

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

Linux

6.0

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: 6.0

Document version: 6.0.3

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Documentation

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

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

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

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

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

Contents

Technical Support	4
Section 1 Installation overview and planning	25
Chapter 1 Introducing Veritas Storage Foundation for Oracle RAC	26
About Veritas Storage Foundation for Oracle RAC	26
Benefits of SF Oracle RAC	27
About SF Oracle RAC components	28
About SF Oracle RAC optional features	29
About VCS notifications	30
About campus clusters	30
About global clusters	30
About Veritas Volume Replicator	31
About database management using SF Oracle RAC	31
About Cluster Manager (Java Console)	31
About Veritas Operations Manager	31
Symantec Operations Readiness Tools	32
SF Oracle RAC cluster setup models	32
Typical configuration of four-node SF Oracle RAC cluster	33
Typical configuration of SF Oracle RAC clusters in secure mode	35
Typical configuration of VOM-managed SF Oracle RAC clusters	36
Typical configuration of SF Oracle RAC campus clusters for disaster recovery	37
Typical configuration of SF Oracle RAC global clusters for disaster recovery	39
Chapter 2 System requirements	41
Important preinstallation information	41
Hardware requirements	42
Supported operating systems	44
I/O fencing requirements	44

	Coordinator disk requirements for I/O fencing	44
	CP server requirements	44
	Supported database software	48
	Supported replication technologies for global clusters	48
	Discovering product versions and various requirement information	49
Chapter 3	Planning to install SF Oracle RAC	50
	Planning your network configuration	50
	Planning the public network configuration for Oracle RAC	51
	Planning the private network configuration for Oracle RAC	51
	Planning the storage	53
	Planning the storage for SF Oracle RAC	54
	Planning the storage for Oracle RAC	54
	Planning volume layout	58
	Planning file system design	59
	About planning to configure I/O fencing	59
	Typical SF Oracle RAC cluster configuration with disk-based I/O fencing	61
	Typical SF Oracle RAC cluster configuration with server-based I/O fencing	62
	Recommended CP server configurations	62
	Planning for cluster management	64
	Planning for disaster recovery	66
	Planning a campus cluster setup	66
	Planning a global cluster setup	67
	Data replication considerations	67
Chapter 4	Licensing SF Oracle RAC	68
	About Veritas SFHA Solutions product licensing	68
	About SF Oracle RAC licenses	69
	Setting or changing the Veritas SFHA Solutions product level for keyless licensing	72
	Installing Veritas SFHA Solutions product license keys	73

Section 2	Installation and configuration of SF Oracle RAC	74
Chapter 5	Preparing to install SF Oracle RAC	75
	Setting the umask before installation	75
	Synchronizing time settings on cluster nodes	75
	Mounting the product disc	76
	Setting the environment variables for SF Oracle RAC	76
	Configuring the I/O scheduler	77
	Optimizing LLT media speed settings on private NICs	78
	Guidelines for setting the media speed of the LLT interconnects	78
	Verifying the systems before installation	78
Chapter 6	Installing SF Oracle RAC	80
	About installing SF Oracle RAC	80
	About installation and configuration methods	80
	Installing SF Oracle RAC using the Veritas script-based installation program	82
	Installing SF Oracle RAC using the Veritas Web-based installation program	85
	Before using the Veritas Web-based installer	86
	Starting the Veritas Web-based installer	86
	Installing products with the Veritas Web-based installer	87
	Installing SF Oracle RAC using Kickstart	88
	Sample Kickstart configuration file	91
	Installing Veritas Storage Foundation for Oracle RAC using yum	93
Chapter 7	Preparing to configure SF Oracle RAC	98
	Setting up the CP server	98
	Planning your CP server setup	99
	Installing the CP server using the installer	100
	Configuring the CP server cluster in secure mode	101
	Setting up shared storage for the CP server database	101
	Configuring the CP server using the configuration utility	102
	Configuring the CP server manually	111
	Verifying the CP server configuration	113

Chapter 8	Configuring SF Oracle RAC	114
	About configuring SF Oracle RAC	114
	Configuring the SF Oracle RAC components using the script-based installer	115
	Configuring the SF Oracle RAC cluster	117
	Creation of SF Oracle RAC configuration files	133
	Stopping and starting SF Oracle RAC processes	133
	Configuring SF Oracle RAC using the Web-based installer	134
Chapter 9	Configuring SF Oracle RAC clusters for data integrity	140
	Setting up disk-based I/O fencing using installsfrac program	140
	Initializing disks as VxVM disks	140
	Identifying disks to use as coordinator disks	141
	Checking shared disks for I/O fencing	141
	Configuring disk-based I/O fencing using installsfrac program	146
	Setting up disk-based I/O fencing manually	149
	Setting up coordinator disk groups	149
	Creating I/O fencing configuration files	150
	Modifying VCS configuration to use I/O fencing	151
	Starting SF Oracle RAC on all nodes	152
	Verifying I/O fencing configuration	154
	Configuring SF Oracle RAC for data integrity using the Web-based installer	155
	Setting up server-based I/O fencing using installsfrac program	157
	Setting up server-based I/O fencing manually	157
	Preparing the CP servers manually for use by the SF Oracle RAC cluster	158
	Configuring server-based fencing on the SF Oracle RAC cluster manually	161
	Configuring CoordPoint agent to monitor coordination points	167
	Verifying server-based I/O fencing configuration	168
Chapter 10	Performing post-installation and configuration tasks	170
	Performing a postcheck on a node	170
	Verifying SF Oracle RAC installation using VCS configuration file	171
	Verifying LLT, GAB, and cluster operation	171
	Verifying LLT	172
	Verifying GAB	174

	Verifying the cluster	175
	Verifying the cluster nodes	176
	About enabling LDAP authentication for clusters that run in secure mode	179
	Enabling LDAP authentication for clusters that run in secure mode	180
	About configuring authentication for SFDB tools	186
	Configuring vxdbd for SFDB tools authentication	186
	Configuring Veritas Volume Replicator	187
	Running SORT Data Collector to collect configuration information	187
Section 3	Upgrade of SF Oracle RAC	189
Chapter 11	About upgrading SF Oracle RAC 6.0	190
	About types of upgrade	190
	Supported upgrade paths	191
Chapter 12	Performing a full upgrade to SF Oracle RAC 6.0	193
	About full upgrades	193
	Preparing to perform a full upgrade to SF Oracle RAC 6.0	194
	Pre-upgrade tasks for migrating the SFDB repository database	196
	Upgrading to SF Oracle RAC 6.0	197
	Upgrading SF Oracle RAC using the Veritas script-based installation program	199
	Upgrading Veritas Storage Foundation for Oracle RAC using the Veritas Web-based installer	201
	Upgrading SF Oracle RAC using a response file	203
Chapter 13	Performing a phased upgrade to SF Oracle RAC 6.0	207
	About phased upgrade	207
	Performing phased upgrade of SF Oracle RAC from version 5.1 Patch 3 and later releases	208
	Step 1: Performing pre-upgrade tasks on the first half of the cluster	209
	Step 2: Upgrading the first half of the cluster	211
	Step 3: Performing pre-upgrade tasks on the second half of the cluster	213
	Step 4: Performing post-upgrade tasks on the first half of the cluster	214
	Step 5: Upgrading the second half of the cluster	215

	Step 6: Performing post-upgrade tasks on the second half of the cluster	216
Chapter 14	Performing a rolling upgrade to SF Oracle RAC 6.0	219
	About rolling upgrades	219
	Preparing to perform a rolling upgrade to SF Oracle RAC 6.0	222
	Performing a rolling upgrade using the installer	224
	Performing a rolling upgrade using the script-based installer	224
	Performing a rolling upgrade of SF Oracle RAC using the Web-based installer	228
Chapter 15	Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0	232
	About upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0	232
	Upgrading from Storage Foundation Cluster File System for Oracle RAC to SF Oracle RAC 6.0	233
	Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0	233
Chapter 16	Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC	236
	Migration overview	236
	Sample configuration before and after migration	237
	Migration requirements	238
	Before you migrate	238
	Migrating to SF Oracle RAC 6.0	239
	Migrating Storage Foundation for Oracle HA to SF Oracle RAC	239
	Migrating a single instance Oracle database to Oracle RAC database	240
	Completing post-migration tasks	243
	Sample configuration files	244
	VCS configuration file for Storage Foundation for Oracle HA	245
	Oracle initialization parameter file for Storage Foundation for Oracle HA	246
	tnsnames.ora for Storage Foundation for Oracle HA	247
	listener.ora for Storage Foundation for Oracle HA	247
	VCS configuration file for SF Oracle RAC	248
	Oracle initialization parameter file for SF Oracle RAC	251

	tnsnames.ora file for SF Oracle RAC	252
	listener.ora file for SF Oracle RAC	252
Chapter 17	Performing post-upgrade tasks	253
	Relinking Oracle RAC libraries with the SF Oracle RAC libraries	253
	Re-joining the backup boot disk group into the current disk group	255
	Reverting to the backup boot disk group after an unsuccessful upgrade	255
	Setting or changing the product license level	256
	Upgrading disk layout versions	256
	Upgrading CVM protocol version and VxVM disk group version	257
	Post upgrade tasks for migrating the SFDB repository database	257
	Migrating from a 5.0 repository database to 6.0	258
	After upgrading from 5.0.x and before migrating SFDB	260
	Verifying the cluster	260
Section 4	Installation and upgrade of Oracle RAC	264
Chapter 18	Before installing Oracle RAC	265
	Important preinstallation information for Oracle RAC	265
	About preparing to install Oracle RAC	265
	Preparing to install Oracle RAC using the SF Oracle RAC installer or manually	266
	Identifying the public virtual IP addresses for use by Oracle	268
	Setting the kernel parameters	268
	Verifying that RPMs and patches required by Oracle are installed	269
	Verifying the user <code>nobody</code> exists	269
	Launching the SF Oracle RAC installer	269
	Creating users and groups for Oracle RAC	271
	Creating storage for OCR and voting disk	275
	Configuring private IP addresses for Oracle RAC	292
	Verifying that multicast is functional on all private network interfaces	311
	Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually	312
	Setting up user equivalence	320
	Verifying the systems for Oracle RAC installation	320

Chapter 19	Installing Oracle RAC	324
	About installing Oracle RAC	324
	Installing the Oracle Clusterware/Grid Infrastructure software	325
	Installing Oracle Clusterware/Grid Infrastructure using the SF	
	Oracle RAC script-based installer	326
	Installing Oracle Clusterware/Grid Infrastructure using the SF	
	Oracle RAC Web-based installer	329
	Installing Oracle Clusterware/Grid Infrastructure using the Oracle	
	Universal Installer	332
	Configuring LLT links in the GPnP profile	333
	Installing the Oracle RAC database software	334
	Installing the Oracle RAC database using the SF Oracle RAC	
	script-based installer	335
	Installing the Oracle RAC database using the SF Oracle RAC	
	Web-based installer	337
	Installing the Oracle RAC database using the Oracle Universal	
	Installer	340
	Verifying the Oracle Clusterware/Grid Infrastructure and database	
	installation	341
Chapter 20	Performing Oracle RAC post-installation tasks	343
	Adding Oracle RAC patches or patchsets	343
	Configuring the CSSD resource	344
	Configuring the CSSD resource using the SF Oracle RAC	
	script-based installer	345
	Configuring the CSSD resource using the SF Oracle RAC	
	Web-based installer	347
	Configuring the CSSD resource manually	348
	Preventing automatic startup of Oracle Clusterware/Grid	
	Infrastructure	350
	Relinking the SF Oracle RAC libraries with Oracle RAC	350
	Relinking the SF Oracle RAC libraries with Oracle RAC using the	
	SF Oracle RAC script-based installer	351
	Relinking the SF Oracle RAC libraries with Oracle RAC using the	
	SF Oracle RAC Web-based installer	352
	Relinking SF Oracle RAC libraries with Oracle RAC	
	manually	353
	Updating the CSS misscount setting	355
	Creating the Oracle RAC database	356
	Configuring VCS service groups for Oracle RAC	356
	Supported types of database management	356
	Sample service group configurations	357

	Configuring VCS service groups manually for Oracle databases	361
	Managing database restart after failure	364
	Location of VCS log files	365
	Preventing automatic database startup	365
	Removing permissions for communication	366
	Configuring the SFDB repository database after installation	366
Chapter 21	Upgrading Oracle RAC	367
	Supported upgrade paths	367
	Preparing to upgrade Oracle RAC	368
	Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g	368
	Upgrading Oracle RAC binaries	369
	Migrating the Oracle RAC database	370
	Performing post-upgrade tasks	370
Section 5	Automated installation using response files	372
Chapter 22	About response files	373
	About response files	373
	Modular deployments using response files	375
	About the <code>-makeresponsefile</code> option	376
	Response file syntax	376
	Guidelines for creating the SF Oracle RAC response file	377
	Installation scenarios for response files	382
Chapter 23	Installing and configuring SF Oracle RAC using a response file	383
	Installing and configuring SF Oracle RAC	383
	Sample response file for installing and configuring SF Oracle RAC	385
Chapter 24	Configuring I/O fencing using a response file	386
	Configuring I/O fencing using response files	386
	Sample response file for configuring disk-based I/O fencing	387
	Sample response file for configuring server-based I/O fencing	387

Chapter 25	Installing Oracle RAC using a response file	389
	About installing Oracle RAC using response files	389
	Before you install	389
	Installing Oracle RAC	391
	Sample response file for installing Oracle RAC	392
Chapter 26	Installing SF Oracle RAC and Oracle RAC using a response file	394
	About installing SF Oracle RAC and Oracle RAC using response files	394
	Information required in the SF Oracle RAC response file	395
	Before you install	396
	Installing SF Oracle RAC and Oracle RAC	398
	Sample response file for installing SF Oracle RAC and Oracle RAC	399
Chapter 27	Response file variable definitions	402
	Response file variables for installing SF Oracle RAC	402
	Response file variables to configure Veritas Storage Foundation for Oracle RAC	405
	Response file variables to configure disk-based I/O fencing	414
	Response file variables to configure server-based I/O fencing	415
	Response file variable definitions for Oracle RAC	416
Section 6	Adding or removing nodes from an SF Oracle RAC cluster	434
Chapter 28	Adding a node to SF Oracle RAC clusters	435
	About adding a node to a cluster	435
	Before adding a node to a cluster	436
	Adding a node to a cluster using the product installer	439
	Adding a node using the Web-based installer	442
	Adding the node to a cluster manually	443
	Starting Volume Manager on the new node	456
	Configuring cluster files on the new node	456
	Configuring cluster files on the new node	447
	Starting fencing on the new node	461
	Configuring CVM and CFS on the new node	465
	Starting VCS after adding the new node	462

