
$ rm libodm11.so
$ cp libodmd11.so libodm11.so
```

For Oracle RAC 10g:

```
$ cd /app/oracle/ASM_HOME/lib
$ rm libodm10.so
$ cp libodmd10.so libodm10.so
```

- 8 Start up the instance:

```
$ srvctl start asm -n galaxy
$ srvctl start asm -n nebula
```

ASM instance startup and shutdown commands

As user oracle use the following commands for normal startup and shutdown of the ASM instance. These commands can be run from any node of the cluster.

For ASM instance startup, type the following commands:

```
galaxy@oracle> srvctl start asm -n galaxy
galaxy@oracle> srvctl start asm -n nebula
```

For ASM instance shutdown, type the following command:

```
galaxy@oracle> srvctl stop asm -n galaxy
```

Note: The ASM instance cannot be shut down when database instances which are using ASM are up and running. Use CVM to shut down all database instances that are accessing the ASM instance before shutting down the ASM instance.

Adding additional storage to an ASM disk group with CVM

This section discusses adding additional storage to an ASM disk group with CVM.

Adding additional storage to an ASM disk group with CVM

- 1 As root user, to add additional storage use the following commands.

```
galaxy@root# vxdiskunsetup -C rhdisk30

galaxy@root# vxdisksetup -i rhdisk30

galaxy@root# vxdg -g ora_asm_dg adddisk rhdisk30

galaxy@root# vxassist -g ora_asm_dg make ora_asm_vol2 2000M
rhdisk30

galaxy@root# vxedit -g ora_asm_dg set group=dba user=oracle \
mode=660 ora_asm_vol2
```

- 2 Perform one of the following tasks to either add a disk to an existing ASM disk group or to extend the existing ASM disk size:

- To add a disk to an existing ASM disk group using SQL, enter the following SQL command:

```
SQL>alter diskgroup asm_diskgroup
add disk ora_asm_vol2
```

- To extend the existing ASM disk size using the root user, enter the following SQL command:

```
SQL>ALTER DISKGROUP asm_diskgroup
RESIZE DISK ora_asm_vol SIZE 4000M
```

Veritas ASM agents

There are two Veritas ASM agents required for an Oracle ASM configuration with the Veritas Cluster Volume Manager:

- ASMinst agent or ASM Instance agent
- ASMDG agent or ASM DG (disk group) agent

ASMinst agent will online, offline, and monitor the ASM instance.

ASMDG agent will mount, dismount, monitor the ASM disk groups required for ASM enabled databases.

Note: For additional information about these Veritas ASM agents, refer to the Veritas Cluster Server Agent for Oracle Installation and Configuration Guide.

Veritas CVM and Oracle service group with ASM

A VCS service group is a collection of resources working together to provide application services to clients. A VCS service group typically includes multiple resources that are both hardware and software based.

For additional details on service group dependencies, refer to the Veritas Cluster Server User's Guide.

Veritas CVM and an Oracle service group with ASM

In a configuration with the VCS Oracle agent, VCS controls the Oracle database. An online local firm dependency exists between the Oracle group and the CVM group.

The ASMInst agent in the asm_grp service group will online, offline, and monitor the ASM instance. The ASMDG agent in the database service group will mount, dismount, and monitor the ASM disk groups required for ASM enabled databases.

During VCS startup, the ASMDG agent will mount the ASM disk group required for the database, after the CVM volume on which the database resides is started. In case of multiple databases running on one node, it will protect them from each other's failures.

[Figure D-3](#) below displays the dependencies between the parallel Oracle RAC 10g service group with ASM and the parallel CVM service group.

Figure D-3

Veritas CVM and Oracle RAC 10g service group with ASM





Figure D-4 below displays the dependencies between the parallel Oracle RAC 11g service group with ASM and the parallel CVM service group.

Figure D-4 Veritas CVM and Oracle RAC 11g service group with ASM



Veritas CVM and ASM main.cf file

The following is an SF Oracle RAC sample configuration for Oracle RAC 10g with ASM.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "OracleASMTypes.cf"

cluster sfrac_asm (
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

system galaxy (
)

system nebula (
)

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

Application cssd (
    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"
)

CFSfsckd vxfsckd (
)
```

```
CVMCluster cvm_clus (  
    CVMClustName = sfrac_asm  
    CVMNodeId = { galaxy = 0, nebula = 1 }  
    CVMTransport = gab  
    CVMTimeout = 200  
)  
  
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  
)  
  
CVMVxconfigd cvm_vxconfigd (  
    Critical = 0  
    CVMVxconfigdArgs = { syslog }  
)  
  
PrivNIC ora_priv (  
    Device@galaxy = { en1 = 0, en2 = 1 }  
    Device@nebula = { en1 = 0, en2 = 1 }  
    Address@galaxy = "192.168.12.1"  
    Address@nebula = "192.168.12.2"  
    NetMask = "255.255.255.0"  
)  
  
cssd requires ocrvote_mnt  
cssd requires oradata_mnt  
cssd requires ora_priv  
ocrvote_mnt requires ocrvote_voldg
```

```
ocrvote_voldg requires cvm_clus
ocrvote_mnt requires vxfsckd
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

group asmgrp (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

ASMInst asm (
    Sid @galaxy = "+ASM1"
    Sid @nebula = "+ASM2"
    Owner = oracle
    Home = "/app/oracle/ASM_HOME"
)

requires group cvm online local firm

group dbgrp (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

ASMDG asmdg (
    Sid @galaxy = "+ASM1"
    Sid @nebula = "+ASM2"
    Owner = oracle
    Home = "/app/oracle/ASM_HOME"
    DiskGroups = { ASM_RAC_DG }
)

CVMVolDg racvoldg (
    CVMDiskGroup = rac_dg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

Oracle asmdb (
    ResourceOwner = oracle
```

```

        Sid @galaxy = vrts1
        Sid @nebula = vrts2
        Owner = oracle
        Home = "/app/oracle/ASM_HOME"
        StartUpOpt = SRVCTLSTART
        ShutDownOpt = SRVCTLSTOP
    )

```

```

requires group asmgrp online local firm
asmdb requires asmdg
asmdg requires racvoldg

```

The following is an SF Oracle RAC sample configuration for Oracle RAC 11g with ASM.

```

include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "OracleASMTypes.cf"

```

```

cluster sfrac_asm (
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

```

```

system galaxy (
)

```

```

system nebula (
)

```

```

group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

```

```
Application cssd (
    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"
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = sfrac_asm
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

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
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

MultiPrivNIC multi_priv (
```

```

Critical = 0
Device @galaxy = { en1 = 0, en2 = 1 }
Device @nebula = { en1 = 0, en2 = 1 }
Address @galaxy = { "192.168.1.1" =0, "192.168.1.2" =0, "192.168.2.1" =1 }
Address @nebula = { "192.168.1.3" =0, "192.168.1.4" =0, "198.168.2.2" =1 }
NetMask = "255.255.255.0"
)

```

```

cssd requires ocrvote_mnt
cssd requires oradata_mnt
cssd requires multi_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_voldg requires cvm_clus
ocrvote_mnt requires vxfsckd
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

```

```

group asmgrp (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

```

```

ASMInst asm (
    Sid @galaxy = "+ASM1"
    Sid @nebula = "+ASM2"
    Owner = oracle
    Home = "/app/oracle/ASM_HOME"
)

```

```

requires group cvm online local firm

```

```

group dbgrp (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

```

```

ASMDG asmdg (
    Sid @galaxy = "+ASM1"

```

```
Sid @nebula = "+ASM2"
Owner = oracle
Home = "/app/oracle/ASM_HOME"
DiskGroups = { ASM_RAC_DG }
)

CVMVolDg racvoldg (
    CVMDiskGroup = rac_dg
    CVMVolume = { oradataVOL }
    CVMActivation = sw
)

Oracle asmdb (
    ResourceOwner = oracle
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Owner = oracle
    Home = "/app/oracle/ASM_HOME"
    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)

requires group asmgrp online local firm
asmdb requires asmdg
asmdg requires racvoldg
```


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.

Note: This appendix only discusses the procedures for creating a test database for Oracle RAC 10g and Oracle RAC 11g. For information about creating a test database for Oracle 9i, refer to earlier SF Oracle RAC documentation.

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 CRS, 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 in a shared raw volume or in a cluster file system (CFS). The following sections describe creating database tablespaces for Oracle on either raw volumes or on a file system:

- [Creating the database tablespaces for Oracle on raw volumes](#)
- [Creating the database tablespaces for Oracle on CFS](#)

Creating the database tablespaces for Oracle on raw volumes

You can create the database tablespaces on shared raw volume.

To create shared raw volume for database tablespaces

- 1 Log in as root user.
- 2 On the master node, create a shared disk group:

```
# vxvg -s init oradatadg HDS0_20
```

- 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 tablespaces for Oracle 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.