	Configuring the ClusterService group for the new node	455
	Starting Volume Manager on the new node	456
	Configuring cluster files on the new node	456
	Setting up the node to run in secure mode	458
	Starting fencing on the new node	461
	Starting VCS after adding the new node	462
	Configuring server-based fencing on the new node	463
	Configuring CVM and CFS on the new node	465
	Preparing the new node for installing Oracle RAC	466
	Preparing the new nodes for installing Oracle RAC using the SF	
	Oracle RAC script-based installer	467
	Preparing the new nodes for installing Oracle RAC using the SF	
	Oracle RAC Web-based installer	474
	Preparing the new node manually for installing Oracle RAC	482
	Adding the new node to Oracle RAC	490
	Adding nodes to a cluster that is using authentication for SFDB	
	tools	490
	Updating the Storage Foundation for Databases (SFDB) repository	
	after adding a node	491
	Sample configuration file for adding a node to the cluster	492
Chapter 29	Removing a node from SF Oracle RAC clusters	496
	About removing a node from a cluster	496
	Removing a node from a cluster	497
	Modifying the VCS configuration files on existing nodes	498
	Removing the node configuration from the CP server	501
	Removing security credentials from the leaving node	502
	Updating the Storage Foundation for Databases (SFDB) repository	
	after removing a node	502
	Sample configuration file for removing a node from the cluster	502
Section 7	Configuration of disaster recovery	
	environments	507
Chapter 30	Configuring a campus cluster setup for disaster	
	recovery	508
	Setting up an SF Oracle RAC campus cluster for disaster	
	recovery	508
	Preparing to set up a campus cluster	509
	Configuring VxVM disk groups for Oracle RAC on a campus	
	cluster	510

	Tuning guidelines for campus clusters	513
	Best practices for a campus cluster	513
Chapter 31	Setting up a replicated global cluster	514
	Replication in the SF Oracle RAC environment	514
	About setting up a global cluster in an SF Oracle RAC environment	515
	Configuring an SF Oracle RAC global cluster at the primary site	517
	Configuring an SF Oracle RAC cluster at the secondary site	519
	Configuring replication for SF Oracle RAC clusters at both sites	523
	Modifying the ClusterService group for a global SF Oracle RAC cluster	524
	Modifying the global SF Oracle RAC cluster configuration using the wizard	525
	Defining the remote SF Oracle RAC cluster and heartbeat objects	526
	Configuring the VCS service groups for global SF Oracle RAC clusters	529
Chapter 32	Configuring a global cluster using VVR	530
	About configuring a global SF Oracle RAC cluster using VVR for replication	530
	Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): setting up the primary site	531
	Creating the data and SRL volumes on the primary site	532
	Setting up the Replicated Volume Group (RVG) on the primary site	533
	Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): setting up the secondary site	534
	Creating the data and SRL volumes on the secondary site	535
	Editing the /etc/vx/vras/.rdg files	535
	Setting up IP addresses for RLINKs on each cluster	536
	Setting up the disk group on secondary site for replication	537
	Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): starting replication	539
	Starting replication using automatic synchronization	539
	Starting replication using full synchronization with Storage Checkpoint	540
	Verifying replication status	540
	Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): configuring cluster resources	541
	About modifying the VCS configuration for replication	542

	Configuration examples before and after modification	543
	Modifying the VCS Configuration on the Primary Site	546
	Modifying the VCS Configuration on the Secondary Site	550
	Managing a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR)	555
	Migrating the role of primary site to the secondary site	557
	Taking over the primary role by the remote cluster	558
Section 8	Uninstallation of SF Oracle RAC	562
Chapter 33	Preparing to uninstall SF Oracle RAC from a cluster	563
	About uninstalling SF Oracle RAC from a cluster	563
	Options for uninstalling SF Oracle RAC	564
	Preparing to uninstall SF Oracle RAC from a cluster	565
	Stopping Oracle instances	566
	Backing up the Oracle database	567
	Unlinking the SF Oracle RAC libraries from Oracle RAC	567
	Uninstalling Oracle RAC (optional)	567
	Removing root disk encapsulation	568
	Stopping the applications that use CVM or CFS (outside of VCS control)	569
	Unmounting CFS file systems (outside of VCS control)	569
	Stopping VCS	569
	Stopping the applications that use VxVM or VxFS (outside of VCS control)	570
	Unmounting VxFS file systems (outside of VCS control)	570
Chapter 34	Uninstalling SF Oracle RAC from a cluster	572
	Uninstalling SF Oracle RAC from a cluster	572
	Removing the CP server configuration using the removal script	573
	Removing the SF Oracle RAC RPMs	577
	Removing other configuration files (optional)	579
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	579
	Uninstalling SF Oracle RAC with the Veritas Web-based installer	580
	Uninstalling SF Oracle RAC using a response file	581
	Response file variables to uninstall SF Oracle RAC	582
	Sample response file for uninstalling SF Oracle RAC	583

Section 9	Installation reference	584
Appendix A	Installation scripts	585
	About Veritas Storage Foundation and High Availability Solutions	
	installation scripts	585
	Installation script options	586
	About using the postcheck option	591
Appendix B	Tunable files for installation	593
	About setting tunable parameters using the installer or a response	
	file	593
	Setting tunables for an installation, configuration, or upgrade	594
	Setting tunables with no other installer-related operations	595
	Setting tunables with an un-integrated response file	596
	Preparing the tunables file	597
	Setting parameters for the tunables file	597
	Tunables value parameter definitions	598
Appendix C	SF Oracle RAC installation and configuration	
	sample values	605
	About the installation and configuration worksheets	605
	SF Oracle RAC worksheet	606
	Veritas Cluster Server component information	609
	I/O fencing information	610
	SF Oracle RAC add user information	611
	Global cluster information	612
	Oracle RAC worksheet	613
	Replicated cluster using VVR worksheet	625
	Replicated cluster using SRDF worksheet	626
	Required installation information for Oracle Clusterware/Grid	
	Infrastructure	628
	Required installation information for Oracle database	629
Appendix D	Sample configuration files	631
	About VCS configuration file	631
	About the LLT and GAB configuration files	632
	About I/O fencing configuration files	634
	Sample configuration files	637
	sfrac02_main.cf file	638
	sfrac03_main.cf file	639

	sfrac04_main.cf file	640
	sfrac05_main.cf file	642
	sfrac06_main.cf file	643
	sfrac07_main.cf and sample08_main.cf files	644
	sfrac09_main.cf and sample10_main.cf files	646
	sfrac11_main.cf file	649
	sfrac12_main.cf and sample13_main.cf files	650
	sfrac14_main.cf file	653
	Sample configuration files for CP server	654
Appendix E	Setting up inter-system communication	661
	About using ssh or rsh with the Veritas installer	661
	Setting up inter-system communication	662
	Setting up ssh on cluster systems	662
Appendix F	Automatic Storage Management	664
	About ASM in SF Oracle RAC environments	664
	ASM configuration with SF Oracle RAC	665
	Configuring ASM in SF Oracle RAC environments	666
	Unlinking the SF Oracle RAC ODM library	667
	Creating database storage on ASM	667
	Creating ASM disk groups and instances	668
	Verifying the ASM setup	669
	Configuring VCS service groups for database instances on ASM	670
	Sample configuration file Veritas CVM and ASM main.cf file	672
Appendix G	Creating a test database	673
	About creating a test database	673
	Creating a database for Oracle	673
	Creating the database storage on raw volumes	674
	Creating the database storage on CFS	675
Appendix H	High availability agent information	676
	About agents	676
	VCS agents included within SF Oracle RAC	677
	VCS agents for Oracle included within SF Oracle RAC	677
	CVMCluster agent	678
	Entry points for CVMCluster agent	678
	Attribute definition for CVMCluster agent	679
	CVMCluster agent type definition	680

CVMCluster agent sample configuration	680
CVMVxconfigd agent	680
Entry points for CVMVxconfigd agent	681
Attribute definition for CVMVxconfigd agent	681
CVMVxconfigd agent type definition	683
CVMVxconfigd agent sample configuration	683
CVMVoIdg agent	683
Entry points for CVMVoIdg agent	683
Attribute definition for CVMVoIdg agent	684
CVMVoIdg agent type definition	685
CVMVoIdg agent sample configuration	686
CFSMount agent	686
Entry points for CFSMount agent	686
Attribute definition for CFSMount agent	687
CFSMount agent type definition	689
CFSMount agent sample configuration	690
CFSfsckd agent	690
Entry points for CFSfsckd agent	690
Attribute definition for CFSfsckd agent	691
CFSfsckd agent type definition	692
CFSfsckd agent sample configuration	693
PrivNIC agent	693
Functions of the PrivNIC agent	693
Attributes of the PrivNIC agent	694
States of the PrivNIC agent	697
Sample service group configuration with the PrivNIC agent	697
Type definition of the PrivNIC resource	698
Sample configuration of the PrivNIC resource	698
MultiPrivNIC agent	699
Managing high availability of private interconnects	699
Functions of the MultiPrivNIC agent	700
Attributes of the MultiPrivNIC agent	700
States of the MultiPrivNIC agent	702
Sample service group configuration with the MultiPrivNIC agent	702
Type definition of the MultiPrivNIC resource	702
Sample configuration of the MultiPrivNIC resource	703
CSSD agent	703
Functions of the CSSD agent	704
Attributes of the CSSD agent	704
States of the CSSD agent	705
Sample service group configurations with the CSSD agent	705
Type definition of the CSSD resource	706

Sample configuration of the CSSD resource	706
VCS agents for Oracle	707
Oracle agent functions	707
Resource type definition for the Oracle agent	713
Netlsnr agent functions	720
Resource type definition for the Netlsnr agent	721
ASMDG agent functions	726
Resource type definition for the ASMDG agent	726
CRSResource agent	729
Functions of the CRSResource agent	729
States of the CRSResource agent	730
Attributes of the CRSResource agent	730
VCS service group dependencies with the CRSResource agent	731
Resource type definition for the CRSResource agent	736
Sample configuration for the CRSResource agent	736

Appendix I	SF Oracle RAC deployment scenarios	738
	SF Oracle RAC cluster with UDP IPC and PrivNIC agent	738
	SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent	740
	SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent	743
	SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent	745
	Configuration diagrams for setting up server-based I/O fencing	747
	Two unique client clusters served by 3 CP servers	747
	Client cluster served by highly available CPS and 2 SCSI-3 disks	748
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	750
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	752
	Deploying Storage Foundation for Databases (SFDB) tools in a Storage Foundation for Oracle RAC environment	754

Appendix J Compatability issues when installing Veritas
 Storage Foundation for Oracle RAC with other
 products 756

 Installing, uninstalling, or upgrading Storage Foundation products when
 other Veritas products are present 756

 Installing, uninstalling, or upgrading Storage Foundation products when
 VOM is already present 757

 Installing, uninstalling, or upgrading Storage Foundation products when
 ApplicationHA is already present 757

 Installing, uninstalling, or upgrading Storage Foundation products when
 NetBackup is already present 758

Index 759

Glossary 767

Installation overview and planning

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

Introducing Veritas Storage Foundation for Oracle RAC

This chapter includes the following topics:

- [About Veritas Storage Foundation for Oracle RAC](#)
- [About SF Oracle RAC components](#)
- [About SF Oracle RAC optional features](#)
- [About Cluster Manager \(Java Console\)](#)
- [About Veritas Operations Manager](#)
- [Symantec Operations Readiness Tools](#)
- [SF Oracle RAC cluster setup models](#)

About Veritas Storage Foundation for Oracle RAC

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

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

Benefits of SF Oracle RAC

SF Oracle RAC provides the following benefits:

- Support for file system-based management. SF Oracle RAC provides a generic clustered file system technology for storing and managing Oracle data files as well as other application data.
- Support for high-availability of cluster interconnects.
 The PrivNIC/MultiPrivNIC agents provide maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
 See the following Technote regarding co-existence of PrivNIC/MultiPrivNIC agents with Oracle RAC 11.2.0.2:
<http://www.symantec.com/business/support/index?page=content&id=TECH145261>
- Use of Cluster File System and Cluster Volume Manager for placement of Oracle Cluster Registry (OCR) and voting disks. These technologies provide robust shared block interfaces for placement of OCR and voting disks. In the absence of SF Oracle RAC, separate LUNs need to be configured for OCR and voting disks.
- Support for a standardized approach toward application and database management. Administrators can apply their expertise of Veritas technologies toward administering SF Oracle RAC.
- Increased availability and performance using Veritas Dynamic Multi-Pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the Host Bus Adapters (HBA), Storage Area Network (SAN) switches, and storage arrays.
- Easy administration and monitoring of multiple SF Oracle RAC clusters using Veritas Operations Manager.
- VCS OEM plug-in provides a way to monitor SF Oracle RAC resources from the OEM console.
 For more information, see the *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases* guide.
- Improved file system access times using Oracle Disk Manager (ODM).
- Ability to configure Oracle Automatic Storage Management (ASM) disk groups over CVM volumes to take advantage of Veritas Dynamic Multi-Pathing (DMP).
- Enhanced scalability and availability with access to multiple Oracle RAC instances per database in a cluster.
- Support for backup and recovery solutions using volume-level and file system-level snapshot technologies, Storage Checkpoints, and Database Storage Checkpoints.

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

- Support for space optimization using periodic deduplication in a file system to eliminate duplicate data without any continuous cost.
For more information, see the Veritas Storage Foundation Administrator's documentation.
- Ability to fail over applications with minimum downtime using Veritas Cluster Server (VCS) and Veritas Cluster File System (CFS).
- Prevention of data corruption in split-brain scenarios with robust SCSI-3 Persistent Group Reservation (PGR) based I/O fencing or Coordination Point Server-based I/O fencing. The preferred fencing feature also enables you to specify how the fencing driver determines the surviving subcluster.
- Support for sharing application data, in addition to Oracle database files, across nodes.
- Support for policy-managed databases in Oracle RAC 11g Release 2.
- Fast disaster recovery with minimal downtime and interruption to users. Users can transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites. If a node fails, clients that are attached to the failed node can reconnect to a surviving node and resume access to the shared database.
- Verification of disaster recovery configuration using fire drill technology without affecting production systems.
- Support for a wide range of hardware replication technologies as well as block-level replication using VVR.
- Support for campus clusters with the following capabilities:
 - Consistent detach with Site Awareness
 - Site aware reads with VxVM mirroring
 - Monitoring of Oracle resources
 - Protection against split-brain scenarios

About SF Oracle RAC components

[Table 1-1](#) lists the components of SF Oracle RAC.

Table 1-1 SF Oracle RAC components

Component	Description
Cluster Volume Manager	Cluster Volume Manager (CVM) enables simultaneous access to the shared volumes that are based on technology from Veritas Volume Manager (VxVM).
Cluster File System	Cluster File System (CFS) enables simultaneous access to the shared file systems that are based on technology from Veritas File System (VxFS).
Veritas Cluster Server	Veritas Cluster Server (VCS) manages Oracle RAC databases and infrastructure components in a clustered environment.
Veritas I/O fencing	Veritas I/O fencing protects the data on shared disks using SCSI-3 Persistent Group Reservations when nodes in a cluster detect a change in the network cluster membership with a potential split brain condition.
Oracle Disk Manager	Oracle Disk Manager (ODM) is a disk and file management interface that is provided by Oracle to improve disk I/O performance. ODM enables Oracle to allocate and release disk space, manage tablespaces, and read/write disk blocks directly. SF Oracle RAC uses a custom driver that enables applications to use ODM for enhanced file system performance and easy file administration.
RAC Extensions	RAC Extensions manage the cluster membership.