- 1 Create a disk group (for example, `oradatadg`):

```
# vxdg -s init oradatadg HDS0_20
```

- 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 are required for the tablespaces, type:

```
# vxassist -g oradatadg make oradatavol 6800M
```

Refer to your Oracle documentation for tablespace sizes.

- 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 -V 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 -V vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol1 \
/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 through step 7.
- 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](#)
- [PrivNIC agent](#)
- [MultiPrivNIC agent](#)
- [CSSD agent](#)
- [VCS agents for Oracle](#)

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 the specific functions that an agent performs.

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 network interface card, a software component such as Oracle RAC 10g 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 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

A SF Oracle RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

You must configure the CVMVolDg agent for each diskgroup that is used by an Oracle service group. Configure a diskgroup for only a single Oracle service group. If the database uses cluster file systems, configure the CFSSMount agent for each volume in the diskgroup.

Use the following information in this appendix about the entry points and attributes of the above listed agents to make necessary configuration changes. Refer to the *Veritas Cluster Server User's Guide* for information on how to modify the VCS configuration.

Other agents included within SF Oracle RAC

SF Oracle RAC includes the following additional agents:

- PrivNIC agent
The PrivNIC agent provides resources that are used to maintain private IP addresses locally that are highly available on LLT Ethernet interfaces.

- **MultiPrivNIC agent**
The MultiPrivNIC agent provides resources for UDP/IP support for Oracle's cache fusion capabilities.
- **CSSD agent**
The CSSD (Cluster Synchronization Services daemon) agent provides the resources to monitor the Oracle RAC 10g or Oracle RAC 11g cssd process, and to ensure that the dependency of cssd on the OCR and VOTE resources and the PrivNIC (optional) resource are satisfied.
- **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.
- **ASMinst agent**
The ASMinst agent manages and controls the ASM instance.
- **ASMDG agent**
The ASMDG agent mounts and unmounts the ASM disk groups onto an ASM instance.

Refer to the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for more information on the Oracle, Netlsnr, ASMinstance, and ASMDG agents.

CVMCluster agent

The CVMCluster agent controls system membership on the cluster port associated with Veritas Volume Manager (VxVM).

Entry points for CVMCluster agent

[Table F-1](#) describes the entry points used by the CVMCluster agent.

Table F-1 CVMCluster agent entry points

Entry Point	Description
Online	Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by autoimporting shared diskgroups.
Offline	Removes a node from the CVM cluster port.
Monitor	Monitors the node's CVM cluster membership state.

Attribute definition for CVMCluster agent

Table F-2 describes the user-modifiable attributes of the CVMCluster resource type.

Table F-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 cluster messaging mechanism. Default = gab Note: Do not change this value.
PortConfigd	integer-scalar	Port number used by CVM for vxconfigd-level communication.
PortKmsgd	integer-scalar	Port number 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 that the `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` attributes are not used in an SF Oracle RAC environment because 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 = RACCluster1
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

```

For a more extensive `main.cf` example that includes the CVMCluster resource:

See [“About sample main.cf files”](#) on page 507.

CVMVxconfigd agent

The CVMVxconfigd agent starts and monitors the vxconfigd daemon. The vxconfigd daemon maintains disk and diskgroup configurations, communicates configuration changes to the kernel, and modifies configuration information 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.

Entry points for CVMVxconfigd agent

[Table F-3](#) describes the entry points for the CVMVxconfigd agent.

Table F-3 Vxconfigd entry points

Entry Point	Description
Online	Starts the vxconfigd daemon
Offline	N/A
Monitor	Monitors whether vxconfigd daemon is running

Attribute definition for CVMVxconfigd agent

Table F-4 describes modifiable attributes of the CVMVxconfigd resource type.

Table F-4 CVMVxconfigd agent attribute

Attribute	Dimension	Description
CVMVxconfigdArgs	keylist	List of arguments sent to the online entry point. Symantec recommends always specifying the syslog option.

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 }
)
.
```

```

.
cvm_clus requires cvm_vxconfigd
// resource dependency tree
//
//     group cvm
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }

```

For a more extensive main.cf that includes the CVMVxconfigd resource:

See [“About sample main.cf files”](#) on page 507.

CVMVolDg agent

The CVMVolDg agent represents and controls CVM diskgroups and CVM volumes within the diskgroups. The global nature of CVM diskgroups and volumes requires importing them only once on the CVM master node.

Configure the CVMVolDg agent for each disk group used by an 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 CFSSMount agent for each volume in the disk group.

Entry points for CVMVolDg agent

[Table F-5](#) describes the entry points used by the CVMVolDg agent.

Table F-5 CVMVolDg agent entry points

Entry Point	Description
Online	Starts all volumes in the shared disk group specified by the CVMVolume attribute. Makes sure that the shared disk groups are imported. Note that the CVMVolDg agent does not import shared disk groups. If the disk group is not imported, the online script returns an error. 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.

Table F-5 CVMVolDg agent entry points *(continued)*

Entry Point	Description
Offline	Sets the diskgroup activation mode to off so that all the volumes in diskgroup are invalid.
Monitor	Monitors specified critical volumes in the diskgroup. The CVMVolume attribute specifies these volumes. SF Oracle RAC requires specifying at least one volume in a diskgroup.
Clean	Cleans up the temporary files created by the online entry point.

Attribute definition for CVMVolDg agent

Table F-6 describes the user-modifiable attributes of the CVMVolDg resource type.

Table F-6 CVMVolDg agent attributes

Attribute	Dimension	Description
CVMDiskGroup	string-scalar	Names the diskgroup.
CVMVolume	string-keylist	Lists critical volumes in the diskgroup. SF Oracle RAC requires specifying at least one volume in the diskgroup.
CVMActivation	string-scalar	Sets the activation mode for the diskgroup. Default = sw (shared-write)

CVMVolDg agent type definition

The CVMTypes.cf file includes the CVMVolDg type definition.

```
type CVMVolDg (
    static keylist RegList = { CVMActivation }
    static str ArgList[] = { CVMDiskGroup, CVMVolume,
        CVMActivation }
    str CVMDiskGroup
    keylist CVMVolume[]
    str CVMActivation
    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 (
    CVMDiskGroup = oradatadg
    CVMVolume = { oradata1, oradata2 }
    CVMActivation = sw
)
```

To see CVMVolDg defined in a more extensive example:

See [“About sample main.cf files”](#) on page 507.

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.

Entry points for CFSMount agent

[Table F-7](#) provides the entry points for the CFSMountpoint agent.

Table F-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.

Attribute definition for CFSMount agent

[Table F-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

Table F-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.
MountOpt (optional)	string-scalar	Options for the mount command. To create a valid MountOpt attribute string: ■ Use the VxFS type-specific options only. ■ Do not use the -o flag to specify the VxFS-specific options. ■ Do not use the -v vxfs file system type option. ■ Be aware the cluster option is not required. ■ Specify options in comma-separated list: ro ro,cluster blkclear,mincache=closesync
Policy (optional)	string-scalar	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.

CFSMount agent type definition

The CFSTypes.cf file includes the CFSMount agent type definition.

```
type CFSMount (
    static keylist RegList = { MountOpt, Policy, NodeList }
    static int FaultOnMonitorTimeouts = 1

    static int InfoTimeout = 0
    static int OnlineRetryLimit = 16

    static int OnlineWaitLimit = 0
    static str ArgList[] = { MountPoint, BlockDevice,
        MountOpt }
```

```

NameRule = resource.MountPoint
str MountPoint
str MountType
str BlockDevice
str MountOpt

keylist NodeList
keylist Policy

temp str Primary
str SetPrimary
str RemountRes
str ForceOff

```

CFSMount agent sample configuration

Each Oracle service group requires a CFSMount resource type to be defined.

```

CFSMount ora_mount (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol1"
    Primary = nebula
)

```

To see CFSMount defined in a more extensive example:

See [“About sample main.cf files”](#) on page 507.

PrivNIC agent

The PrivNIC resource can be used to maintain a private IP address that is locally highly available on LLT Ethernet interfaces. Private IP addresses are required by CRS daemons in Oracle RAC 10g to provide communication between the cluster nodes.

The PrivNIC agent relies on LLT to monitor the LLT Ethernet interfaces. The PrivNIC agent queries LLT to count the number of visible nodes on each of the LLT Ethernet interfaces.

Entry point for PrivNIC agent

[Table F-9](#) describes the PrivNIC agent's monitor entry point.

Note: Because the resource is persistent, only the monitor entry point is required.

Table F-9 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.

Attribute definition for PrivNIC agent

[Table F-10](#) describes the user-modifiable attributes of the PrivNIC resource type.