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

About SF Oracle RAC optional features

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

- VCS notifications
See [“About VCS notifications”](#) on page 30.
- Campus clusters
See [“About campus clusters”](#) on page 30.
- Global clusters
See [“About global clusters”](#) on page 30.

- Storage Foundation Database Management tools
See [“About database management using SF Oracle RAC”](#) on page 31.
- Veritas Volume Replicator
See [“About Veritas Volume Replicator”](#) on page 31.

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

About VCS notifications

You can configure both Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) notifications for VCS. Symantec recommends you to configure at least one of these notifications. You have the following options:

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

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

About campus clusters

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

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

About global clusters

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

About Veritas Volume Replicator

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

About database management using SF Oracle RAC

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

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

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers administration capabilities for your cluster. Use the different views in the Java Console to monitor and manage clusters and Veritas Cluster Server (VCS) objects, including service groups, systems, resources, and resource types. You cannot manage the new features of this release using the Java Console.

See *Veritas Cluster Server Administrator's Guide*.

If you want to manage a single cluster using Cluster Manager (Java Console), the latest version is available for download from https://www4.symantec.com/Vrt/offer?a_id=89446. You will need a (free) SymAccount for downloading.

The Veritas Cluster Server Management Console is deprecated. Symantec recommends using Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

About Veritas Operations Manager

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

Veritas Operations Manager provides a centralized management console for Veritas Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports.

You can download Veritas Operations Manager at no charge at <http://go.symantec.com/vom>.

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

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. If you want to continue using VEA, a software version is available for download from <http://www.symantec.com/operations-manager/support>. Veritas Storage Foundation Management Server is deprecated.

If you want to manage a single cluster using Cluster Manager (Java Console), a version is available for download from https://www4.symantec.com/Vrt/offer?a_id=89446. You cannot manage the new features of this release using the Java Console. Veritas Cluster Server Management Console is deprecated.

Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

Among its broad set of features, SORT lets you do the following:

- Generate server-specific reports that describe how to prepare your servers for installation or upgrade of Symantec enterprise products.
- Access a single site with the latest production information, including patches, agents, and documentation.
- Create automatic email notifications for changes in patches, documentation, and array-specific modules.

To access SORT, go to:

<https://sort.symantec.com>

SF Oracle RAC cluster setup models

SF Oracle RAC supports a variety of cluster configurations.

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

- Basic setup
See [“Typical configuration of four-node SF Oracle RAC cluster”](#) on page 33.
- Secure setup


See [“Typical configuration of SF Oracle RAC clusters in secure mode”](#) on page 35.

- Central management setup
See [“Typical configuration of VOM-managed SF Oracle RAC clusters”](#) on page 36.
- Campus cluster setup
See [“Typical configuration of SF Oracle RAC campus clusters for disaster recovery”](#) on page 37.
- Global cluster setup
See [“Typical configuration of SF Oracle RAC global clusters for disaster recovery”](#) on page 39.

Typical configuration of four-node SF Oracle RAC cluster

[Figure 1-1](#) depicts a high-level view of a basic SF Oracle RAC configuration for a four-node cluster.

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



A basic topology has the following layout and characteristics:

- Multiple client applications that access nodes in the cluster over a public network.
- Nodes that are connected by at least two private network links (also called cluster interconnects) using 100BaseT or gigabit Ethernet controllers on each system.
If the private links are on a single switch, isolate them using VLAN.
- Nodes that are connected to iSCSI or Fibre Channel shared storage devices over SAN.
All shared storage must support SCSI-3 PR.


- Nodes that are connected with private network links using similar network devices and matching port numbers.
For example, if you use eth1 on one end of a link, it is recommended that you use eth1 on the other end too.
- The Oracle Cluster Registry and vote disks configured on the shared storage that is available to each node. The shared storage for Oracle Cluster Registry and vote disks can be a cluster file system or ASM disk groups created using raw VxVM volumes.
- Three or more odd number of disks or LUNs used as coordinator disks for I/O fencing.
- VCS manages the resources that are required by Oracle RAC. The resources must run in parallel on each node.

Typical configuration of SF Oracle RAC clusters in secure mode

Enabling secure mode for SF Oracle RAC guarantees that all inter-system communication is encrypted and that security credentials of users are verified.

[Figure 1-2](#) illustrates typical configuration of SF Oracle RAC clusters in secure mode.

Figure 1-2 Typical configuration of SF Oracle RAC clusters in secure mode




Typical configuration of VOM-managed SF Oracle RAC clusters

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

See [“About Veritas Operations Manager”](#) on page 31.

[Figure 1-3](#) illustrates a typical setup of SF Oracle RAC clusters that are centrally managed using Veritas Operations Manager.

Figure 1-3 Typical configuration of VOM-managed clusters



Typical configuration of SF Oracle RAC campus clusters for disaster recovery

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


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

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

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

Figure 1-4 illustrates a basic campus cluster setup.

Figure 1-4 Basic campus cluster setup




For more information, see the chapter *Configuring a campus cluster setup for disaster recovery*.

Typical configuration of SF Oracle RAC global clusters for disaster recovery

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

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

Figure 1-5 Global clusters



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

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

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

For information on supported replication technologies:

See [“Supported replication technologies for global clusters”](#) on page 48.

System requirements

This chapter includes the following topics:

- [Important preinstallation information](#)
- [Hardware requirements](#)
- [Supported operating systems](#)
- [I/O fencing requirements](#)
- [Supported database software](#)
- [Supported replication technologies for global clusters](#)
- [Discovering product versions and various requirement information](#)

Important preinstallation information

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

- Hardware compatibility list for information about supported hardware:
<http://www.symantec.com/docs/TECH170013>
- Latest information on support for Oracle database versions:
www.symantec.com/docs/DOC4039
- General information regarding the release, installation notes, known issues, and fixed issues:
See Veritas Storage Foundation for Oracle RAC Release Notes.
- Oracle documentation for additional requirements pertaining to your version of Oracle.

Hardware requirements

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

For basic clusters See [Table 2-1](#) on page 42.

For campus clusters See [Table 2-2](#) on page 43.

Table 2-1 Hardware requirements for basic clusters

Item	Description
SF Oracle RAC systems	Two to sixteen systems with two or more CPUs. For details on the additional requirements for Oracle, see the Oracle documentation.
DVD drive	A DVD drive on one of the nodes in the cluster.
Disks	SF Oracle RAC requires that all shared storage disks 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. The minimum size required for a coordinator disk is 128 MB.
Disk space	You can evaluate your systems for available disk space by running the product installation program. Navigate to the product directory on the product disc and run the following command: <code># ./installsf rac -precheck node_name</code> You can also use the Veritas Web-based installation program to determine the available disk space. For details on the additional space that is required for Oracle, see the Oracle documentation.
RAM	Each SF Oracle RAC system requires at least 2 GB. For Oracle RAC requirements, see the Oracle Metalink document: 169706.1
Swap space	See the Oracle Metalink document: 169706.1

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

Item	Description
Network	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. Oracle requires that all nodes use the IP addresses from the same subnet.
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

[Table 2-2](#) lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

Table 2-2 Hardware requirements for campus clusters

Item	Description
Storage	■ The storage switch (to which each host on a site connects) must have access to storage arrays at all the sites. ■ Volumes must be mirrored with storage allocated from at least two sites. ■ DWDM links are recommended between sites for storage links. DWDM works at the physical layer and requires multiplexer and de-multiplexer devices. ■ The storage and networks must have redundant-loop access between each node and each storage array to prevent the links from becoming a single point of failure.
Network	■ Oracle requires that all nodes use the IP addresses from the same subnet. ■ Symantec recommends a common cross-site physical infrastructure for storage and LLT private networks.
I/O fencing	I/O fencing requires placement of a third coordinator point at a third site. The DWDM can be extended to the third site or the iSCSI LUN at the third site can be used as the third coordination point. Alternatively Coordination Point Server can be deployed at the third remote site as an arbitration point.

Supported operating systems

For information on supported operating systems, see the *Veritas Storage Foundation for Oracle RAC Release Notes*.

I/O fencing requirements

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

- Coordinator disks
See [“Coordinator disk requirements for I/O fencing”](#) on page 44.
- CP servers
See [“CP server requirements”](#) on page 44.

Note: Irrespective of whether you use coordinator disks or CP server for I/O fencing, ensure that the shared storage supports SCSI-3 persistent reservations.

Coordinator disk requirements for I/O fencing

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

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

CP server requirements

SF Oracle RAC 6.0 clusters (application clusters) support coordination point servers (CP servers) which are hosted on the following VCS and SFHA versions:

- VCS 6.0, 5.1SP1, or 5.1 single-node cluster

Single-node VCS clusters with VCS 5.1 SP1 RP1 and later or VCS 6.0 and later that hosts CP server does not require LLT and GAB to be configured.

- SFHA 6.0, 5.1SP1, or 5.1 cluster

Warning: Before you upgrade 5.1 CP server nodes to use VCS or SFHA 6.0, you must upgrade all the application clusters that use this CP server to version 6.0. Application clusters at version 5.1 cannot communicate with CP server that runs VCS or SFHA 5.1 SP1 or later.

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

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

Note: SF Oracle RAC requires at least 3 coordination points for I/O fencing.

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

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

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

Table 2-3 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var ■ 10 MB in /etc (for the CP server database)
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.

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

Hardware required	Description
RAM	Each CP server requires at least 512 MB.
Network	Network hardware capable of providing TCP/IP connection between CP servers and SF Oracle RAC clusters (application clusters).

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

Table 2-4 CP server supported operating systems and versions

CP server	Operating system and version
CP server hosted on a VCS single-node cluster or on an SFHA cluster	CP server supports any of the following operating systems: ■ AIX 6.1 and 7.1 ■ HP-UX 11i v3 ■ Linux: ■ RHEL 5 ■ RHEL 6 ■ SLES 10 ■ SLES 11 ■ Solaris 10 Review other details such as supported operating system levels and architecture for the supported operating systems. See the <i>Veritas Cluster Server Release Notes</i> or the <i>Veritas Storage Foundation High Availability Release Notes</i> for that platform.

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 14250. This is the default port that can be changed during a CP server configuration.

Symantec recommends that you configure multiple network paths to access a CP server. If a network path fails, CP server does not require a restart and continues to listen on one of the other available virtual IP addresses.

- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is being hosted.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops and network latency from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to the number of hops between the nodes.

For secure communication between the SF Oracle RAC cluster (application cluster) and the CP server, review the following support matrix:

	CP server in secure mode	CP server in non-secure mode
SF Oracle RAC cluster in secure mode	Yes	Yes
SF Oracle RAC cluster in non-secure mode	Yes	Yes
CP server cluster in secure mode	Yes	No
CP server cluster in non-secure mode	No	Yes

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

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

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

Supported database software

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

The following database versions are supported:

- Oracle RAC 11g Release 2

Note: If you are running SLES 10 SP4, install the Oracle patch 12311357.

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

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

Support for minor database versions is also documented in the afore-mentioned Technical Support TechNote.

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

Supported replication technologies for global clusters

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

Hardware-based replication	■ EMC SRDF ■ Hitachi TrueCopy ■ IBM Metro Mirror ■ IBM SAN Volume Controller (SVC) ■ EMC MirrorView
Software-based replication	■ Veritas Volume Replicator ■ Oracle Data Guard

Discovering product versions and various requirement information

Symantec provides several methods to check the Veritas product you have installed, plus various requirement information.

You can check the existing product versions using the `installer` command with the `-version` option before or after you install. After you have installed the current version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find version information.

Information the `version` option or the `showversion` script discovers on systems includes the following:

- The installed version of all released Storage Foundation and High Availability Suite of products
- The required RPMs or patches (if applicable) that are missing
- The available updates (including patches or hotfixes) from Symantec Operations Readiness Tools (SORT) for the installed products

To run the version checker

- 1 Mount the media.
- 2 Start the installer with the `-version` option.

```
# ./installer -version system1 system2
```

Planning to install SF Oracle RAC

This chapter includes the following topics:

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

Planning your network configuration

The following practices are recommended for a resilient network setup:

- Configure the private cluster interconnect over multiple dedicated gigabit Ethernet links. All single point of failures such as network interface cards (NIC), switches, and interconnects should be eliminated.
- The NICs used for the private cluster interconnect should have the same characteristics regarding speed, MTU, and full duplex on all nodes. Do not allow the NICs and switch ports to auto-negotiate speed.
- Configure non-routable IP addresses for private cluster interconnects.
- The default value for LLT peer inactivity timeout is 16 seconds. The value should be set based on service availability requirements and the propagation delay

between the cluster nodes in case of campus cluster setup. The LLT peer inactivity timeout value indicates the interval after which SF Oracle RAC on one node declares the other node in the cluster dead, if there is no network communication (heartbeat) from that node.

The default value for the CSS miss-count in case of SF Oracle RAC is 600 seconds. The value of this parameter is much higher than the LLT peer inactivity timeout so that the two clusterwares, VCS and Oracle Clusterware, do not interfere with each other's decisions on which nodes should remain in the cluster in the event of network split-brain. Veritas I/O fencing is allowed to decide on the surviving nodes first, followed by Oracle Clusterware. The CSS miss-count value indicates the amount of time Oracle Clusterware waits before evicting another node from the cluster, when it fails to respond across the interconnect. For more information, see the Oracle Metalink document: 782148.1

Planning the public network configuration for Oracle RAC

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

Additionally, you need a Single Client Access Name (SCAN) registered in Enterprise DNS that resolves to three IP addresses (recommended).

Oracle Clusterware/Grid Infrastructure manages the virtual IP addresses.

Planning the private network configuration for Oracle RAC

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

You must use UDP IPC for the database cache fusion traffic. The Oracle UDP IPC protocol requires an IP address. Depending on your deployment needs, this IP address may be a dedicated IP address or one that is shared with Oracle Clusterware.

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

The following practices provide a resilient private network setup:

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

In an SF Oracle RAC cluster, the Oracle Clusterware heartbeat link **MUST** be configured as an LLT link. If Oracle Clusterware and LLT use different links for their communication, then the membership change between VCS and Oracle Clusterware is not coordinated correctly. For example, if only the Oracle Clusterware links are down, Oracle Clusterware kills one set of nodes after the expiry of the `css-miscount` interval and initiates the Oracle Clusterware and database recovery, even before CVM and CFS detect the node failures. This uncoordinated recovery may cause data corruption.

- Oracle Clusterware interconnects need to be protected against NIC failures and link failures. The PrivNIC or MultiPrivNIC agent can be used to protect against NIC failures and link failures, if multiple links are available. Even if link aggregation solutions in the form of bonded NICs are implemented, the PrivNIC or MultiPrivNIC agent can be used to provide additional protection against the failure of the aggregated link by failing over to available alternate links. These alternate links can be simple NIC interfaces or bonded NICs.

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

See “[High availability solutions for Oracle RAC private network](#)” on page 52.

- Configure Oracle Cache Fusion traffic to take place through the private network. Symantec also recommends that all UDP cache-fusion links be LLT links.