Table F-10 Required attributes for PrivNIC agent

Attribute	Dimension	Description
Device	string - association	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. At least one interface device must be specified. Example:

Table F-10 Required attributes for PrivNIC agent (*continued*)

Attribute	Dimension	Description
Address	string-scalar	The numerical private IP address. Checks are performed to determine if this is a valid IP address. IP address validation is performed on the numerical private IP address to determine if this is an IPV4 or IPV6 address. Additional checks are performed to determine if this is a valid IP address. The following is an example of an IPv4 address: <pre>Address = "192.168.12.1"</pre>
NetMask	string - association	The numerical netmask for the private IP address. For example: <pre>Address = "255.255.240.0"</pre>

Optional Attributes

Table F-11 Optional attributes for PrivNIC agent

Attribute	Dimension	Description
DeviceTag	string - association	Associates an LLT device "tag" with device via the network-id. 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. Example: <code>GabPort = "b"</code>

Table F-11 Optional attributes for PrivNIC agent (*continued*)

Attribute	Dimension	Description
UseVirtualIP	integer-scalar	The default is 0, which specifies that the agent use the physical interface for configuring the private IP address when possible. The value 1 specifies that the agent always use the virtual interface for configuring the private IP address. 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.

PrivNIC agent type definition

The following shows the content of 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[]
)
```

PrivNIC agent sample configuration

The following is a sample configuration using the PrivNIC agent (IPv4 address).

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1 } // XXX
    Device@nebula = { en1 = 0, en2 = 1 } // XXX
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.240.0"
)
```

For a more extensive example:

See [“About sample main.cf files”](#) on page 507.

```
Device = { en1 = 0, en2 = 1, en3 = 2 }
```

MultiPrivNIC agent

The MultiPrivNIC agent provides resources for UDP/IP support for Oracle RAC 11g cache fusion capabilities.

Note: Private IP address types are required by Oracle for cache fusion traffic. In addition to the MultiPrivNIC agent and UDP/IP support for Oracle RAC 11g cache fusion, Veritas Cluster Inter-Process Communication (VCSIPC) supports Oracle RAC 10g cache fusion and provides an API for Oracle RAC 10g cache fusion. The API is provided through the libskgxp library. The libskgxp library makes calls into the LLT Multiplexer (LMX). LMX uses Low Latency Transport LLT to send data across the private interconnects. LMX also provides link aggregation and redundancy of private links.

MultiPrivNIC agent entry point

Table F-12 below describes the MultiPrivNIC agent's monitor entry point.

Note: Because the resource is persistent, only the monitor entry point is required.

Table F-12 MultiPrivNIC agent entry point

Entry point	Description
Monitor	The MultiPrivNIC agent queries LLT to create a list of the visible nodes on each LLT network interface. The MultiPrivNIC agent then applies various filters to this list to determine a most desired failover decision and it 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" device are different, then the MultiPrivNIC agent initiates a failover from the currently active device to the "winner" device. All things being equal, the IP address is configured on the preferred device.

MultiPrivNIC agent attribute definition

Table F-13 below describes the user-modifiable attributes of the MultiPrivNIC resource type.

Table F-13 MultiPrivNIC agent attribute definitions

Attribute	Dimension	Description
Device	string-association	The device attribute specifies the network interface device displayed by the <code>ifconfig</code> command and the network-id associated with the interface. The 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. An example of the device attribute is as follows: On galaxy: en1 en2 On nebula: en1 en2

Table F-13 MultiPrivNIC agent attribute definitions (*continued*)

Attribute	Dimension	Description
Address	string-association	The numerical private IP address. Checks are performed to determine if this is a valid IP address. The following is an example of an IPv4 address: <pre>Address @galaxy = { "192.168.1.1" =0, "192.168.1.2" =0, "192.168.2.1" =1 }</pre>
Netmask	string-association	The netmask attribute is the numerical netmask for the private IP address. For example, Address = "255.255.255.0".

Resource type definition for MultiPrivNIC agent

The following is an example of the contents of 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}
  static str Operations = None
  str Device{}
  str DeviceTag{}
  str Address{}
  str NetMask
  int UseVirtualIP
  str GabPort = o
  int UseSystemList
  int ExcludeNode[]
)
```

Sample configuration for MultiPrivNIC agent

The following is a sample MultiPrivNIC agent configuration for IPv4:

```
MultiPrivNIC multi_priv (
  Critical = 0
```

```
Device @galaxy = { en1 = 0, en2 = 1 }
Device @nebula = { en1 = 0, en2 = 1 }
Address @galaxy = { "192.168.1.1" =0, "192.168.1.2" =0, "192.168.2.1" =1 }
Address @nebula = { "192.168.1.3" =0, "192.168.1.4" =0, "192.168.2.2" =1 }
NetMask = "255.255.255.0"
)
```

For a more extensive example:

See [“About sample main.cf files”](#) on page 507.

In the above example, the interface next to the IP address is the preferred device for that particular IP address. All links being equal (same number of visible nodes on each of the LLT interfaces), the IP addresses are configured on the preferred interfaces. The MultiPrivNIC agent relies on LLT to monitor the cluster interfaces. The MultiPrivNIC agent 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 next available private link, which has the maximum nodes. When the preferred links comes back up again, the IP address corresponding to this link reverts back to it.

Note: The user configures the priority of the devices during SF Oracle RAC installation and configuration.

CSSD agent

The `cssd` resource monitors the Oracle RAC 10g `cssd` process. The purpose of the `cssd` resource is to ensure that the dependency of `cssd` on the OCR and VOTE resources and the PrivNIC (optional) resource are satisfied. If the `cssd` resource is online and any of its dependencies are brought offline, the machine will reboot. This agent allows this behavior to be avoided since the dependencies will be enforced by VCS.

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

For Oracle RAC 10g R2, to disable CRS from starting at boot enter the following command:

```
# $CRS_HOME/bin/crsctl disable crs
```

For Oracle RAC 10g R2, to enable CRS enter the following command:

```
# $CRS_HOME/bin/crsctl enable crs
```

The `cssd` resource should use the Application agent. The user determines the name of the resource.

[Table F-14](#) lists the required attributes of the `cssd` resource:

Table F-14 Required attributes for CSSD resource

Attribute Name	Required Value
Critical	0
OnlineRetryLimit	20
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

An example `main.cf` entry is as follows:

```
Application cssd-resource (
    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
)
```

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 640.
See [“Oracle resource type”](#) on page 646.
- The Netlsnr agent monitors the listener process.

See [“Netlsnr agent functions”](#) on page 653.

See [“Netlsnr resource type”](#) on page 654.

- The ASMinst agent monitors the Oracle ASM instance.
See [“ASMinst agent functions”](#) on page 657.
See [“ASMinst resource type”](#) on page 658.
- The ASMDG agent monitors the Oracle ASM disk groups.
See [“ASMDG agent functions”](#) on page 660.
See [“ASMDG resource type”](#) on page 661.

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 F-15](#) lists the Oracle agent operations.

Table F-15 Oracle agent operations

Agent operation	Description
Online	Starts the Oracle database by using the following <code>svrmgrl</code> or <code>sqlplus</code> command: <pre>startup force pfile=\$PFile</pre> The default Startup option is <code>STARTUP_FORCE</code>. You can also configure the agent to start the database using different Startup options for Oracle. See “Startup and shutdown options” on page 641.
Offline	Stops the Oracle database with the specified options by using the following <code>svrmgrl</code> or <code>sqlplus</code> command: <pre>shutdown immediate</pre> The default Shutdown option is <code>IMMEDIATE</code>. You can also configure the agent to stop the database using different Shutdown options for Oracle. See “Startup and shutdown options” on page 641.
Monitor	Verifies the status of the Oracle processes. The Oracle agent provides two levels of monitoring: basic and detail. See “Monitor options for Oracle agent” on page 643.

Table F-15 Oracle agent operations (*continued*)

Agent operation	Description
Clean	Forcibly stops the Oracle database by using the following <code>svrmgrl</code> or <code>sqlplus</code> command: <pre>shutdown abort</pre> If the process does not respond to the <code>shutdown</code> command, then the agent does the following: ■ Scans the process table for the processes that are associated with the configured instance ■ Kills the processes that are associated with the configured instance
Info	Provides the static and dynamic information about the state of the database. See “Info entry point” on page 644.
Action	Performs the predefined actions on a resource. See “Action entry point” on page 645.

Startup and shutdown options

You can specify Startup and Shutdown options for the Oracle instances that are configured.

[Table F-16](#) lists the startup options that the agent supports.

Table F-16 Startup options

Option	Description
STARTUP_FORCE (Default)	Runs the command <code>startup force</code> <code>pfile='location_of_pfile'</code> if the pfile is configured. If the pfile is not configured, the agent runs <code>startup force</code>. It picks up the default parameter files from their default locations.
STARTUP	Runs the command <code>startup pfile='location_of_pfile'</code> if the pfile is configured. If the pfile is not configured, the agent picks up the default parameter files from their default locations and runs <code>startup</code>.
RESTRICTED	Starts the database in the RESTRICTED mode.
RECOVERDB	Performs a database recovery on instance startup.

Table F-16 Startup options (*continued*)

Option	Description
CUSTOM	Uses a predefined SQL script (<code>start_custom_\$\$SID.sql</code>) and runs custom startup options. The script must be in the <code>/opt/VRTSagents/ha/bin/Oracle</code> directory and must have access to the Oracle Owner OS user. If the file is not present, the agent logs an error message. With a custom script, the agent takes the following action: <pre>sqlplus /nolog <<! connect / as sysdba; @start_custom_\$\$SID.sql exit; !</pre>
SRVCTLSTART	Uses the <code>srvctl</code> utility to start an instance of the database. For RAC clusters, you must manually set the default startup option as <code>SRVCTLSTART</code>.

[Table F-17](#) lists the shutdown options that the agent supports.

Table F-17 Shutdown options

Option	Description
IMMEDIATE (Default)	Shuts down the Oracle instance by running <code>shutdown immediate</code> .
TRANSACTIONAL	Runs the <code>shutdown transactional</code> command. This option is valid only for the database versions that support this option.
CUSTOM	Uses a predefined SQL script (<code>shut_custom_\$\$SID.sql</code>) and runs custom shutdown options. The script must be in the <code>/opt/VRTSagents/ha/bin/Oracle</code> directory and must have access to the Oracle Owner OS user. If the file is not present, the agent shuts the agent down with the default option.
SRVCTLSTOP	Uses the <code>srvctl</code> utility to stop an instance of the database. For RAC clusters, you must manually set the default option as <code>SRVCTLSTOP</code>.

Monitor options for 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 the agent must perform basic monitoring in Process check or Health check mode.

[Table F-18](#) describes the basic monitoring options.

Table F-18 Basic monitoring options

Option	Description
0 (Default)	Process check The agent scans the process table for the ora_dbw, ora_smon, ora_pmon, and ora_lgwr processes to verify that Oracle is running.
1	Health check (supported on Oracle 10g and later) The agent uses the Health Check APIs from Oracle to monitor the SGA and retrieve the information about the instance. If you want to use the Oracle agent's intentional offline functionality, you must enable Health check monitoring.

Review the following considerations if you want to configure basic monitoring:

- Basic monitoring of Oracle processes is user-specific. As a result, an Oracle instance started under the context of another user cannot be detected as online. For example, if an Oracle instance is started under the user "oraVRT" and the agent is configured for a user "oracle", the agent will not detect the instance started by "oraVRT" as online.
This could lead to situations where issuing a command to online a resource on a node might online an already running instance on that node (or any other node).
So, Symantec recommends that instances started outside SF Oracle RAC control be configured with the correct Owner attribute corresponding to the OS user for that instance.
- Within a failover service group, when the administrator onlines 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 DetailMonitor attribute of the Oracle resource determines whether the the agent must perform detail monitoring.

Info entry point

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

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 F-20](#) on page 646.

[Table F-19](#) describes the agent's predefined actions.

Table F-19 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; <i>actionargs</i> contains name of the tablespace to be backed up.

[Table F-20](#) lists the virtual fire drill actions of the Veritas Cluster Server agent for Oracle lets you run infrastructure checks and fix specific errors.

Table F-20 Predefined virtual fire drill actions

Virtual fire drill action	Description
getid (Oracle agent)	Verifies that the Oracle Owner exists on the node.
home.vfd (Oracle agent)	Verifies the following: ■ ORACLE_HOME is mounted on the node and corresponding entry is in the fstab. If the ORACLE_HOME is not mounted, the action entry point checks if any other resource has already mounted ORACLE_HOME. ■ Pfile is provided and it exists on the node. ■ Password file from \$ORACLE_HOME/dbs/orapw[SID] is present.
owner.vfd (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 (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 (Netlsnr agent)	Checks if listener.ora file is present. If listener.ora file, it checks if ORACLE_HOME is mounted and displays appropriate messages.

Oracle resource type

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, EnvFile, AutoEndBkup, DetailMonitor,
```

```

        User, Pword, Table, MonScript, AgentDebug, Encoding,
        MonitorOption }

    str Sid
    str Owner
    str Home
    str Pfile
    str StartUpOpt = STARTUP_FORCE
    str ShutDownOpt = IMMEDIATE
    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 int IntentionalOffline = 1

)

```

Oracle attribute definitions

Review the description of the Oracle agent attributes. The agent attributes are classified as required, optional, and internal.

[Table F-21](#) lists the required attributes. You must assign values to the required attributes.

Table F-21 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, as the defined owner of executables and database files in /etc/passwd.

Table F-21 Required attributes for Oracle agent (*continued*)

Required attributes	Type and dimension	Definition
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path.

[Table F-22](#) lists the optional attributes for Oracle agent. You can configure the optional attributes if necessary.

Table F-22 Optional attributes for Oracle agent

Optional Attributes	Type and Dimension	Definition
StartUpOpt	string-scalar	Startup options for the Oracle instance. This attribute can take the following values: ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ RECOVERDB ■ SRVCTLSTART ■ CUSTOM Default is STARTUP_FORCE. See “Startup and shutdown options” on page 641.
ShutDownOpt	string-scalar	Shut down options for the Oracle instance. This attribute can take the following values: ■ IMMEDIATE ■ TRANSACTIONAL ■ SRVCTLSTOP ■ CUSTOM Default is IMMEDIATE. See “Startup and shutdown options” on page 641.

Table F-22 Optional attributes for Oracle 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. 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.
Pfile	string-scalar	The name of the initialization parameter file with the complete path of the startup profile. 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.
AutoEndBkup	integer-scalar	Setting the AutoEndBkup attribute to a non-zero value takes the datafiles in the database out of the backup mode, during Online. Default = 1
MonitorOption	integer-scalar	Monitor options for the Oracle instance. This attribute can take values 0 or 1. ■ 0 - Process check monitoring (recommended) ■ 1 - Health check monitoring Default = 0 See “Monitor options for Oracle agent” on page 643.

Table F-22 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
DetailMonitor	integer-scalar	Setting this flag to a non-zero enables detail monitoring for Oracle. 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 monitor intervals. Note: If you set the AutoEndBkup attribute value to 0, then you must set the DetailMonitor attribute value to 1. Default = 0
MonScript	string-scalar	Pathname to the script provided for detail monitoring. The default (basic monitoring) is to monitor the database PIDs only. Note: Detail monitoring is disabled if the value of the attribute MonScript is invalid or is set to an empty string. The pathname to the supplied detail monitor script is /opt/VRTSagents/ha/bin/Oracle/SqlTest.pl. MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Oracle/SqlTest.pl.
User	string-scalar	Internal database user. Connects to the database for detail monitoring.

Table F-22 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
Pword	string-scalar	Encrypted password for internal database-user authentication. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.
Table	string-scalar	Table for update by <code>User/Pword</code> .
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".

Table F-22 Optional attributes for Oracle agent *(continued)*

Optional Attributes	Type and Dimension	Definition
IntentionalOffline	static-int	Defines how VCS reacts when Oracle is intentionally stopped outside of VCS control. If you stop Oracle out of VCS control, the agent behavior is as follows: ■ 0—the Oracle agent registers a fault and initiates the failover of the service group. ■ 1—the Oracle agent takes the Oracle resource offline when Health check monitoring is enabled. If Health check monitoring is not enabled, the agent registers a fault and initiates the failover of the service group. Note: 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. Default = 1 See <i>Veritas Cluster Server User's Guide</i>.
AgentDebug	boolean-scalar	Additional debug messages are logged when this flag is set. Default = 0

Table F-23 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 F-23 Internal attributes for Oracle agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	Specifies the location of binaries, scripts, and other files related to the Oracle agent. Default is /opt/VRTSagents/ha/bin/Oracle.

Netlsnr agent functions

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 F-24](#) lists the Netlsnr agent operations.

Table F-24 Netlsnr agent operations

Agent operation	Description
Online	Starts the listener process by using the following command: <code>lsnrctl start \$LISTENER</code>
Offline	Stops the listener process by using the following command: <code>lsnrctl stop \$LISTENER</code> 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. The Netlsnr agent provides two levels of monitoring, basic and detail: ■ 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. ■ In the detail monitoring mode, the agent uses the <code>lsnrctl status \$LISTENER</code> command to verify the status of the Listener process. (Default)
Clean	Scans the process table for <code>tnslsnr \$Listener</code> and kills it.
Action	Performs the predefined actions on a resource. See “Action entry point” on page 645.