Note: The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2. For more information, see

<http://www.symantec.com/business/support/index?page=content&id=TECH145261>

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

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

High availability solutions for Oracle RAC private network

[Table 3-1](#) lists the high availability solutions that you may adopt for your private network.

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

Options	Description
Using link aggregation/ NIC bonding for Oracle Clusterware	Use a native NIC bonding solution to provide redundancy, in case of NIC failures. Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link. When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership: ■ Configure an additional NIC under LLT in addition to the bonded NIC. ■ Add the following line in the <code>/etc/llttab</code> file: <pre>set-dbg-minlinks 2</pre>
Using PrivNIC/MultiPrivNIC agents	Use the PrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability using multiple network interfaces. Use the MultiPrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces. Note: See the following Technote regarding co-existence of PrivNIC and MultiPrivNIC agents with Oracle RAC: http://www.symantec.com/business/support/index?page=content&id=TECH145261 For more deployment scenarios that illustrate the use of PrivNIC/MultiPrivNIC deployments, see the appendix "SF Oracle RAC deployment scenarios" in this document.

Planning the storage

SF Oracle RAC provides the following options for shared storage:

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

The following recommendations ensure better performance and availability of storage.

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

Planning the storage for SF Oracle RAC

[Table 3-2](#) lists the type of storage required for SF Oracle RAC.

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

SF Oracle RAC files	Type of storage
SF Oracle RAC binaries	Local
SF Oracle RAC fencing coordinator disks	Shared
SF Oracle RAC database storage management repository	Shared

Planning the storage for Oracle RAC

Review the storage options and guidelines for Oracle RAC:


- Storage options for OCR and voting disk
See [“Planning the storage for OCR and voting disk”](#) on page 55.
- Storage options for the Oracle RAC installation directories (ORACLE_BASE, CRS_HOME or GRID_HOME (depending on Oracle RAC version), and ORACLE_HOME)
See [“Planning the storage for Oracle RAC binaries and data files”](#) on page 56.

Planning the storage for OCR and voting disk

You can place the OCR and voting disk information on a clustered file system or on ASM disk groups created using CVM raw volumes.

[Figure 3-1](#) illustrates the options for storing OCR and voting disk information.

Figure 3-1 OCR and voting disk storage options



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

For both option 1 and option 2:

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

Note: Retain the default disk detach policy setting (`global`) for OCR and voting disk.

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

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

Planning the storage for Oracle RAC binaries and data files

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

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

Consider the following points while planning the installation:

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

[Table 3-3](#) lists the type of storage for Oracle RAC binaries and data files.

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

Oracle RAC files	Type of storage
Oracle base	Local
Oracle Clusterware/Grid Infrastructure binaries	Local Placing the Oracle Grid Infrastructure binaries on local disks enables rolling upgrade of the cluster.
Oracle database binaries	Local Placing the Oracle database binaries on local disks enables rolling upgrade of the cluster.

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

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

Planning for Oracle ASM over CVM

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

Supported by ASM	ASM provides storage for data files, control files, Oracle Cluster Registry devices (OCR), voting disk, online redo logs and archive log files, and backup files.
Not supported by ASM	ASM does not support Oracle binaries, trace files, alert logs, export files, tar files, core files, and application binaries.

The following practices offer high availability and better performance:

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

Recovery Area, archive logs, and a second set of redo logs and control files on the second disk group.

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

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

Planning volume layout

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

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

Planning file system design

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

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

About planning to configure I/O fencing

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


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

The coordination points in server-based fencing can include only CP servers or a mix of CP servers and coordinator disks.

Note: Irrespective of whether you use coordinator disks or CP servers, SF Oracle RAC requires at least 3 coordination points.

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

Figure 3-2 Workflow to configure I/O fencing



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

Using the installsfrac program	See “Setting up disk-based I/O fencing using installsfrac program” on page 140. See “Setting up server-based I/O fencing using installsfrac program” on page 157.
Using the Web-based installer	See “Configuring SF Oracle RAC for data integrity using the Web-based installer” on page 155.
Using response files	See “Response file variables to configure disk-based I/O fencing” on page 414. See “Response file variables to configure server-based I/O fencing” on page 415. See “Configuring I/O fencing using response files” on page 386.
Manually editing configuration files	See “Setting up disk-based I/O fencing manually” on page 149. See “Setting up server-based I/O fencing manually” on page 157.


You can also migrate from one I/O fencing configuration to another.

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

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

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


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



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

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

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



Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
 See Figure 3-5 on page 63.
- Multiple application clusters use a single CP server and multiple pairs of coordinator disks (two) as their coordination points
 See Figure 3-6 on page 64.

Although the recommended CP server configurations use three coordination points, you can use more than three coordination points for I/O fencing. Ensure that the total number of CP servers you use is an odd number. In a configuration where multiple application clusters share a common set of CP server coordination points,

the application cluster as well as the CP server use a Universally Unique Identifier (UUID) to uniquely identify an application cluster.

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

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





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

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



See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 747.

Planning for cluster management

[Table 3-4](#) lists the various agents supported in SF Oracle RAC installations for effective cluster management.

Table 3-4 List of agents

Agent	Description
VCS agent for Oracle	Oracle database management The VCS Oracle agent is recommended for managing Oracle databases. VCS controls the Oracle database in this configuration. The configuration without VCS Oracle agent may be used only in a single database setup.
VCS agent for ASMDG	Oracle ASM disk group monitoring The ASMDG agent monitors the Oracle ASM disk groups.

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

Agent	Description
VCS agents for CVM	Volume management An SF Oracle RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource. You must configure the CVMVolDg agent for each shared disk group.
VCS agents for CFS	File system management If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.
CSSD agent (Application agent for Oracle Clusterware)	Oracle Clusterware management The CSSD agent starts, stops, and monitors Oracle Clusterware. It ensures that the OCR, the voting disk and the private IP address resources required by Oracle Clusterware are online before Oracle Clusterware starts. Note: It is mandatory to use CSSD agent in SF Oracle RAC installations to ensure adequate handling of inter-dependencies and thereby prevent the premature startup of Oracle Clusterware.
PrivNIC agent	High availability for a private IP address The PrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability using multiple network interfaces.
MultiPrivNIC agent	High availability for multiple private IP addresses The MultiPrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces.
CRSResource agent	Oracle Clusterware resource monitoring. The CRSResource agent monitors the Oracle Clusterware resources such as the virtual IP address, listener, and the Oracle database instance. It provides an alternative mechanism for monitoring the Oracle database in the absence of the VCS Oracle agent. It is useful in scenarios where the database is not managed by VCS and the applications need to be started using VCS after Oracle Clusterware starts the database.

Note: Make sure that you do not configure the following Oracle resources under VCS: SCAN IP, virtual IP, and listener.

Planning for disaster recovery

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

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

Planning a campus cluster setup

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

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

- Campus cluster sites are typically connected using a redundant high-capacity network that provides access to storage and private network communication between the cluster nodes. A single DWDM link can be used for both storage and private network communication.
- Tag the disks or enclosures that belong to a site with the corresponding VxVM site name. VxVM allocates storage from the correct site when creating or resizing a volume and when changing a volume's layout if the disks in the VxVM disk group that contain the volume are tagged with the site name.
- Tag each host with the corresponding VxVM site name. Make sure the read policy of the volumes is set to `SITEREAD`. This setting ensures that the reads on the volumes are satisfied from the local site's plex.
- Turn on the `allsites` attribute for all volumes that have data required by the application, to make sure they are evenly mirrored. Each site must have at least one mirror of all volumes hosting application data, including the FlashSnap log volume.
- Turn on the `siteconsistent` attribute for the disk groups and the volumes to enable site-aware plex detaches. Snapshot volumes need not be site-consistent.
- In the case of a two-site campus cluster, place the third coordinator disk on the third site. You may use iSCSI disk on the third site as an alternative to Dark Fiber connected FC-SAN or a Coordination Point Server (CPS), as a third coordination point.

- Make sure that a DCO log version 20 or higher is attached to the volumes to enable Fast Resync operations.
- Set the CVM disk detach policy as `global` for all disk groups containing data volumes as well as disk groups containing OCR and voting disk.

Planning a global cluster setup

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

Global clustering involves two steps:

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

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

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

Data replication considerations

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

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

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

Licensing SF Oracle RAC

This chapter includes the following topics:

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

About Veritas SFHA Solutions product licensing

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

www.symantec.com/techsupp/

During the installation, you can choose to either:

- Install a license key for the product and features that you want to install.
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server, such as Veritas Operations Manager (VOM). If you do not comply with the above terms,

continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

For more information about keyless licensing, see the following URL:

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

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

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

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

See [About Veritas Storage Foundation and High Availability Solutions 6.0](#)

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.

See [“Setting or changing the Veritas SFHA Solutions product level for keyless licensing”](#) on page 72.

See the `vxkeyless(1m)` manual page.

- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.

See [“Installing Veritas SFHA Solutions product license keys”](#) on page 73.

See the `vxlicinst(1m)` manual page.

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

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

About SF Oracle RAC licenses

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

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

License	Description	Features enabled
SFRACENT	SF Oracle RAC Enterprise Edition	The license enables the following features: ■ Veritas Volume Manager ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services
SFRACENT_VR	SF Oracle RAC Enterprise Edition with VR (Veritas Replicator)	The license enables the following features: ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services
SFRACENT_VFR	SF Oracle RAC Enterprise Edition with VFR (Veritas File Replicator)	The license enables the following features: ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services

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

License	Description	Features enabled
SFRACENT_GCO	SF Oracle RAC Enterprise Edition with GCO (Global Cluster Option)	The license enables the following features: ■ Veritas Volume Manager ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services
SFRACENT_VR_GCO	SF Oracle RAC Enterprise Edition with VR and GCO	The license enables the following features: ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services
SFRACENT_VFR_GCO	SF Oracle RAC Enterprise Edition with VFR and GCO	The license enables the following features: ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas Storage and Availability Management Tools for Oracle databases ■ Veritas Extension for ODM ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services

Setting or changing the Veritas SFHA Solutions product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 Show your current working directory:

```
# pwd
```

Output resembles:

```
/opt/VRTSvlic/bin
```

- 2 View the current setting for the product level.

```
# ./vxkeyless -v display
```

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

```
# ./vxkeyless set prod_levels
```

where *prod_levels* is a comma-separated list of keywords. The keywords are the product levels as shown by the output of step 3.

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas SFHA Solutions product license keys

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

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

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

To install a new license

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

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

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

Installation and configuration of SF Oracle RAC

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

Preparing to install SF Oracle RAC

This chapter includes the following topics:

- [Setting the umask before installation](#)
- [Synchronizing time settings on cluster nodes](#)
- [Mounting the product disc](#)
- [Setting the environment variables for SF Oracle RAC](#)
- [Configuring the I/O scheduler](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Verifying the systems before installation](#)

Setting the umask before installation

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

```
# umask 0022
```

Synchronizing time settings on cluster nodes

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

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

Use the following command on each system to synchronize the time settings with the NTP server:

- On RHEL 5, use the `rdate` command:

```
# rdate -s timesrv
```

where **timesrv** is the name of the NTP server

- On SLES 10 and SLES 11, use the `yast` command:

```
# yast ntp-client
```

For instructions, see the operating system documentation.

Mounting the product disc

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

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

To mount the product disc

- 1 Log in as the superuser to a cluster node or a remote node in the same subnet as the cluster nodes.
- 2 Insert the product disc with the SF Oracle RAC software into a drive that is connected to the system.

The disc is automatically mounted.

- 3 If the disc does not automatically mount, then enter:

```
# mount -o ro /dev/dvdrom /dvd_mount
```

Setting the environment variables for SF Oracle RAC

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

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

```
For sh, ksh, or bash    # export MANPATH=$MANPATH:\n                        /opt/VRTS/man
```

```
For csh                 # setenv MANPATH $MANPATH:/opt/VRTS/man
```

Some terminal programs may display garbage characters when you view the man pages. Set the environment variable `LC_ALL=C` to resolve this issue.

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

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

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

Configuring the I/O scheduler

Symantec recommends using the Linux 'deadline' I/O scheduler for database workloads. Configure your system to boot with the 'elevator=deadline' argument to select the 'deadline' scheduler.

For information on configuring the 'deadline' scheduler for your Linux distribution, see the operating system documentation.

To determine whether a system uses the deadline scheduler, look for "elevator=deadline" in `/proc/cmdline`.

To configure a system to use the deadline scheduler

- 1 Include the `elevator=deadline` parameter in the boot arguments of the GRUB or ELILO configuration file. The location of the appropriate configuration file depends on the system's architecture and Linux distribution. For x86_64, the configuration file is `/boot/grub/menu.lst`
 - For GRUB configuration files, add the `elevator=deadline` parameter to the kernel command. For example, change:

```
title Red Hat Enterprise Linux Server (2.6.18-92.el5)
  root (hd1,1)
  kernel /boot/vmlinuz-2.6.18-92.el5 ro root=/dev/sdb2
  initrd /boot/initrd-2.6.18-92.el5.img
```

To

```
title Red Hat Enterprise Linux Server (2.6.18-92.el5)
  root (hd1,1)
  kernel /boot/vmlinuz-2.6.18-92.el5 ro root=/dev/sdb2 elevator=deadline
  initrd /boot/initrd-2.6.18-92.el5.img
```

- A setting for the elevator parameter is always included by SUSE in its ELILO and its GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.
- 2 Reboot the system once the appropriate file has been modified.

See the operating system documentation for more information on I/O schedulers.

Optimizing LLT media speed settings on private NICs

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

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

Guidelines for setting the media speed of the LLT interconnects

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

- Symantec recommends that you manually set the same media speed setting on each Ethernet card on each node.
If you use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.
- If you have hubs or switches for LLT interconnects, then set the hub or switch port to the same setting as used on the cards on each node.

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

Verifying the systems before installation

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

- Option 1: Run Symantec Operations Readiness Tools (SORT).
For information on downloading and running SORT:
<https://sort.symantec.com>
- Option 2: Run the `installsfrac` program with the `"-precheck"` option as follows:

Navigate to the directory that contains the `installsfrac` program.

The program is located in the `storage_foundation_for_oracle_rac` directory.

Start the preinstallation check:

```
# ./installsfrac -precheck node_1 node_2
```

where `node_1`, `node_2` are the names of the nodes in the cluster.

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

- Option 3: Run the Veritas Web-based installation program as follows:
 - Navigate to the directory that contains the Web-based installation program and start the installer:

```
# ./webinstaller start
```

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

- On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.
Select the product from the **Product** drop-down list, and click **Next**.
Enter the names of the systems that you want to verify.
The installer performs the precheck and displays the results.

Installing SF Oracle RAC

This chapter includes the following topics:

- [About installing SF Oracle RAC](#)
- [About installation and configuration methods](#)
- [Installing SF Oracle RAC using the Veritas script-based installation program](#)
- [Installing SF Oracle RAC using the Veritas Web-based installation program](#)
- [Installing SF Oracle RAC using Kickstart](#)
- [Installing Veritas Storage Foundation for Oracle RAC using yum](#)

About installing SF Oracle RAC

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

By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.

The product installer has improved ability to recover from failed installations. By default, the installer generates a response file that you can reuse or customize for silent installations on other systems.