Netlsnr resource type

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 }  
    str Owner  
    str Home  
    str TnsAdmin  
    str Listener  
    str EnvFile  
    str MonScript = "./bin/Netlsnr/LsnrTest.pl"  
    str LsnrPwd  
    boolean AgentDebug = 0  
    str Encoding  
    static int IntentionalOffline = 0  
  
)
```

Netlsnr attribute definitions

Review the description of the Netlsnr agent attributes. The agent attributes are classified as required, optional, and internal.

[Table F-25](#) lists the required attributes for Netlsnr agent. You must assign values to the required attributes.

Table F-25 Required attributes for Netlsnr agent

Required attributes	Type and dimension	Definition
Owner	string-scalar	Oracle user, as the defined owner of executables and database files in /etc/passwd.

Table F-25 Required attributes for Netlsnr agent (*continued*)

Required attributes	Type and dimension	Definition
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Do not append a slash (/) at the end of the path.

[Table F-26](#) lists the optional attributes for Netlsnr agent. You can configure the optional attributes if necessary.

Table F-26 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). 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. 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. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.

Table F-26 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
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. 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.
MonScript	string-scalar	Pathname to the script provided for detail monitoring. By default, the detail monitoring is enabled to monitor the listener process. Note: If the value of the attribute MonScript is set to an empty string, the agent disables detail monitoring. The pathname to the supplied detail monitoring script is /opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl. MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Netlsnr/LsnrTest.pl.
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".

Table F-26 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
IntentionalOffline	static-int	For future use. Do not change the value of this attribute. Default = 0
AgentDebug	boolean	Additional debug messages are logged when this flag is set. Default = 0

[Table F-27](#) 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 F-27 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. Default is /opt/VRTSagents/ha/bin/Netlsnr.

ASMIInst agent functions

The ASMIInst agent monitors the processes of ASM instance.

[Table F-28](#) lists the ASMIInst agent operations.

Table F-28 ASMIInst agent operations

Agent operation	Description
Online	Starts the Oracle ASM instance by using the following SQL command: <code>startup nomount</code> The Online operation starts the ASM instance without mounting any disk groups.
Offline	Stops the Oracle ASM instance by using the following SQL command: <code>shutdown immediate</code>

Table F-28 ASMInst agent operations (continued)

Agent operation	Description
Monitor	Verifies the status of the ASM instance. The ASMInst agent monitors the ASM instance using the Health check monitoring method. If the Health check monitoring fails, the agent does Process check monitoring. The agent also checks if the ocssd.bin process is running. The agent returns offline for the following conditions: ■ The process is not running. ■ The process is restarted. Note: Make sure that the OCSSD process is running. The ASMInst agent only monitors the OCSSD process. The agent does not start or stop the process.
Clean	Forcibly stops the Oracle ASM instance by using the following SQL command: <pre>shutdown abort</pre> If the process does not respond to the <code>shutdown</code> command, the agent kills the process using the <code>SIGTERM</code> or <code>SIGKILL</code> commands.

ASMInst resource type

The ASMInst agent of the Veritas Cluster Server agent for Oracle is represented by the ASMInst resource type in SF Oracle RAC.

```
type ASMInst (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/ASMInst"
    static str ArgList[] = { Sid, Owner, Home, Pfile,
    EnvFile, Encoding }
    str Sid
    str Owner
    str Home
    str Pfile
    str EnvFile
    str Encoding
)
```

ASMInst attribute definitions

Review the description of the ASMInst agent attributes. The agent attributes are classified as required, optional, and internal.

[Table F-29](#) lists the required attributes. You must assign values to the required attributes.

Table F-29 Required attributes for ASMinst agent

Required attributes	Type and dimension	Definition
Sid	string-scalar	The variable \$ORACLE_SID that represents the ASM instance. The Sid is considered case-sensitive by the ASMinst agent.
Owner	string-scalar	The Oracle user, as the defined owner of \$ORACLE_HOME of ASM instance and in /etc/passwd.
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. Note: Do not append a slash (/) at the end of the path.

[Table F-30](#) lists the optional attributes for ASMinst agent. You can configure the optional attributes if necessary.

Table F-30 Optional attributes for ASMinst agent

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. 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 F-30 Optional attributes for ASMInst agent (*continued*)

Optional Attributes	Type and Dimension	Definition
Pfile	string-scalar	The name of the initialization parameter file of ASM instance with the complete path of the startup profile. 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.
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".

[Table F-31](#) lists the internal attribute for ASMInst agent. This attribute is for internal use only. Symantec recommends not to modify the value of this attribute.

Table F-31 Internal attributes for ASMInst agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	Specifies the location of binaries, scripts, and other files related to the ASMInst agent. Default is /opt/VRTSagents/ha/bin/ASMInst.

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 DiskGroup attribute of the ASMDG agent.

[Table F-32](#) lists the ASMDG agent operations.

Table F-32 ASMIInst agent operations

Agent operation	Description
Online	Mounts the specified Oracle ASM disk groups to an ASM instance by using the following SQL command: <pre>alter diskgroup dg_name1, dg_name2 mount</pre>
Offline	Unmounts the specified Oracle ASM disk groups from an ASM instance by using the following SQL command: <pre>alter diskgroup dg_name1, dg_name2 dismount</pre> Note: The following Oracle message appears in the VCS log when an ASM instance with no ASM disk groups mounted is shut down: ORA-15100: invalid or missing diskgroup name
Monitor	Verifies the status of the specified ASM disk groups. The disk groups can be in one of the following states: ■ mounted ■ dismounted ■ unknown ■ broken ■ connected 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.
Clean	Forcibly unmounts the Oracle ASM disk groups by using the following SQL command: <pre>alter diskgroup dg_name1, dg_name2 dismount force</pre>

ASMDG resource type

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, DiskGroups,
        EnvFile, Encoding }
    str Sid
    str Owner
    str Home
```

```
keylist DiskGroups
str EnvFile
str Encoding
)
```

ASMDG attribute definitions

Review the description of the ASMDG agent attributes. The agent attributes are classified as required, optional, and internal.

Table F-33 lists the required attributes. You must assign values to the required attributes.

Table F-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. The Sid is considered case-sensitive by the ASMinst agent.
Owner	string-scalar	The Oracle user, as the defined owner of \$ORACLE_HOME of ASM instance and in /etc/passwd.
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. Note: Do not append a slash (/) at the end of the path.

Table F-34 lists the optional attributes for ASMDG agent. You can configure the optional attributes if necessary.

Table F-34 Optional attributes for ASMDG agent

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. 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 F-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 F-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. Default is /opt/VRTSagents/ha/bin/ASMDG.

Glossary

Agent	A process that starts, stops, and monitors all configured resources of a type, and reports their status to VCS.
Authentication Broker	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.
Cluster	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.
CVM (cluster volume manager)	The cluster functionality of Veritas Volume Manager.
Disaster Recovery	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.
disk array	A collection of disks logically arranged into an object. Arrays tend to provide benefits such as redundancy or improved performance.
DMP (Dynamic Multipathing)	A feature of 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.
DST (Dynamic Storage Tiering)	A feature with which administrators of multi-volume 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.

See also Veritas File System (VxFS)

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 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.

resources	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.
Resource Dependency	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.
Resource Types	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.
root broker	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.
SAN (storage area network)	A networking paradigm that provides easily reconfigurable connectivity between any subset of computers, disk storage and interconnecting hardware such as switches, hubs and bridges.
Service Group	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.
Service Group Dependency	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.
Shared Storage	Storage devices that are connected to and used by two or more systems.
shared volume	A volume that belongs to a shared disk group and is open on more than one node at the same time.
SFCFS (Storage Foundation Cluster File System)	The Veritas Cluster File System is an extension of the industry standard Veritas File System (VxFS). CFS allows the same file system to be simultaneously mounted on multiple nodes. Unlike other clustered file systems, CFS is a true SAN file system. All I/O takes place over the storage area network.
SNMP Notification	Simple Network Management Protocol (SNMP) developed to manage nodes on an IP network.
State	The current activity status of a resource, group or system. Resource states are given relative to both systems.
Storage Checkpoint	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.

System	The physical system on which applications and service groups reside. When a system is linked by VCS, it becomes a node in a cluster. See Node
types.cf	A file that describes standard resource types to the VCS engine; specifically, the data required to control a specific resource.
VCS (Veritas Cluster Server)	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.
Virtual IP Address	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.
VxFS (Veritas File System)	A component of the Veritas 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.
VxVM (Veritas Volume Manager)	A Symantec product installed on storage clients that enables management of physical disks as logical devices. enhances data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems.
VVR (Veritas Volume Replicator)	A data replication tool designed to contribute to an effective disaster recovery plan.