About installation and configuration methods

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

Table 6-1 Installation and configuration methods

Method	Description
Interactive installation and configuration using the script-based installer Note: If you obtained SF Oracle RAC from an electronic download site, you must use the <code>installsfrac program</code> script instead of the <code>installer</code> script.	You can use one of the following script-based installers: ■ Common product installer script: <code>installer</code> The common product installer script provides a menu that simplifies the selection of installation and configuration options. ■ Product-specific installation script: <code>installsfrac program</code> ■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsfrac program</code> script is identical to specifying SF Oracle RAC from the <code>installer</code> script. Use this method to install or configure only SF Oracle RAC.
Silent installation using the response file	The response file automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. You can use the script-based installers with the response file to install silently on one or more systems. See "About response files" on page 373.
Web-based installer	The Web-based installer provides an interface to manage the installation and configuration from a remote site using a standard Web browser. <code>webinstaller</code> See "Installing SF Oracle RAC using the Veritas Web-based installation program" on page 85.
Kickstart	The Kickstart enables the system administrator automatically to install systems based on predefined customized configurations. Kickstart is an automatic operating system installation method available for the Red Hat Linux operating system.

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

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

See “[SF Oracle RAC worksheet](#)” on page 606.

During the installation, the installer performs the following tasks:

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

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

- Installs the SF Oracle RAC 6.0 RPMs.

For Storage Foundation for Oracle RAC:

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

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

To install SF Oracle RAC

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

```
SF Oracle   Run the program:
RAC installer # ./installsfrac galaxy nebula
```

Common
product
installer

Navigate to the directory that contains the installation program.

Run the program:

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

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

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

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

- 4 Set up the systems so that commands between systems execute without prompting for passwords or confirmations.

Would you like the installer to setup ssh or rsh communication automatically between the systems?

Superuser passwords for the systems will be asked. [y,n,q] (y)

Enter the superuser password for system vcslx613:

- 1) Setup ssh between the systems
- 2) Setup rsh between the systems
- b) Back to previous menu

Select the communication method [1-2,b,q,?] (1)

- 5 If you had quit the installer in the process of an active installation, the installer discovers that installer process and provides the option of resuming the installation or starting a new installation. Provide a suitable response.

The installer has discovered an existing installer process.

The process exited while performing configure of SF Oracle RAC on galaxy.

Do you want to resume this process? [y,n,q,?] (y) **n**

- 6 Enter **y** to agree to the End User License Agreement (EULA).
- 7 Select the type of installation—Minimal, Recommended, All. Each option displays the disk space that is required for installation.

Symantec recommends you to choose the option **Install all RPMs**.

- 1) Install minimal required RPMs
- 2) Install recommended RPMs
- 3) Install all RPMs
- 4) Display RPMs to be installed for each option

Select the RPMs to be installed on all systems?

[1-4,q,?] (2) **3**

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

The installer installs the SF Oracle RAC RPMs and patches.

8 Select the appropriate license option.

1) Enter a valid license key

2) Enable keyless licensing and complete system licensing later

How would you like to license the systems? [1-2,q]

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

Enter a SF Oracle RAC license key:

xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-x

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

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

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

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

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

Would you like to enable the
Veritas Volume Replicator? [y,n,q] (n)
Would you like to enable the
Veritas File Replicator? [y,n,q] (n)
Would you like to enable the
Global Cluster Option? [y,n,q] (n)

The installer registers the license.

- 9 Verify that the installation process completed successfully. Review the output at the end of the installation and note the location of the summary and log files for future reference.
- 10 Would you like to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] (y) **y**
- 11
- 12 Enter **y** to configure SF Oracle RAC:

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

 If you plan to set up server-based fencing, make sure that you complete the preparatory tasks before starting the SF Oracle RAC configuration. For instructions, see the chapter "Preparing to configure SF Oracle RAC" in this document.
- 13 Enter **y** if you want to send the installation information to Symantec.

 Would you like to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] (y) **y**
- 14 Enter **y** if you want to view the summary file.

 Would you like to view the summary file? [y,n,q] (n) **y**

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

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

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

- Review the system configuration requirements for a Web-based installation. See [“Before using the Veritas Web-based installer”](#) on page 86.
- Start the Veritas Web installer. See [“Starting the Veritas Web-based installer”](#) on page 86.
- Install SF Oracle RAC. See [“Installing products with the Veritas Web-based installer”](#) on page 87.

During the installation, the installer performs the following tasks:

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

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

- Installs the SF Oracle RAC 6.0 RPMs.

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

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 6-2 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Veritas products.	Must be a supported platform for SF Oracle RAC 6.0.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and must be at one of the supported operating system update levels.
Administrative system	The system where you run the Web browser to perform the installation.	Must have a Web browser. Supported browsers: ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1wid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

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

The default listening port is 14172. If you have a firewall that blocks port 14172, use the `-port` option to use a free port instead.

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

```
Secure Connection Failed
```

Obtain a security exception for your browser.

When prompted, enter `root` and root's password of the installation server.

- 5 Log in as superuser.

Installing products with the Veritas Web-based installer

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

To install SF Oracle RAC

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 2 On the Select a task and product page, select the installation from the **Task** drop-down list.
- 3 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list, and click **Next**.
- 4 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 5 Choose minimal, recommended, or all packages. Click **Next**.
- 6 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.

- 7 After the validation completes successfully, click **Next** to install SF Oracle RAC on the selected system.

- 8 After the installation completes, you must choose your licensing method.

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

- Keyless licensing

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

For more information, go to the following website:

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

Complete the following information:

Choose whether you want to enable Global Cluster option.

Click **Register**.

- Enter license key

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

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

For instructions, see the chapter, "Preparing to Configure SF Oracle RAC."

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

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

If the installer prompts you to reboot the systems, do so.

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

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

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

Installing SF Oracle RAC using Kickstart

You can install SF Oracle RAC using Kickstart. Kickstart is supported for Red Hat Enterprise Linux (RHEL5) and Red Hat Enterprise Linux 6 (RHEL6).

To install SF Oracle RAC using Kickstart

- 1 Create a directory for the Kickstart configuration files.

```
# mkdir /kickstart_files/
```

- 2 Generate the Kickstart configuration files. The configuration files have the extension `.ks`. Do one of the following:

- To generate configuration files, enter the following command:

```
# ./installer -kickstart /kickstart_files/
```

The system lists the files.

- To only generate the configuration file for SFRAC, enter the following command:

```
# ./installsfrac -kickstart /kickstart_files/
```

The command output includes the following:

```
The kickstart script for SFRAC60 is generated at
/kickstart_files/kickstart_sfrac60.ks
```

- 3 Setup an NFS exported location which the Kickstart client can access. For example, if `/nfs_mount_kickstart` is the directory which has been NFS exported, the NFS exported location may look similar to the following:

```
# cat /etc/exports
/nfs_mount_kickstart * (rw,sync,no_root_squash)
```

- 4 Copy the `rpms` directory from the installation media to the NFS location.
- 5 Verify the contents of the directory.

```
# ls /nfs_mount_kickstart/
```

- 6 In the SF Oracle RAC Kickstart configuration file, modify the `BUILDSRC` variable to point to the actual NFS location. The variable has the following format:

```
BUILDSRC="hostname_or_ip:/nfs_mount_kickstart"
```

- 7 Append the entire modified contents of the Kickstart configuration file to the operating system `ks.cfg` file.
- 8 Launch the Kickstart installation for the operating system.

- 9 After the operating system installation is complete, check the file `/var/tmp/kickstart.log` for any errors related to the installation of Veritas RPMs and Veritas product installer scripts.

- 10 Verify that all the product RPMs have been installed. Enter the following command:

```
# rpm -qa | grep -i vrts
```

- 11 If you do not find any installation issues or errors, configure the product stack. Enter the following command:

```
# /opt/VRTS/install/installsfrac -configure node1 node2
```

- 12 Verify that all the configured Ilt links and gab ports have started. Enter the following:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen    b66f01 membership 01
Port b gen    b66f03 membership 01
Port d gen    b66f07 membership 01
Port f gen    b66f0d membership 01
Port h gen    b66f05 membership 01
Port v gen    b66f09 membership 01
Port u gen    b66f08 membership 01
Port w gen    b66f0b membership 01
Port y gen    b66f0c membership 01
```

13 Verify that the product is configured properly. Enter the following:

```
# hastatus -summ
-- SYSTEM STATE
-- System                               State                               Frozen

A   swlxl3                             RUNNING                             0
A   swlxl4                             RUNNING                             0

-- GROUP STATE
-- Group                               System                               Probed                               AutoDisabled                               State

B   cvm                                swlxl3                             Y                                    N                                    ONLINE
B   cvm                                swlxl4                             Y                                    N                                    ONLINE
```

Note that the cvm service group comes online when the SFCFSHA stack is configured.

14 If you configure the node in a secured mode, verify the VxSS service group status. For example:

```
# haclus -value SecureClus
1
```

Sample Kickstart configuration file

The following is a sample RedHat Enterprise Linux 5 (RHEL5) Kickstart configuration file.

```
# The packages that follow are required. These packages are automatically installed
# from the operating system's installation media.
```

```
%packages
libattr.i386
libacl.i386
```

```
%post --nochroot
# Add any necessary scripts or commands here.
# This generated kickstart file is only for
# the installation of products in the installation media.
```

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

```
#
# Notice:
# * Modify the BUILDSRC below according to your environment
# * The location that you specify in BUILDSRC should let the
#   Kickstart Server be NFS accessible.
# * Copy the rpms directory from installation media to the
#   BUILDSRC location.
#

BUILDSRC="<hostname_or_ipv4address_or_ipv6address>:/path/to/rpms"

#
# Notice:
# * You do not have to change the following scripts.
#

# Define the path variables.
ROOT=/mnt/sysimage
BUILDDIR="${ROOT}/build"
RPMDIR="${BUILDDIR}/rpms"

# Define the log path.
KSLOG="${ROOT}/var/tmp/kickstart.log"

echo "==== Executing kickstart post section: ====" >> ${KSLOG}

mkdir -p ${BUILDDIR}
mount -t nfs -o nolock,vers=3 ${BUILDSRC} ${BUILDDIR} >> ${KSLOG} 2>&1

# Install the RPMs in the following order.
for RPM in VRTSvlic VRTSperl VRTSspt VRTSvxvm VRTSaslapm VRTSob
VRTSlmconv VRTSsfmh VRTSvxfs VRTSfsadv VRTSfssdk VRTSllt VRTSgab
VRTSvxfen VRTSamf VRTSvcs VRTScps VRTSvcsg VRTSvcscr VRTSvcsea
VRTSdbed VRTSglm VRTScavf VRTSgms VRTSodm VRTSdbac VRTSsfcp160

do
    echo "Installing package -- $RPM" >> ${KSLOG}
    rpm -U -v --root ${ROOT} ${RPMDIR}/${RPM}-* >> ${KSLOG} 2>&1
done

umount ${BUILDDIR}
```

```
${ROOT}/opt/VRTS/install/bin/UXRT60/add_install_scripts >> ${KSLOG} 2>&1

echo "==== Completed kickstart file =====" >> ${KSLOG}

exit 0
```

Installing Veritas Storage Foundation for Oracle RAC using yum

You can install SF Oracle RAC using yum. yum is supported for Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6.

To install SF Oracle RAC using yum

- 1 Run the `installsfrac -pkginfo` command to get SF Oracle RAC RPMs.

```
# ./installsfrac -pkginfo
```

- 2 Add the SF Oracle RAC RPMs into the yum repository. You can add SF Oracle RAC RPMs into either a new repository or an existing repository with other RPMs. Use the `createrepo` command to create or update the repository. The operating system package `createrepo-ver-rel.noarch.rpm` provides the command.

- **To create the new repository */path/to/new/repository/* for SF Oracle RAC RPMs**

1. Create an empty directory, for example: */path/to/new/repository*. The yum client systems should be able to access the directory with the HTTP, FTP, or file protocols.

```
# rm -rf /path/to/new/repository
# mkdir -p /path/to/new/repository
```

2. Copy all the SF Oracle RAC RPMs into */path/to/new/repository/*.

```
# cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcp160-* \
/path/to/new/repository
```

3. Use the `createrepo` command to create the repository.

```
# /usr/bin/createrepo /path/to/new/repository
```

Output resembles:

```
27/27 - VRTSsfcp60-6.0.000.000-GA_GENERIC.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
```

4. The metadata for this repository is created in `/path/to/new/repository/repodata`.

■ To use an existing repository in `/path/to/existing/repository/` for SF Oracle RAC RPMs

1. Copy all the SF Oracle RAC RPMs into `/path/to/existing/repository/`. The yum client systems should be able to access the directory with the HTTP, FTP, or file protocols.

```
# cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcp60-* \
/path/to/existing/repository
```

2. Use the `createrepo` command with the `--update` option to update the repository's metadata.

```
# createrepo --update /path/to/existing/repository
```

Output resembles:

```
27/27 * VRTSsfcp60-6.0.000.000-GA_GENERIC.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
```

3. The metadata in `/path/to/existing/repository/repodata` is updated for the newly added RPMs.

■ To create a package group for SF Oracle RAC RPMs when the repository is created or updated (optional)

1. Create an XML file, which you can name SF Oracle RAC_group.xml in the repository directory. In the file specify the name, the id, the package list, and other information for the group. An example of this XML file for SF Oracle RAC is:

```
# cat SF Oracle RAC_group.xml
<comps>
  <group>
    <id>SF Oracle RAC60</id>
    <name>SF Oracle RAC60</name>
    <default>true</default>
    <description>RPMs of SF Oracle RAC 6.0</description>
    <uservisible>true</uservisible>
    <packagelist>
      <packagereq type="default">VRTSvlic</packagereq>
      <packagereq type="default">VRTSperl</packagereq>
      ... [other RPMs for SF Oracle RAC]
      <packagereq type="default">VRTSsfcp60</packagereq>
    </packagelist>
  </group>
</comps>
```

2. Create the group when the repository is created or updated. You can generate this XML file using the installer with the option `-yumgroupxml`.

```
# createrepo -g SF Oracle RAC_group.xml /path/to/new/repository/
```

Or

```
# createrepo -g SF Oracle RAC_group.xml --update /path/to/existing\
/repository/
```

Refer to the *Red Hat Enterprise Linux Deployment Guide* for more information on yum repository configuration.

3. Configure a yum repository on a client system.

- Create a `.repo` file under `/etc/yum.repos.d/`. An example of this `.repo` file for SF Oracle RAC is:

```
# cat /etc/yum.repos.d/SF Oracle RAC.repo
[repo-SF Oracle RAC]
name=Repository for SF Oracle RAC
baseurl=file:///path/to/repository/
enabled=1
gpgcheck=0
```

The values for the `baseurl` attribute can start with `http://`, `ftp://`, or `file://`. The URL you choose needs to be able to access the `repodata` directory. It also needs to access all the SF Oracle RAC RPMs in the repository that you create or update.

- Check the yum configuration. List SF Oracle RAC RPMs.

```
# yum list 'VRTS*'
Available Packages
VRTSperl.x86_64          5.10.0.7-RHEL5.3          repo-SF Oracle RAC
VRTSsfcp160.noarch      6.0.000.000-GA_GENERIC    repo-SF Oracle RAC
VRTSvlic.x86_64         3.02.52.000-0             repo-SF Oracle RAC
...
```

The SF Oracle RAC RPMs may not be visible immediately if:

- the repository was visited before the SF Oracle RAC RPMs were added, and
- the local cache of its metadata has not expired.

To eliminate the local cache of the repositories' metadata and get the latest information from the specified `baseurl`, run the following commands:

```
# yum clean expire-cache
# yum list 'VRTS*'
```

Refer to the *Red Hat Enterprise Linux Deployment Guide* for more information on yum repository configuration.

4 Install the RPMs on the target systems.

- **To install all the RPMs**

1. Specify each RPM name as its yum equivalent. For example:

```
# yum install VRTSvlic VRTSperl ... VRTSsfcp160
```

2. Specify all of SF Oracle RAC's RPMs by its package glob. For example:

```
# yum install 'VRTS*'
```


3. Specify the group name if a group is configured for SF Oracle RAC's RPMs. In this example, the group name is *SF Oracle RAC60*:

```
# yum install @SF Oracle RAC60
```

Or

```
# yum groupinstall SF Oracle RAC60
```

■ To install one RPM at a time

1. Run the `installsfrac -pkginfo` command to determine package installation order.
2. Use the same order as the output from the `installsfrac -pkginfo` command:

```
# yum install VRTSvlic
# yum install VRTSperl
...
# yum install VRTSsfcp60
```

- 5 After you install all the RPMs, use the `/opt/VRTS/install/installsfrac` script to license, configure, and start the product.

If the `VRTSsfcp60` package is installed before you use `yum` to install SF Oracle RAC, this package is not reinstalled. The `/opt/VRTS/install/installsfrac` script may not be created. If the script is not created, use the `/opt/VRTS/install/bin/UXRT60/add_install_scripts` script to create the `installsfrac` or `uninstallsfrac` scripts after all the other SF Oracle RAC RPMs are installed.

```
# /opt/VRTS/install/bin/UXRT60/add_install_scripts
Creating install/uninstall scripts for installed products
Creating /opt/VRTS/install/installsfrac for UXRT60
Creating /opt/VRTS/install/uninstallsfrac for UXRT60
```

Preparing to configure SF Oracle RAC

This chapter includes the following topics:

- [Setting up the CP server](#)

Setting up the CP server

[Table 7-1](#) lists the tasks to set up the CP server for server-based I/O fencing.

Table 7-1 Tasks to set up CP server for server-based I/O fencing

Task	Reference
Plan your CP server setup	See “Planning your CP server setup” on page 99.
Install the CP server	See “Installing the CP server using the installer” on page 100.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 101.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 101.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 102. See “Configuring the CP server manually” on page 111.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 113.

Planning your CP server setup

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

To plan your CP server setup

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

Symantec recommends hosting the CP server on an SFHA cluster to make the CP server highly available.

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

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

Symantec recommends using at least three coordination points.

- 3 Decide whether you want to configure the CP server cluster in secure mode.

Symantec recommends configuring the CP server cluster in secure mode to secure the communication between the CP server and its clients (SF Oracle RAC clusters). It also secures the HAD communication on the CP server cluster.

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

See [“CP server requirements”](#) on page 44.

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

- Name for the CP server
The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
- Port number for the CP server
Allocate a TCP/IP port for use by the CP server.
Valid port range is between 49152 and 65535. The default port number is 14250.
- Virtual IP address, network interface, netmask, and networkhosts for the CP server
You can configure multiple virtual IP addresses for the CP server.

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

Installing the CP server using the installer

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

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

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

CP server setup uses a single system

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

During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs.

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

Proceed to configure the CP server.

See “[Configuring the CP server using the configuration utility](#)” on page 102.

See “[Configuring the CP server manually](#)” on page 111.

CP server setup uses multiple systems

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

Meet the following requirements for CP server:

- During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs.
- During configuration, configure disk-based fencing (scsi3 mode).

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

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

Configuring the CP server cluster in secure mode

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

This step secures the HAD communication on the CP server cluster.

Note: If you already configured the CP server cluster in secure mode during the VCS configuration, then skip this section.

To configure the CP server cluster in secure mode

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

If you have VCS installed on the CP server, run the following command:

```
# installvcs -security
```

If you have SFHA installed on the CP server, run the following command:

```
# installsfha -security
```

Setting up shared storage for the CP server database

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

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

To set up shared storage for the CP server database

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

For example:

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

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

For example:

```
# vxdg import cps_dg
```

3 Create a mirrored volume over the disk group.

For example:

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

4 Create a file system over the volume.

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

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

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

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

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

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

Configuring the CP server using the configuration utility

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

Perform one of the following procedures:

For CP servers on single-node VCS cluster: See [“To configure the CP server on a single-node VCS cluster”](#) on page 102.

For CP servers on an SFHA cluster: See [“To configure the CP server on an SFHA cluster”](#) on page 106.

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

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

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

- 3 Enter **1** at the prompt to configure CP server on a single-node VCS cluster.

The configuration utility then runs the following preconfiguration checks:

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

- 4 Enter the name of the CP server.

```
Enter the name of the CP Server: mycps1
```

- 5 Enter valid virtual IP addresses on which the CP server process should depend on:

- Enter the number of virtual IP addresses you want to configure:

```
Enter the number of virtual IP(s) to configure : 2
```

- Enter valid virtual IP addresses:

```
Enter a valid IP address for Virtual IP - 1 which the CP Server  
process should depend on : 10.209.83.85
```

```
Enter a valid IP address for Virtual IP - 2 which the CP Server  
process should depend on : 10.209.83.87
```

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

```
Enter a port number for virtual IP 10.209.83.85 in range [49152,  
65535], or press enter for default port (14250) :
```

```
Using default port: 14250
```

```
Enter a port number for virtual IP 10.209.83.87 in range  
[49152, 65535], or press enter for default port (14250) :
```

```
Using default port: 14250
```

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

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

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

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

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

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

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

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

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

- 9 Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```
-----
(a)CP Server Name: mycps1
(b)CP Server Virtual IP(s): 10.209.83.85 10.209.83.87
(c)CP Server Port(s): 14250 14250
(d)CP Server Security : 1
(e)CP Server Database Dir: /etc/VRTScps/db
-----
```

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

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

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

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

- 11** Enter the number of NIC resources that you want to configure. You must use a public NIC.

```
Enter how many NIC resources you want to configure [1 to 2]: 2
```

Answer the following questions for each NIC resource that you want to configure.

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

```
Enter a valid network interface for virtual IP 10.209.83.85
on mycps1.symantecexample.com: eth1
Enter a valid network interface for virtual IP 10.209.83.87
on mycps1.symantecexample.com: eth1
```

- 13** Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.85 [1 to 2] : 1
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.87 [1 to 2] : 2
```

- 14** Enter networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure
NIC resource to be online always.
Do you want to add NetworkHosts attribute for the NIC device
eth1 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC eth1 on system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

15 Enter the netmask for each virtual IP address. For example:

```
Enter the netmask for virtual IP 10.209.83.85 :  
255.255.252.0  
Enter the netmask for virtual IP 10.209.83.87 :  
255.255.252.0
```

16 After the configuration process has completed, a success message appears. For example:

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

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

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

For example:

```
# hagrp -state CPSSG
```

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

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 654.

To configure the CP server on an SFHA cluster

- 1 Verify that the VRTScps RPM is installed on each node.
- 2 Make sure that you have configured passwordless ssh or rsh on the CP server cluster nodes.

- 3 Run the CP server configuration script on any node in the cluster:

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

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

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

The configuration utility then runs the following preconfiguration checks:

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

- 5 Enter the name of the CP server.

```
Enter the name of the CP Server: mycps1
```

- 6 Enter valid virtual IP addresses on which the CP server process should depend on:

- Enter the number of virtual IP addresses you want to configure:

```
Enter the number of virtual IP(s) to configure : 2
```

- Enter valid virtual IP addresses:

```
Enter a valid IP address for Virtual IP - 1 which the CP Server  
process should depend on : 10.209.83.85
```

```
Enter a valid IP address for Virtual IP - 2 which the CP Server  
process should depend on : 10.209.83.87
```

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

Enter a port number for virtual IP 10.209.83.85 in range [49152, 65535], or press enter for default port (14250) :

Using default port: 14250

Enter a port number for virtual IP 10.209.83.87 in range [49152, 65535], or press enter for default port (14250) :

Using default port: 14250

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

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

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

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

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

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

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

Note: As the CP Server is being configured on SFHA cluster, the database should reside on shared storage with vxfs file system.

Please refer to documentation for information on setting up of shared storage for CP server database.

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

10 Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```
-----
(a)CP Server Name: mycps1
(b)CP Server Virtual IP(s): 10.209.83.85 10.209.83.87
(c)CP Server Port(s): 14250 14250
(d)CP Server Security : 1
(e)CP Server Database Dir: /etc/VRTScps/db
-----
```

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

11 The configuration utility proceeds with the configuration process, and creates a vxcps.conf configuration file on each node.

The following output is for one node:

```
Successfully generated the /etc/vxcps.conf
configuration file.
Successfully created directory /etc/VRTScps/db.
Creating mount point /etc/VRTScps/db on
mycps1.symantecexample.com.
Copying configuration file /etc/vxcps.conf to
mycps1.symantecexample.com
```

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

12 Enter the number of NIC resources that you want to configure. You must use a public NIC.

Enter how many NIC resources you want to configure [1 to 2]: 2

Answer the following questions for each NIC resource that you want to configure.

13 Confirm whether you use the same NIC name for the virtual IP on all the systems in the cluster.

Is the name of network interfaces for NIC resource - 1
same on all the systems?[y/n] : y

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

```
Enter a valid interface for virtual IP 10.209.83.85
on all the systems : eth1
```

- 15** Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.85 [1 to 2] : 1
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.87 [1 to 2] : 2
```

- 16** Enter networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure
NIC resource to be online always.
Do you want to add NetworkHosts attribute for the NIC device
eth1 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC eth1 on system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

- 17** Enter the netmask for each virtual IP address.

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

- 18** Enter the name of the disk group for the CP server database.

```
Enter the name of diskgroup for cps database :
cps_dg
```

- 19** Enter the name of the volume that is created on the above disk group.

```
Enter the name of volume created on diskgroup cps_dg :
cps_volume
```

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

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

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

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

For example:

```
# hagrp -state CPSSG
```

#Group	Attribute	System	Value
CPSSG	State	mycps1	ONLINE
CPSSG	State	mycps2	OFFLINE

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 654.

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

To manually configure the CP server

- 1 Stop VCS on each node in the CP server cluster using the following command:

```
# hastop -local
```

- 2 Edit the `main.cf` file to add the CPSSG service group on any node. Use the CPSSG service group in the `main.cf` as an example:

See [“Sample configuration files for CP server”](#) on page 654.

Customize the resources under the CPSSG service group as per your configuration.

- 3 Verify the `main.cf` file using the following command:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

If successfully verified, copy this `main.cf` to all other cluster nodes.

- 4 Create the `/etc/vxcps.conf` file using the sample configuration file provided at `/etc/vxcps/vxcps.conf.sample`.

Based on whether you have configured the CP server cluster in secure mode or not, do the following:

- For a CP server cluster which is configured in secure mode, edit the `/etc/vxcps.conf` file to set `security=1`.
- For a CP server cluster which is not configured in secure mode, edit the `/etc/vxcps.conf` file to set `security=0`.

Symantec recommends enabling security for communication between CP server and the application clusters.

- 5 Start VCS on all the cluster nodes.

```
# hastart
```

- 6 Verify that the CP server service group (CPSSG) is online.

```
# hagrps -state CPSSG
```

Output similar to the following appears:

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

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

- 1 Verify that the following configuration files are updated with the information you provided during the CP server configuration process:
 - `/etc/vxcps.conf` (CP server configuration file)
 - `/etc/VRTSvcs/conf/config/main.cf` (VCS configuration file)
 - `/etc/VRTSvcs/db` (default location for CP server database)
- 2 Run the `cpsadm` command to check if the `vxcpserv` process is listening on the configured Virtual IP.

```
# cpsadm -s cp_server -a ping_cps
```

where `cp_server` is the virtual IP address or the virtual hostname of the CP server.

Configuring SF Oracle RAC

This chapter includes the following topics:

- [About configuring SF Oracle RAC](#)
- [Configuring the SF Oracle RAC components using the script-based installer](#)
- [Configuring SF Oracle RAC using the Web-based installer](#)

About configuring SF Oracle RAC

You need to configure SF Oracle RAC when:

- You have completed installation of SF Oracle RAC on your systems.
- You want to reconfigure an existing SF Oracle RAC cluster.

Note: Before you reconfigure a cluster, make sure that you stop any running applications that use VxFS/CFS. Then, unmount the VxFS/CFS mounts.

SF Oracle RAC configuration involves the following high-level tasks:

- Starting the product installer (if you quit the installer after installation or want to reconfigure the cluster)
- Configuring the SF Oracle RAC components—VCS, CVM, and CFS
- Configuring the SF Oracle RAC clusters for data integrity

During the configuration process, the installer performs the following tasks:

- Verifies the cluster information.
- Stops SF Oracle RAC processes.
- Creates SF Oracle RAC configuration files.

- Starts SF Oracle RAC processes.
- Creates a new directory with a log file that contains any system commands executed, and their output, a response file that can be used with the `-responsefile` option of the installer, and a summary file that contains the output of the install scripts. The location of the files is indicated by the installer.

Configuring the SF Oracle RAC components using the script-based installer

Make sure that you have performed the necessary pre-configuration tasks if you want to configure the cluster in secure mode.

Start the `installsf` or `installer` program if you quit the installer after installation.

By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.

At the end of the configuration, the VCS, CVM, and CFS components are configured to provide a cluster-aware environment.

Note: If you want to reconfigure SF Oracle RAC, before you start the installer you must stop all the resources that are under VCS control using the `hasstop` command or the `hagrp -offline` command.

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

To configure the SF Oracle RAC components

- 1 Log in as the superuser on any of the nodes in the cluster.
- 2 Start the configuration program.

SF Oracle RAC installer	Run the program:
	<code># cd /opt/VRTS/install</code>
	<code># ./installsfrac -configure <i>galaxy nebula</i></code>

Common product installer	Run the program:
	<code># ./installer -configure <i>galaxy nebula</i></code>
	Choose Veritas Storage Foundation for Oracle RAC to configure SF Oracle RAC.

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

- 3 Enter **1** to select the option **Configure SF Oracle RAC sub-components**.

```

1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1)

```

- 4 If you had quit the installer in the process of an active configuration, the installer discovers that installer process and provides the option of resuming the configuration or starting a new configuration. Provide a suitable response.