Index

Symbols

`$CRS_HOME`
 required information 304
 setting up 218, 303
`$ORACLE_BASE`
 setting up 220, 306
`$ORACLE_HOME`
 setting up 218, 303

A

adding
 users 99
agent operations
 Netlsnr agent 653
 Oracle agent 640
agents
 about 621
 ASM 605
 CFSMount 629
 CSSD 638
 CVMCluster 623
 CVMVolDg 627
 CVMVxconfigd 625
 log files 263, 359
 MultiPrivNIC 635
 of VCS 622
 PrivNIC 631
AIX
 operating systems supported 39
 required filesets 39
ASM
 about 593
 adding storage 605
 configuring with CVM
 about 597
 configuring disk group 598
 creating disk groups 601
 creating shared raw volumes 600
 installing ASM 599
 workflow 598
 dynamic multipathing 594

ASM (*continued*)
 integration with CVM 594
 main.cf 609
 monitoring utilities 597
 on CVM, rules 594
 requirements 597
 SF Oracle RAC configuration 595
 shutdown command 604
 startup command 604
 unlinking ODM 602
 verifying installation 601
 Veritas agents 605
 with CVM 606
 with Oracle service group 606
 with service groups
 Oracle RAC 10g 607
 Oracle RAC 11g 608
ASM disk groups
 adding storage 605
 creating
 Oracle DBCA procedure 601
ASMDG agent
 attribute definitions 662
 resource type 661
ASMDG agent attributes
 AgentDirectory 662
 DiskGroups 662
 Encoding 662
 EnvFile 662
 Home 662
 Owner 662
 Sid 662
ASMInst agent
 attribute definitions 658
 resource type 658
ASMInst agent attributes
 AgentDirectory 658
 Encoding 658
 EnvFile 658
 Home 658
 Owner 658
 Pfile 658

ASMinst agent attributes *(continued)*

Sid 658

attribute definitions

ASMDG agent 662

ASMinst 658

Netlsnr agent 654

Oracle agent 647

attributes

about 622

CFSMount agent 629

CVMCluster agent 624

CVMVolDg agent 624, 628

CVMVxconfigd agent 626

MultiPrivNIC agent 636

PrivNIC agent

optional 633

required 632

UseFence 127

Automatic Storage Management. *See* ASM

B

backup

of databases 30

basic monitoring 643

health check 643

process 643

C

CFS

creating databases 619

stopping applications 486, 498

unmounting file systems 486, 498

CFSMount agent 629

attributes 629

entry points 629

sample configuration 630–631

type definition 630

CFSTypes.cf 630

cluster

verifying operation 140

cluster configuration 96

Cluster File System. *See* CFS

cluster management

about 29

console, setting up 48

Cluster Volume Manager. *See* CVM

clusters

advanced setup 44

clusters *(continued)*

basic setup 42

four-node illustration 43

global setup 45

optional setup 44

replication 45

commands

hastatus 140

hasys 141

lltconfig 131

lltstat 136

vradmin 461

vxassist 454, 456, 619

vxdisksetup (initializing disks) 119

vxedit 619

vxprint 463

vxvol 454, 619

communication

inter-system 68

configuration

required information

for CVM 562

for global clusters 565

for I/O fencing 566

for secure clusters 563

for SF Oracle RAC, 559

for SMTP 564

for SNMP 565

for VCS 561

for VVR 566

VVR ports 106

VVR statistics collector 107

configuration files

removing 490, 502

See also main.cf samples

configuring VCS

adding users 99

event notification 100, 102

global clusters 103

secure mode 97

coordinator disks

DMP devices 115

for I/O fencing 115

setting up 125

CSSD agent 638

attributes 639

CVM

configuration worksheet 571

configuring 105

CVM (*continued*)

- configuring CVM groups 357
 - configuring groups 260
 - CVMTypes.cf file 624
 - installation worksheet 571–572
 - replication
 - primary site, main.cf 529
 - secondary site, main.cf 534
 - required information
 - for configuration 562
 - service group resources 259, 356
 - setting up default disk group 105
 - setting up naming scheme 105
 - upgrading protocol version 163
- CVM configuration** 105
- CVMCluster agent** 623
- attributes 624
 - entry points 623
 - sample configuration 625
 - type definition 624
- CVMTypes.cf**
- definition, CVMCluster agent 624
 - definition, CVMVolDg agent 628
 - definition, CVMVxconfigd agent 626
- CVMVolDg agent** 627
- attributes 628
 - entry points 627
 - sample configuration 629
 - type definition 628
- CVMVxconfigd agent** 625
- attributes 626
 - CVMTypes.cf 626
 - entry points 625
 - sample configuration 626
 - type definition 626

D

- data disks
 - for I/O fencing 114
- database. *See* repository database
- database administration
 - setting permissions 104
- Database FlashSnap 31
- database management 30
- database migration
 - from Oracle RAC 10g R1
 - to Oracle RAC 10g R2 269
 - from Oracle RAC 9i R2
 - to Oracle RAC 10g 266

database upgrade. *See* database migration

databases

- backing up 30
 - creating 617
 - creating tablespaces
 - on CFS 619
 - on raw volumes 618
 - recovering 30
 - snapshots 31
 - upgrade paths 265
 - upgrading to Oracle RAC 11g 336
- detail monitoring 643
- disk space
 - for SF Oracle RAC 35
- disks
 - adding and initializing 119
 - coordinator 125
 - testing with vxfsentsthdw 122
 - verifying node access 122
- Dynamic Storage Tiering 31

E

- environment variables
 - MANPATH 74
 - PATH 74

F

- files. *See* configuration files
- filesets
 - AIX 5.3, required 39

G

- GAB
 - port memberships 139
- GAB ports 140
- gabtab file
 - verifying after installation 131
- global clusters
 - about 29
 - adding VVR types 447
 - configuration 103
 - using VVR 448
 - configuration examples 465
 - configuration, using VVR 448
 - configuring VCS service groups 442
 - illustration of dependencies 463
 - overview 434
 - replicating database volumes 463

global clusters *(continued)*

- required information
 - for configuration 565
- requirements 432
- testing replication 442
- VVR preparation 446

groups

- log owner 464
- RVG 464

H

HADR

- installing license 78

hardware replication

- primary site
 - main.cf, on Oracle RAC 10g 538
- secondary site
 - main.cf, on Oracle RAC 10g 543

hastatus -summary command 140

hasys -display command 141

health check APIs 643

health check monitoring 643

I

I/O fencing

- checking disks 122
- required information
 - for configuration 566
- setting up 125
- shared storage 122

installation

- of Oracle RAC 10g 195
- of Oracle RAC 11g 275
- of Oracle RAC 11g Clusterware 321
- preparation 567
- required information
 - for SF Oracle RAC 561

installation programs 81

installation worksheets 567

installing

- Root Broker 62

installmp

- upgrading SF Oracle RAC 152

installsfrac

- configuring SF Oracle RAC 90
- installing Oracle RAC 10g 198
- installing Oracle RAC 11g 278
- installing Oracle RAC 11g Clusterware 321

installsfrac *(continued)*

- installing SF Oracle RAC 83
- response file 576
- upgrading SF Oracle RAC 154

L

license upgrades

- from Storage Foundation Cluster File System 188
- from Storage Foundation for Oracle 187
- from Storage Foundation HA for Oracle 187

licenses

- global clusters 77
- HADR 78
- obtaining 75
- removing 76
- removing keys 77
- SF Oracle RAC 75
- VVR 78

links

- private network 131

LLT

- optimizing settings 93
- preparing private interconnects 93
- settings, guidelines for 94
- verifying 136

lltconfig command 131

llthosts file

- verifying after installation 131

lltstat command 136

llttab file

- verifying after installation 131

M

main.cf

- after SF Oracle RAC installation 133–134

main.cf samples

- for CSSD resource 639
- for Oracle RAC 10g
 - with Oracle agent 511
 - with Oracle agent and VCSMC 514
 - without Oracle agent 508

for Oracle RAC 11g

- with Oracle agent 525
- without Oracle agent 522

hardware-based replication

- primary site, Oracle RAC 10g 538
- secondary site, Oracle RAC 10g 543

- main.cf samples *(continued)*
 - modifying VVR configuration 449
 - Oracle replication configurations 529
 - replication
 - primary site 529
 - secondary site 534
 - without Oracle agent, with VxSS 518
- MANPATH environment variable 74
- migration, of Oracle 362
- monitoring
 - basic 643
 - detail 643
- MultiPrivNIC
 - configuring 339
 - for UDP 299
- MultiPrivNIC agent 635
 - attributes 636
 - entry point 636
 - type definition 637
- MultiPrivNIC.cf 637
 - sample configuration 637

N

- Netlsnr agent
 - attribute definitions 654
 - operations 653
 - resource type 654
- Netlsnr agent attributes
 - AgentDebug 654
 - AgentDirectory 654
 - EnvFile 654
 - Home 654
 - Listener 654
 - LsnrPwd 654
 - MonScript 654
 - Owner 654
 - TnsAdmin 654
- nodes
 - adding Oracle RAC 10g nodes
 - about 371
 - adding systems to the cluster 373
 - configuring GAB 377
 - configuring LLT 377
 - configuring VCSMM 377
 - configuring VXFEN 377
 - enabling dlpi driver 372
 - installing SF Oracle RAC 373
 - starting Volume Manager 375
 - system requirements 373