```

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

```

- 5 Configure the Veritas Cluster Server component to set up the SF Oracle RAC cluster.

See [“Configuring the SF Oracle RAC cluster”](#) on page 117.

- 6 Add VCS users.
See [“Adding VCS users”](#) on page 127.
- 7 Configure SMTP email notification.
See [“Configuring SMTP email notification”](#) on page 128.
- 8 Configure SNMP trap notification.
See [“Configuring SNMP trap notification”](#) on page 130.
- 9 Configure global clusters, if you chose to enable GCO during the installation.
See [“Configuring global clusters”](#) on page 132.
- 10 Stop the SF Oracle RAC resources.
See [“Stopping and starting SF Oracle RAC processes”](#) on page 133.

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 or an advanced cluster. A basic cluster configuration requires the cluster name and ID and the private heartbeat links for LLT. The remaining configuration options presented by the installer are optional and may be used if you plan to configure an advanced cluster.

Refer to the *Veritas Cluster Server Installation Guide* for more information.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

- 1 Review the configuration instructions that the installer presents.
- 2 Enter a unique cluster name.

```
Enter the unique cluster name: [q,?] rac_cluster101
```

Configuring private heartbeat links

You now configure the private heartbeats that LLT uses. VCS provides the option to use LLT over Ethernet or over UDP (User Datagram Protocol). Symantec recommends that you configure heartbeat links that use LLT over Ethernet, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

The following procedure helps you configure LLT over Ethernet.

To configure private heartbeat links

- 1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or UDP.
 - Option 1: LLT over Ethernet (answer installer questions)
 Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.
 Skip to step [2](#).
 - Option 2: LLT over UDP (answer installer questions)
 Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
 Skip to step [3](#).
 - Option 3: Automatically detect configuration for LLT over Ethernet
 Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.
 Make sure that you activated the NICs for the installer to be able to detect and automatically configure the heartbeat links.

Skip to step 5.

- 2 If you chose option 1, enter the network interface card details for the private heartbeat links.

The installer discovers and lists the network interface cards.

You must not enter the network interface card that is used for the public network (typically eth0.)

Enter the NIC for the first private heartbeat link on galaxy:

[b,q,?] **eth1**

eth1 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use eth1 for the first private heartbeat link? [y,n,q,b,?] (n) **y**

Would you like to configure a second private heartbeat link?

[y,n,q,b,?] (y)

Enter the NIC for the second private heartbeat link on galaxy:

[b,q,?] **eth2**

eth2 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use eth2 for the second private heartbeat link? [y,n,q,b,?] (n) **y**

Would you like to configure a third private heartbeat link?

[y,n,q,b,?] (n)

- 3 If you chose option 2, enter the NIC details for the private heartbeat links. This step uses examples such as *private_NIC1* or *private_NIC2* to refer to the available names of the NICs.

```

Enter the NIC for the first private heartbeat
link on galaxy: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on galaxy: [b,q,?] (50000) ?
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat
link on galaxy: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on galaxy: [b,q,?] (50001) ?
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) y
Enter the NIC for the low priority heartbeat
link on galaxy: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on galaxy: [b,q,?] (50004)

```

- 4 Choose whether to use the same NIC details to configure private heartbeat links on other systems.

```

Are you using the same NICs for private heartbeat links on all
systems? [y,n,q,b,?] (y)

```

If you want to use the NIC details that you entered for galaxy, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

For LLT over UDP, if you want to use the same NICs on other systems, you still must enter unique IP addresses on each NIC for other systems.

If the NIC device names are different on some of the systems, enter **n**. Provide the NIC details for each system as the program prompts.

- 5 If you chose option 3, the installer detects NICs on each system and network links, and sets link priority.

If the installer fails to detect heartbeat links or fails to find any high-priority links, then choose option 1 or option 2 to manually configure the heartbeat links.

See step 2 for option 1, or step 3 for option 2.

- 6 Enter a unique cluster ID:

```
Enter a unique cluster ID number between 0-65535: [b,q,?] (60842)
```

The cluster cannot be configured if the cluster ID 60842 is in use by another cluster. Installer performs a check to determine if the cluster ID is duplicate. The check takes less than a minute to complete.

```
Would you like to check if the cluster ID is in use by another
cluster? [y,n,q] (y)
```

- 7 Verify and confirm the information that the installer summarizes.

A basic cluster is now configured. The remaining configuration settings are optional.

Note: You can proceed through the subsequent screens by just accepting the default value **n**.

Configuring the virtual IP of the cluster

You can configure the virtual IP of the cluster to use to connect from the Cluster Manager (Java Console), Veritas Operations Manager (VOM), or to specify in the RemoteGroup resource.

See the *Veritas Cluster Server Administrator's Guide* for information on the Cluster Manager.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for information on the RemoteGroup agent.

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 When the system prompts whether you want to configure the virtual IP, enter **y**.
- 3 Confirm whether you want to use the discovered public NIC on the first system. Do one of the following:
 - If the discovered NIC is the one to use, press **Enter**.

- If you want to use a different NIC, type the name of a NIC to use and press Enter.

```
Active NIC devices discovered on galaxy: eth1
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (eth1)
```

4 Confirm whether you want to use the same public NIC on all nodes.

Do one of the following:

- If all nodes use the same public NIC, enter *y*.
- If unique NICs are used, enter *n* and enter a NIC for each node.

```
Is eth1 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

5 Enter the virtual IP address for the cluster.

```
Enter the Virtual IP address for the Cluster:
[b,q,?] 192.168.1.16
```

6 Confirm the default netmask or enter another one:

```
Enter the netmask for IP 192.168.1.16: [b,q,?] (255.255.240.0)
```

7 Verify and confirm the Cluster Virtual IP information.

Cluster Virtual IP verification:

```
NIC: eth0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

Configuring the cluster in secure mode

The installer prompts whether you want to configure a secure cluster.

```
Would you like to configure the SF Oracle RAC cluster in secure mode?
[y,n,q,?] (n)
```

To configure a secure cluster, enter *y*.

If you want to confirm that the configured cluster is in secure mode, verify that the output of the following command is 1.

```
# haclust -value SecureClus
```

1

Setting up trust relationships for your SF Oracle RAC cluster

If you need to use an external authentication broker for authenticating VCS users, you must set up a trust relationship between VCS and the broker. For example, if Veritas Operations Manager (VOM) is your external authentication broker, the trust relationship ensures that VCS accepts the credentials that VOM issues.

Perform the following steps to set up a trust relationship between your SF Oracle RAC cluster and a broker.

To set up a trust relationship

- 1 Ensure that you are logged in as superuser on one of the nodes in the cluster.
- 2 Enter the following command:

```
# /opt/VRTS/install/installsfrac -securitytrust
```

The installer specifies the location of the log files. It then lists the cluster information such as cluster name, cluster ID, node names, and service groups.

- 3 When the installer prompts you for the broker information, specify the IP address, port number, and the data directory for which you want to establish trust relationship with the broker.

Input the broker name or IP address: **15.193.97.204**

Input the broker port: (14545)

Specify a port number or press Enter to accept the default port.

Input the data directory to setup trust with: (/var/VRTSvcs/vcsauth/data/HAD)

Specify a valid data directory or press Enter to accept the default directory.

- 4 The installer performs one of the following actions:

- If you specified a valid directory, the installer prompts for a confirmation.

Are you sure that you want to setup trust for the VCS cluster with the broker 15.193.97.204 and port 14545? [y,n,q] **y**

The installer sets up trust relationship with the broker for all nodes in the cluster and displays a confirmation.

```
Setup trust with broker 15.193.97.204 on cluster node1
.....Done

Setup trust with broker 15.193.97.204 on cluster node2
.....Done
```

The installer specifies the location of the log files, summary file, and response file and exits.

- If you entered incorrect details for broker IP address, port number, or directory name, the installer displays an error. It specifies the location of the log files, summary file, and response file and exits.

Configuring a secure cluster node by node

For environments that do not support passwordless ssh or passwordless rsh, you cannot use the `-security` option to enable secure mode for your cluster. Instead, you can use the `-securityonnode` option to configure a secure cluster node by node.

[Table 8-1](#) lists the tasks that you must perform to configure a secure cluster.

Table 8-1 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 124.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 125.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 126.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

- 1 Ensure that you are logged in as superuser.
- 2 Enter the following command:

```
# /opt/VRTS/install/installsfrac -securityonnode
```

The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure.

```
VCS is not running on all systems in this cluster. All VCS systems
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y
```

```
1) Perform security configuration on first node and export
security configuration files.
```

```
2) Perform security configuration on remaining nodes with
security configuration files.
```

```
Select the option you would like to perform [1-2,q,?] 1
```

Warning: All configurations about cluster users are deleted when you configure the first node. You can use the `/opt/VRTSvcs/bin/hauser` command to create cluster users manually.

- 3 The installer completes the secure configuration on the node. It specifies the location of the security configuration files and prompts you to copy these files to the other nodes in the cluster. The installer also specifies the location of log files, summary file, and response file.
- 4 Copy the security configuration files from the `/var/VRTSvcs/vcsauth/bkup` directory to temporary directories on the other nodes in the cluster.

Configuring the remaining nodes

On each of the remaining nodes in the cluster, perform the following steps.

To configure security on each remaining node

- 1 Ensure that you are logged in as superuser.
- 2 Enter the following command:

```
# /opt/VRTS/install/installsfrac -securityonnode
```

The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure. Enter **2**.

```
VCS is not running on all systems in this cluster. All VCS systems
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y
```

```
1) Perform security configuration on first node and export
security configuration files.
```

```
2) Perform security configuration on remaining nodes with
security configuration files.
```

```
Select the option you would like to perform [1-2,q.?] 2
```

The installer completes the secure configuration on the node. It specifies the location of log files, summary file, and response file.

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

- 1 On the first node, freeze all service groups except the ClusterService service group.

```
# /opt/VRTSvcs/bin/haconf -makerw
```

```
# /opt/VRTSvcs/bin/hagrp -list Frozen=0
```

```
# /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent
```

```
# /opt/VRTSvcs/bin/haconf -dump -makero
```

- 2 On the first node, stop the VCS engine.

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

- 3 On all nodes, stop the CmdServer.

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

- 4 On the first node, edit the `/etc/VRTSvcs/conf/config/main.cf` file to resemble the following:

```
cluster clus1 (
SecureClus = 1
)
```

- 5 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

- 6 On the first node, start VCS. Then start VCS on the remaining nodes.

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

- 7 On all nodes, start CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer
```

- 8 On the first node, unfreeze the service groups.

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hagrp -list Frozen=1
# /opt/VRTSvcs/bin/hagrp -unfreeze groupname -persistent
# /opt/VRTSvcs/bin/haconf -dump -makero
```

Adding VCS users

If you have enabled a secure VCS cluster, 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 wish to accept the default cluster credentials of
'admin/password'? [y,n,q] (y) n
Enter the user name: [b,q,?] (admin)
Enter the password:
Enter again:
```

- 3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

- 4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [b,q,?] smith
Enter New Password:*****

Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [b,q,?] a
```

- 5 Enter **n** at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```

- 6 Review the summary of the newly added users and confirm the information.

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.
- 2 Specify whether you want to configure the SMTP notification.

Do you want to configure SMTP notification? [y,n,q,?] (n) **y**

If you do not want to configure the SMTP notification, you can skip to the next configuration option.

See [“Configuring SNMP trap notification”](#) on page 130.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SMTP server's host name.

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

- Enter the email address of each recipient.

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

- Enter the minimum security level of messages to be sent to each recipient.

```
Enter the minimum severity of events for which mail should be
sent to ozzie@example.com [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] w
```

- 4 Add more SMTP recipients, if necessary.

- If you want to add another SMTP recipient, enter **y** and provide the required information at the prompt.

Would you like to add another SMTP recipient? [y,n,q,b] (n) **y**

```
Enter the full email address of the SMTP recipient
```

```
(example: user@yourcompany.com): [b,q,?] harriet@example.com
```

```
Enter the minimum severity of events for which mail should be
sent to harriet@example.com [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] E
```

- If you do not want to add, answer **n**.

```
Would you like to add another SMTP recipient? [y,n,q,b] (n)
```

5 Verify and confirm the SMTP notification information.

```
NIC: eth0
```

```
SMTP Address: smtp.example.com
```

```
Recipient: ozzie@example.com receives email for Warning or
higher events
```

```
Recipient: harriet@example.com receives email for Error or
higher events
```

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

Configuring SNMP trap notification

You can choose to configure VCS to send event notifications to SNMP management consoles. You need to provide the SNMP management console name to be notified and message severity levels.

Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure the SNMP trap notification

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

```
Do you want to configure SNMP notification? [y,n,q,?] (n) y
```

If you skip this option and if you had installed a valid HA/DR license, the installer presents you with an option to configure this cluster as global cluster. If you did not install an HA/DR license, the installer proceeds to configure SF Oracle RAC based on the configuration details you provided.

See [“Configuring global clusters”](#) on page 132.

3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SNMP trap daemon port.

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

- Enter the SNMP console system name.

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

- Enter the minimum security level of messages to be sent to each console.

```
Enter the minimum severity of events for which SNMP traps
should be sent to saturn [I=Information, W=Warning, E=Error,
S=SevereError]: [b,q,?] E
```

4 Add more SNMP consoles, if necessary.

- If you want to add another SNMP console, enter `y` and provide the required information at the prompt.