- nodes *(continued)*
 - adding Oracle RAC 11g nodes
 - about 399
 - adding systems to cluster 401
 - configuring GAB 405
 - configuring LLT 405
 - configuring VCSMM 405
 - configuring VXFEN 405
 - enabling dlpi driver 400
 - installing SF Oracle RAC 401
 - starting Volume Manager 403
 - system requirements 401
 - adding to SFDB 419
 - adding to SFDB repository 391
 - preparing Oracle RAC 10g nodes
 - \$CRS_HOME, required information 386
 - about 381
 - configuring CVM 388
 - configuring private IP addresses 383
 - configuring shell limits 383
 - creating file system 386
 - creating Oracle disk groups 385
 - creating Oracle users and groups 382
 - identifying public virtual IP addresses 384
 - linking Veritas membership library 389
 - maximum processes per user 383
 - preparing \$CRS_HOME 385
 - preparing \$ORACLE_HOME 385
 - preparing OCR mount points 388
 - preparing VOTE disk mount points 388
 - relinking SF Oracle RAC libraries 391
 - setting oracle user capabilities 382
 - setting system parameters 382
 - setting up user equivalence 382
 - tuning VMM 383
 - preparing Oracle RAC 11g nodes
 - \$CRS_HOME, required information 414
 - about 409
 - configuring CVM 416
 - configuring kernel parameters 411
 - configuring private IP addresses 411
 - configuring shell limits 411
 - creating file system 414
 - creating Oracle disk groups 413
 - creating Oracle user and groups 410
 - identifying public virtual IP addresses 413
 - linking Veritas membership library 417
 - OCR mount point 416
 - preparing \$CRS_HOME 413

nodes *(continued)*

- preparing Oracle RAC 11g nodes *(continued)*
 - preparing \$ORACLE_BASE 415
 - preparing \$ORACLE_HOME 413
 - relinking SF Oracle RAC libraries 419
 - setting oracle user capabilities 410
 - setting system parameters 410
 - setting up user equivalence 410
 - tuning VMM 411
 - VOTE disk mount point 416
- rebooting 491, 503
- removing Oracle RAC 10g nodes
 - editing VCS configuration files 396
 - modifying VCS configuration 397
 - Oracle procedure 394
 - uninstalling SF Oracle RAC 394
 - workflow 393
- removing Oracle RAC 11g nodes
 - editing VCS configuration files 424
 - GAB configuration 425
 - LLT configuration 424
 - modifying VCS configuration 425
 - Oracle procedure 422
 - uninstalling SF Oracle RAC 422
 - workflow 421
- requirements, for adding 373
- nodes, removing 393, 421

O

OCR

- creating file system 223
 - editing main.cf 311
 - using CLI 308
- editing main.cf 226
- on CFS
 - required information 222, 308

operating systems

- supported 39

operations

- Netlsnr agent 653
- Oracle agent 640

options

- SF Oracle RAC configuration 87

Oracle

- agent log file 263, 359
- configuration worksheet 568
- configuring service groups 260, 357
- creating databases 618–619
- database replication 460

Oracle *(continued)*

- installation worksheet 568
- required patches 38
- shutdown options 641
- startup options 641
- supported versions 38

Oracle agent

- attribute definitions 647
- operations 640
- resource type 646

Oracle agent attributes

- AgentDebug 647
- AgentDirectory 647
- AutoEndBkup 647
- DetailMonitor 647
- EnvFile 647
- Home 647
- MonitorOption 647
- MonScript 647
- Owner 647
- Pfile 647
- Pword 647
- ShutDownOpt 647
- Sid 647
- StartUpOpt 647
- Table 647
- User 647

Oracle Cluster Registry

- OCR 222

Oracle RAC 10g

- agent log file 263
- configuring service groups
 - configuring VCS resources 260
 - dependencies 254
 - preventing automatic starting of
 - database 259
 - with Oracle agent 255
 - without Oracle agent 256
- workflow 253

- creating \$CRS_HOME 217

installation

- installing database 238
- verifying 243
- workflow 231

- installing Oracle Clusterware 233

- main.cf, with Oracle agent 511

- main.cf, with Oracle agent and VCSMC 514

- main.cf, without Oracle agent 508

Oracle RAC 10g *(continued)*

- migration from Oracle RAC 10g R1
 - about 269
 - installing Oracle RAC 10g R2 270
 - of existing database 271
 - post-upgrade tasks 271
 - pre-upgrade tasks 269
 - requirements 269
 - migration from Oracle RAC 9i R1
 - requirements 266
 - migration from Oracle RAC 9i R2
 - about 266
 - installing Oracle RAC 10g 267
 - of existing database 268
 - post-upgrade tasks 269
 - pre-upgrade tasks 266
 - post-installation
 - adding patchsets 244
 - relinking SF Oracle RAC libraries 249
 - pre-installation
 - configuring maximum processes per user 208
 - configuring private IP addresses 210
 - configuring shell values 207
 - creating OCR 221
 - creating Oracle disk groups 216
 - creating Oracle users and groups 201
 - creating VOTE disk volumes 221
 - identifying public virtual IP addresses 215
 - preparing \$CRS_HOME 218
 - preparing \$ORACLE_BASE 220
 - preparing \$ORACLE_HOME 218
 - setting system parameters 197
 - setting up Oracle RAC 10g Clusterware user 207
 - setting up Oracle user equivalence 206
 - setting up storage 198
 - starting installsfrac 198
 - tuning VMM 209
 - verifying GAB port membership 228
 - verifying SF Oracle RAC resources 228
 - workflow 195
 - removing database 484
 - replication, main.cf 529, 534
 - stopping instances 483
 - uninstalling 484
 - using Oracle interface 236
- ## Oracle RAC 10g Clusterware
- configuring private IP addresses 210

Oracle RAC 10g Clusterware *(continued)*

- creating file system 219
 - installation 233
 - installation workflow 218
 - installation, using Oracle interface 236
 - setting user capabilities 207
 - verifying installation 243
- ## Oracle RAC 11g
- agent log file 359
 - configuring service groups
 - configuring VCS resources 357
 - dependencies 352
 - preventing automatic starting of database 356
 - with Oracle agent 353
 - without Oracle agent 354
 - workflow 351
 - creating \$CRS_HOME 302
 - creating OCR 302
 - creating VOTE-disk volumes 302
 - installation
 - configuring service groups 336
 - installing database 323
 - of Oracle Clusterware 317
 - setting environment variables 317
 - verifying 328
 - workflow 315
 - main.cf, with Oracle agent 525
 - main.cf, without Oracle agent 522
 - manual installation
 - adding Oracle RAC 11g Clusterware to VCS 347
 - adding patches 346
 - configuring cluster interconnects 348
 - configuring MultiPrivNIC 339
 - configuring service groups 349
 - creating database 348
 - installing database binaries 345
 - installing Oracle Clusterware 343
 - linking ODM library 346
 - linking Veritas Membership library 341
 - setting environment variables 339
 - upgrading databases 349
 - workflow 337
 - post-installation
 - adding patches 329
 - adding UDP IP addresses 334
 - creating database 334
 - relinking SF Oracle RAC libraries 330

Oracle RAC 11g (*continued*)post-installation (*continued*)

upgrading databases 336

pre-installation

configuring MultiPrivNIC 299

configuring Networks Tuning

Parameters 290

configuring private IP addresses 291

configuring shell values 288

creating OCR 306

creating Oracle disk groups 302

creating Oracle users and groups 281

creating VOTE disk volumes 306

identifying public virtual IP addresses 301

maximum processes per user 289

preparing \$CRS_HOME 303

preparing \$ORACLE_BASE 306

preparing \$ORACLE_HOME 303

setting system parameters 277

setting up Oracle RAC 11g Clusterware
user 288

setting up Oracle user equivalence 286

setting up storage 278

starting installsfrac 278

tuning VMM 290

verifying GAB port membership 313

verifying SF Oracle RAC resources 313

workflow 275

removing database 496

setting environment variables 317

stopping instances 495

uninstalling 496

Oracle RAC 11g Clusterware

adding to VCS 347

configuring private IP addresses 291

creating file system 305

installation workflow 303

installing 321

setting user capabilities 288

verifying installation 328

Oracle RAC 11g, manual installation 338

P

patches

adding 329

for Oracle 38

patchsets

adding 244, 329

applying Oracle RAC 10g R1 244

patchsets (*continued*)

applying Oracle RAC 10g R2 246

PATH environment variable 74

PATH variable

VCS commands 136

performance management 31

permissions

setting 104

ports

GAB 140

primary site

creating SRL volume 454

setting up 435

setting up replication objects 455

VCS configuration 468, 472

PrivNIC agent 631

attributes 632

entry point 631

sample configuration 635

type definition 634

PrivNIC.cf 634

process monitoring 643

processes per user

for Oracle RAC 10g 208

for Oracle RAC 11g 289

R

raw volumes

creating

Oracle DBCA procedure 600

reinstallation

using response file 576

replication. *See* cluster replication

automatic synchronization 461

configuring on both sites 441

full synchronization with Checkpoint 461

modifying VCS configuration 464

options 433

primary sites

main.cf, Oracle RAC 10g 538

secondary sites

main.cf, Oracle RAC 10g 543

setting up primary site 435

setting up secondary site 438

supported hardware 432

supported software 432

supported technologies 37, 432

testing 442

using VVR 454

- replication (*continued*)
 - verifying status 462
- repository database
 - removing 497
- resource type
 - ASMDG 661
 - ASMInst 658
 - Netlsnr 654
 - Oracle 646
- resources
 - CVMVolDg 465
 - RVGSharedPri 465
- response file
 - about 575
 - configuration, sample 584
 - for reinstallation 576
 - install, sample 583
 - installation on other clusters 580
 - installsfrac 576
 - syntax 583
 - variables 585
 - VCS 576
- Root Broker
 - installing 62

S

- sample configuration files. *See* main.cf samples
- SCSI-3 persistent reservations
 - verifying 125
- secondary site
 - configuring replication 456
 - creating SRL volume 456
 - setting up 438
 - setting up disk groups 458
 - setting up RLINKs 457
- secure clusters
 - adding users 563
 - required information
 - for configuration 563
- service groups
 - configuring 336, 349
 - dependencies 254, 352
 - VCS, for global clusters 442
 - with Oracle agent 255, 353
 - without Oracle agent 256, 354
 - workflow 253, 351
- setting umask, before installing 67
- setting up inter-system communication 68

- SF Oracle RAC
 - about 25
 - coordinator disks 125
 - database components 30
 - high-level view 42
 - illustration of global cluster dependencies 463
 - optional components 28
 - overview, tasks 46
 - required components 27
- SF Oracle RAC configuration
 - about 87
 - checking systems 92
 - configuring clusters 96
 - configuring CVM 105
 - configuring VVR 106
 - creating configuration files 108
 - Ethernet settings 94
 - installers 89
 - LLT 93
 - of components 95
 - optimizing LLT settings 93
 - options 87
 - preparation
 - worksheets 567
 - required information 559
 - selecting options 91
 - setting permissions 104
 - specifying systems 91
 - starting processes 109
 - using installsfrac 90
 - using product installer 90
 - using response file 576
 - verifying 139
- SF Oracle RAC installation
 - applying license 84
 - installing filesets 86
 - mounting product disc 82
 - on other clusters
 - using response file 580
 - pre-installation tasks
 - mounting product disc 73
 - obtaining license 75
 - removing license 76
 - setting MANPATH 74
 - setting PATH 74
 - setting umask 67
 - setting up inter-system communication 68
 - setting up shared storage 73
 - synchronizing time settings 68

- SF Oracle RAC installation *(continued)*
 - pre-installation tasks *(continued)*
 - verifying systems 78
 - workflow 58
 - preinstallation information 33
 - preparation
 - worksheets 567
 - required information 561
 - requirements
 - hardware 34
 - miscellaneous 54
 - operating systems 39
 - software 38
 - specifying systems 84
 - using installsfrac 83
 - using product installer 83
 - using response file 576
 - verifying
 - cluster operations 136
 - GAB operations 136
 - LLT operations 136
 - VCS configuration file 134
- SF Oracle RAC uninstallation
 - preparation
 - removing database 484
 - removing repository database 485
 - stopping applications, CFS 486
 - stopping applications, VxFS 487
 - stopping Oracle instances 483
 - stopping VCS 487
 - uninstalling Oracle RAC 10g 484
 - unmounting CFS file systems 486
 - unmounting VxFS file systems 487
 - rebooting nodes 491
 - removing configuration files 490
 - using uninstallfrac 488
 - workflow 481
- SF Oracle RAC uninstallation, Oracle RAC 10g
 - removing filesets 488
- SF Oracle RAC uninstallation, Oracle RAC 11g
 - preparation
 - removing database 496
 - removing repository database 497
 - stopping applications, CFS 498
 - stopping applications, using VxFS 499
 - stopping instances 495
 - stopping VCS 499
 - uninstalling Oracle RAC 10g 496
 - unmounting CFS file systems 498
- SF Oracle RAC uninstallation, Oracle RAC 11g *(continued)*
 - preparation *(continued)*
 - unmounting VxFS file systems 499
 - rebooting nodes 503
 - removing configuration files 502
 - removing filesets 500
 - using uninstallfrac 500
 - workflow 493
- SF Oracle RAC upgrade
 - about 145
 - from pre-5.0 versions 154
 - from version 5.0 152
 - using RSH 152
 - using SSH 152
 - post-upgrade tasks
 - creating skgxn library 157
 - relinking Oracle RAC 10g 159
 - relinking Oracle RAC 11g 159
 - relinking Oracle RAC 9i 161
 - resetting VCS resource attribute 161
 - starting resources 160
 - upgrading CVM protocol version 163
 - upgrading disk group version 163
 - preparation 147
 - restarting systems 152
 - restoring configuration files 152, 154
 - stopping cluster resources 148
 - stopping CRS 149
 - stopping Oracle RAC 148
 - tasks, overview 146
 - upgrade paths 146
 - upgrading licenses 187
 - from Storage Foundation Cluster File System 188
 - from Storage Foundation for Oracle 187
 - from Storage Foundation HA for Oracle 187
 - using installmp 152
 - using installsfrac 154
 - verifying cluster protocol version 164
 - with minimum downtime 164
 - for Oracle RAC 10g clusters 164
 - for Oracle RAC 11g clusters 164
 - for Oracle RAC 9i clusters 178
- shared storage
 - setting up 73
- shell values
 - for oracle user 208
 - for users 289

- shutdown options 641
- SMTP
 - required information
 - for configuration 564
- SMTP email notification 100
- snapshots
 - of databases 31
- SNMP
 - required information
 - for configuration 565
- SNMP trap notification 102
- startup options 641
- Storage Checkpoint 30
- storage management 31
- Storage Mapping 31
- Storage Rollback 30
- Symantec Product Authentication Service 62, 97
 - about 29
 - setting up 48
- synchronizing time settings, before installing 68
- system check
 - SF Oracle RAC configuration 92
- system state attribute value 140

U

- uninstallation
 - of SF Oracle RAC 488
 - of SF Oracle RAC, on Oracle RAC 11g 500
 - Oracle RAC 10g 484
 - Oracle RAC 11g 496
- uninstallsfrac
 - removing filesets 488, 500
 - removing SF Oracle RAC, Oracle RAC 10g 394
 - removing SF Oracle RAC, Oracle RAC 11g 422
- upgrade paths, for SF Oracle RAC 146
- users
 - creating 201, 281
 - for SF Oracle RAC, adding 563

V

- variables
 - response files 585
- VCS
 - adding VVR resource types 447
 - agent log file 263, 359
 - command directory path variable 136
 - configuration files
 - GAB configuration 397

- VCS (*continued*)
 - configuration files (*continued*)
 - LLT configuration 396
 - configuration, for database volume replication 463
 - configuring resources 260, 357
 - configuring service groups 253, 351, 442
 - managing Oracle RAC 11g Clusterware 347
 - required information
 - for configuration 561
 - stopping 487, 499
- VCS configuration
 - for replication 464
- VCS notifications
 - setting up 49
- Veritas File System
 - stopping applications 487, 499
 - unmounting 487, 499
- Veritas Volume Replicator. *See* VVR
- virtual fire drill 645
- Virtual Memory Manager. *See* VMM
- VMM
 - configuring 209
 - configuring settings 290
- VOTE disks
 - creating file system 223
 - editing main.cf 311
 - using CLI 308
 - editing main.cf 226
 - on CFS
 - required information 222, 308
- vradm command 461
- VVR
 - about 29
 - adding resource types 447
 - configuration examples 465
 - configuration worksheet 571
 - configuring global clusters 448
 - configuring on both sites 441
 - configuring statistics collector 107
 - configuring tunable parameters 108
 - database volume replication
 - configuring VCS 463
 - defining heartbeat cluster objects 451
 - defining remote clusters 451
 - global cluster overview 445
 - installation worksheet 571–572
 - installing license 78
 - modifying configuration 449

VVR (*continued*)

- Oracle RAC database replication 460
 - port settings 106
 - primary site
 - creating SRL volume 454
 - setting up replication objects 455
 - replication
 - main.cf 529, 534
 - using automatic synchronization 461
 - using full synchronization with
 - Checkpoint 461
 - replication agents 433
 - replication, preparing 446
 - required information
 - for configuration 566
 - secondary site
 - configuring replication 456
 - creating SRL volume 456
 - setting up disk groups 458
 - setting up RLINKs 457
 - setting up primary site 435
 - setting up replication 454
 - setting up secondary site 438
 - starting VCS 476
 - testing replication 442
 - types of replication 433
 - VCS configuration
 - CVMolDG resource 465
 - log owner group 464
 - Oracle RAC database service group 465
 - primary site 468
 - RVG group 464
 - RVGSharedPri resource 465
 - secondary site 472
 - verifying replication 462
- VVR configuration 106
- VVR replication
- primary site
 - main.cf 529
 - secondary site
 - main.cf 534
- vxassist command 454, 456, 619
- vxdisksetup command 119
- VxFS. *See* Veritas File System
- vxprint command 463
- vxvol command 454, 619

W

- worksheets
 - for CVM 571
 - for Oracle RAC 568
 - for replication
 - SRDF 572
 - VVR 571
 - for SF Oracle RAC 567
 - for VVR 571