```
Would you like to add another SNMP console? [y,n,q,b] (n) y
Enter the SNMP console system name: [b,q,?] jupiter
Enter the minimum severity of events for which SNMP traps
should be sent to jupiter [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] S
```

- If you do not want to add, answer `n`.

Would you like to add another SNMP console? [y,n,q,b] (n)

5 Verify and confirm the SNMP notification information.

NIC: eth0

SNMP Port: 162

Console: saturn receives SNMP traps for Error or higher events

Console: jupiter receives SNMP traps for SevereError or higher events

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds with other set of questions for CVM and CFS.

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the VCS configuration file. You must perform additional configuration tasks to set up a global cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to set up SF Oracle RAC global clusters.

Note: If you installed a HA/DR license to set up campus cluster, skip this installer option.

To configure the global cluster option

- 1 Review the required information to configure the global cluster option.
- 2 Specify whether you want to configure the global cluster option.

Do you want to configure the Global Cluster Option? [y,n,q] (n) **y**

If you skip this option, the installer proceeds to configure VCS based on the configuration details you provided.

3 Provide information to configure this cluster as global cluster.

The installer prompts you for a NIC, a virtual IP address, and value for the netmask.

If you had entered virtual IP address details, the installer discovers the values you entered. You can use the same virtual IP address for global cluster configuration or enter different values.

4 Verify and confirm the configuration of the global cluster.

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

```
NIC: eth0
```

```
IP: 192.168.1.16
```

```
Netmask: 255.255.240.0
```

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

Creation of SF Oracle RAC configuration files

The program consolidates all the information gathered in the preceding configuration tasks and creates configuration files.

If you chose to configure the cluster in secure mode, the installer also configures the Symantec Product Authentication Service, which creates an Authentication Broker with root and authentication mode.

Review the output as the configuration program starts VCS, creates VCS configuration files, and copies the files to each node.

Stopping and starting SF Oracle RAC processes

The installer stops and starts SF Oracle RAC processes and configures the SF Oracle RAC agents.

Note: Do not opt to start SF Oracle RAC now if you want to configure private heartbeats to use aggregated interfaces that the installer has not discovered or to use aggregated interfaces on nodes that run SLES.

To stop SF Oracle RAC processes

- 1 Enter **y** to stop SF Oracle RAC processes.

Do you want to stop SF Oracle RAC processes now? [y,n,q,?] (y)

- 2 Review the output as the installer stops and starts the SF Oracle RAC processes.

Note that SF Oracle RAC configuration program starts I/O fencing feature in disabled mode. SF Oracle RAC requires you to configure and enable I/O fencing feature.

Configuring SF Oracle RAC using the Web-based installer

Before you begin to configure SF Oracle RAC using the Web-based installer, review the configuration requirements.

By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.

You can click **Quit** to quit the Web-installer at any time during the configuration process.

To configure SF Oracle RAC on a cluster

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 86.

- 2 On the Select a task and a product page, select the task and the product as follows:

Task	Configure a Product
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 3 On the Select Systems page, enter the system names where you want to configure SF Oracle RAC, and click **Next**.

Example: **galaxy nebula**

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

Click **Next** after the installer completes the system verification successfully.

- 4 In the Confirmation dialog box that appears, choose whether or not to configure I/O fencing.

To configure I/O fencing, click **Yes**.

To configure I/O fencing later, click **No**. You can configure I/O fencing later using the Web-based installer.

See [“Configuring SF Oracle RAC for data integrity using the Web-based installer”](#) on page 155.

You can also configure I/O fencing later using the `installsfrac -fencing` command, the response files, or manually configure.

- 5 On the Set Cluster Name/ID page, specify the following information for the cluster.

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID. Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments.
Check duplicate cluster ID	Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete.
LLT Type	Select an LLT type from the list. You can choose to configure LLT over UDP or over Ethernet. If you choose Auto detect over Ethernet, the installer auto-detects the LLT links over Ethernet. Verify the links and click Yes in the Confirmation dialog box. Skip to step 7. If you click No, you must manually enter the details to configure LLT over Ethernet.
Number of Heartbeats	Choose the number of heartbeat links you want to configure.
NICs	Choose the NICs that you want to configure.
Additional Low Priority Heartbeat NIC	Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link.
Unique Heartbeat NICs per system	For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems. For LLT over UDP, this check box is selected by default.

Click **Next**.

- 6 On the Set Cluster Heartbeat page, select the heartbeat link details for the LLT type you chose on the Set Cluster Name/ID page.

For **LLT over Ethernet**: Do the following:

- If you are using the same NICs on all the systems, select the NIC for each private heartbeat link.
- If you had selected **Unique Heartbeat NICs per system** on the Set Cluster Name/ID page, provide the NIC details for each system.

For **LLT over UDP**: Select the NIC, Port, and IP address for each private heartbeat link. You must provide these details for each system.

Click **Next**.

- 7 On the Optional Configuration page, decide the optional VCS features that you want to configure. Click the corresponding tab to specify the details for each option:

Virtual IP

- Select the **Configure Virtual IP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select the interface on which you want to configure the virtual IP.
- Enter a virtual IP address and value for the netmask.

VCS Users

- Reset the password for the Admin user, if necessary.
- Select the **Configure VCS users** option.
- Click **Add** to add a new user.
Specify the user name, password, and user privileges for this user.

SMTP

- Select the **Configure SMTP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SMTP Server** box, enter the domain-based hostname of the SMTP server. Example: smtp.yourcompany.com
- In the **Recipient** box, enter the full email address of the SMTP recipient. Example: user@yourcompany.com.
- In the **Event** list box, select the minimum security level of messages to be sent to each recipient.
- Click **Add** to add more SMTP recipients, if necessary.

SNMP

- Select the **Configure SNMP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SNMP Port** box, enter the SNMP trap daemon port: (162).
- In the **Console System Name** box, enter the SNMP console system name.
- In the **Event** list box, select the minimum security level of messages to be sent to each console.
- Click **Add** to add more SNMP consoles, if necessary.

GCO

If you installed a valid HA/DR license, you can now enter the wide-area heartbeat link details for the global cluster that you would set up later.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for instructions to set up SF Oracle RAC global clusters.

- Select the **Configure GCO** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select a NIC.
- Enter a virtual IP address and value for the netmask.

Security

To configure a secure SF Oracle RAC cluster, select the **Configure secure cluster** check box.

If you want to perform this task later, do not select the **Configure secure cluster** check box. You can use the `-security` option of the `installsrac` program.

Click **Next**.

- 8 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.
- 9 On the Start Processes page, click **Next** after the installer performs the configuration based on the details you provided and starts all the processes successfully.

If you did not choose to configure I/O fencing in step 4, then skip to step 11. Go to step 10 to configure fencing.

- 10 On the Select Fencing Type page, choose the type of fencing configuration:

Configure Coordination Point client based fencing Choose this option to configure server-based I/O fencing.

Configure disk based fencing Choose this option to configure disk-based I/O fencing.

Based on the fencing type you choose to configure, follow the installer prompts.

See [“Configuring SF Oracle RAC for data integrity using the Web-based installer”](#) on page 155.

- 11 Click **Next** to complete the process of configuring SF Oracle RAC.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

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

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

Configuring SF Oracle RAC clusters for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using installsfrac program](#)
- [Setting up disk-based I/O fencing manually](#)
- [Configuring SF Oracle RAC for data integrity using the Web-based installer](#)
- [Setting up server-based I/O fencing using installsfrac program](#)
- [Setting up server-based I/O fencing manually](#)

Setting up disk-based I/O fencing using installsfrac program

You can configure I/O fencing using the `-fencing` option of the `installsfrac` program.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

- 1 List the new external disks or the LUNs as recognized by the operating system.
On each node, enter:


```
# fdisk -l
```
- 2 To initialize the disks as VxVM disks, use one of the following methods:
 - Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks.

For more information see the *Veritas Storage Foundation 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 sdr
```

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

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 140.

Review the following procedure to identify disks to use as coordinator disks.

To identify the coordinator disks

- 1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# vxdisk -o all dgs list
```

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

See [“Checking shared disks for I/O fencing”](#) on page 141.

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure SF Oracle RAC meets the I/O fencing requirements. You can test the shared disks using the `vxfcntlshdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

Make sure to test the disks that serve as coordinator disks.

The `vxfststhdw` utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

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

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

- Verifying the Array Support Library (ASL)
See [“Verifying Array Support Library \(ASL\)”](#) on page 142.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 143.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfststhdw utility”](#) on page 144.

Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

To verify Array Support Library (ASL)

- 1 If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

The ASL for the supported storage device that you add is available from the disk array vendor or Symantec technical support.

- 2 Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

The following output is a sample:

```
# vxddladm listsupport all
```

LIBNAME	VID	PID
libvxhitachi.so	HITACHI	DF350, DF400, DF400F, DF500, DF500F
libvxxp1281024.so	HP	All
libvxxp12k.so	HP	All
libvxddns2a.so	DDN	S2A 9550, S2A 9900, S2A 9700
libvxpurple.so	SUN	T300
libvxxiotechE5k.so	XIOTECH	ISE1400
libvxcopan.so	COPANSYS	8814, 8818
libvxibmds8k.so	IBM	2107

- 3 Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

```
# vxdisk scandisks
```

See the Veritas Volume Manager documentation for details on how to add and configure disks.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the vxfcntlshdw 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.

```
# vxfenadm -i diskpath
```

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

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

From node A, enter:

```
# vxfenadm -i /dev/sdx
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

```
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/sdy` path.

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

```
# vxfenadm -i /dev/sdz
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

```
Vendor id      : HITACHI
```

```
Product id     : OPEN-3
```

```
Revision       : 0117
```

```
Serial Number  : 0401EB6F0002
```

Testing the disks using vxfcntlthdw utility

This procedure uses the `/dev/sdx` disk in the steps.

If the utility does not show a message that states a disk is ready, the verification has failed. Failure of verification can be the result of an improperly configured disk array. The failure can also be due to a bad disk.

If the failure is due to a bad disk, remove and replace it. The `vxfcntlsthaw` utility indicates a disk can be used for I/O fencing with a message resembling:

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

For more information on how to replace coordinator disks, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

To test the disks using `vxfcntlsthaw` utility

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

See [“Setting up inter-system communication”](#) on page 662.

- 2 From one node, start the utility.

Run the utility with the `-n` option if you use `rsh` for communication.

```
# vxfcntlsthaw [-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:
for dmp: /dev/vx/rdmp/sdx
for raw: /dev/sdx
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdr
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
nebula in the format:
for dmp: /dev/vx/rdmp/sdx
for raw: /dev/sdx
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdr
```

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 reports its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node galaxy.

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

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

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

Configuring disk-based I/O fencing using installsfrac program

Note: The installer stops and starts SF Oracle RAC to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop SF Oracle RAC.

To set up disk-based I/O fencing using the installsfrac program

- 1 Start the installsfrac program with `-fencing` option.

```
# /opt/VRTS/install/installsfrac -fencing
```

The installsfrac program starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether SF Oracle RAC 6.0 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **2** to configure disk-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-4,b,q] 2
```

- 4 Review the output as the configuration program checks whether VxVM is already started and is running.

- If the check fails, configure and enable VxVM before you repeat this procedure.
- If the check passes, then the program prompts you for the coordinator disk group information.

- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.
The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
 - Enter the number corresponding to the **Create a new disk group** option.
The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.

Symantec recommends that you use three disks as coordination points for disk-based I/O fencing.

If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
- Enter the disk group name.

6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the `vxfsntsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 141.

7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

8 Enter the I/O fencing disk policy that you chose to use. For example:

```
Enter disk policy for the disk(s) (raw/dmp): [b,q,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
- Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information

9 Verify and confirm the I/O fencing configuration information that the installer summarizes.

10 Review the output as the configuration program does the following:

- Stops VCS and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the VCS configuration file `main.cf` if necessary.
- Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
- Starts VCS on each node to make sure that the SF Oracle RAC is cleanly configured to use the I/O fencing feature.

- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
- 12 Configure the Coordination Point agent to monitor the coordinator disks.
See [“Configuring CoordPoint agent to monitor coordination points”](#) on page 167.

Setting up disk-based I/O fencing manually

[Table 9-1](#) lists the tasks that are involved in setting up I/O fencing.

Table 9-1 Tasks to set up I/O fencing manually

Task	Reference
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 140.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 141.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 141.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 149.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 150.
Modifying SF Oracle RAC configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 151.
Configuring CoordPoint agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 167.
Starting SF Oracle RAC on all nodes	See “Starting SF Oracle RAC on all nodes” on page 152.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 154.

Setting up coordinator disk groups

From one node, create a disk group named vxfencoorddg. This group must contain three disks or LUNs. You must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

Note that if you create a coordinator disk group as a regular disk group, you can turn on the coordinator attribute in Volume Manager.

Refer to the *Veritas Storage Foundation Administrator's Guide* for details on how to create disk groups.

The following example procedure assumes that the disks have the device names `sdx`, `sdz`, and `sdz`.

To create the `vxfencoorddg` disk group

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

```
# vxdg init vxfencoorddg sdx sdz sdz
```

- 2 Set the coordinator attribute value as "on" for the coordinator disk group.

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

After you set up the coordinator disk group, you must do the following to configure I/O fencing:

- Create the I/O fencing configuration file `/etc/vxfendg`
- Update the I/O fencing configuration file `/etc/vxfenmode`

To update the I/O fencing files and start I/O fencing

- 1 On each nodes, type:

```
# echo "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.

- 2 On all cluster nodes depending on the SCSI-3 mechanism, type one of the following selections:

- For DMP configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
```

- For raw device configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_raw /etc/vxfenmode
```

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

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

- 4 Edit the following file on each node in the cluster to change the values of the VXFEN_START and the VXFEN_STOP environment variables to 1:

```
/etc/sysconfig/vxfen
```

Modifying VCS configuration to use I/O fencing

After you add coordination points and configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the VCS configuration file /etc/VRTSvcs/conf/config/main.cf.

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

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

```
# /etc/init.d/vxfen 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
)
```

Regardless of whether the fencing configuration is disk-based or server-based, the value of the cluster-level attribute UseFence is set to SCSI3.

- 6 Save and close the file.
- 7 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 8 Using rcp or another utility, copy the VCS configuration file from a node (for example, galaxy) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp galaxy:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

Starting SF Oracle RAC on all nodes

You must start SF Oracle RAC on each node to bring up the cluster configuration with I/O fencing enabled.

Starting I/O fencing, VCS, CVM, and CFS

You must start I/O fencing, VCS, CVM, and CFS on all nodes in the cluster.

To start VCS, CVM, and CFS on a node

- 1 Start the I/O fencing driver. Run the following command on each node:

```
# /etc/init.d/vxfen 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 174.

To verify GAB port membership

- ◆ Run the `gabconfig -a` command.

For example:

```
galaxy# gabconfig -a
GAB Port Memberships
=====
Port a gen ada401 membership 01
Port b gen ada40d membership 01
Port d gen ada409 membership 01
Port f gen ada41c membership 01
Port h gen ada40f membership 01
Port o gen ada406 membership 01
Port u gen ada41a membership 01
Port v gen ada416 membership 01
Port w gen ada418 membership 01
Port y gen ada424 membership 01
```

Verifying the CVM group is online

Make sure the cvm group is in the online state.

To verify CVM group

- 1 On any node, run the following command to verify that the cvm group is online:

```
# hagrps -state cvm
```

- 2 On any node, run the following command to make sure CVM has started properly:

```
# vxclustadm nidmap
```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the SCSI-3 disk policy reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- 1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is SCSI3 and the SCSI3 disk policy is dmp:

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: dmp
```

```
Cluster Members:
```

```
* 0 (galaxy)
```

```
1 (nebula)
```

```
RFSM State Information:
```

```
node 0 in state 8 (running)
```

```
node 1 in state 8 (running)
```

- 2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

Configuring SF Oracle RAC for data integrity using the Web-based installer

After you configure SF Oracle RAC, you must configure the cluster for data integrity. Review the configuration requirements.

See [“Configuring SF Oracle RAC using the Web-based installer”](#) on page 134.

See [“About planning to configure I/O fencing”](#) on page 59.

To configure SF Oracle RAC for data integrity

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 86.

- 2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

- 4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

- 5 On the Select Fencing Type page, choose whether to configure disk-based fencing or server-based fencing.

If you chose to configure disk-based fencing, go to step [6](#).

If you chose to configure server-based fencing, go to step [9](#).

- 6 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

Click **Next**.

- 7 On the Configure Fencing page, specify the following information:

- Select a Disk Group** Select the **Create a new disk group** option or select one of the disk groups from the list.
- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
Go to step [15](#).
 - If you selected the **Create a new disk group** option, make sure you have SCSI-3 PR enabled disks, and click **Yes** in the confirmation dialog box.
Click **Next**. Go to step [8](#).

8 On the Create New DG page, specify the following information:

- New Disk Group Name** Enter a name for the new coordinator disk group you want to create.
- Select Disks** Select at least three disks to create the coordinator disk group.

If you want to select more than three disks, make sure to select an odd number of disks.
- Fencing Disk Policy** Choose the fencing disk policy for the disk group.

Go to step [15](#).

9 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.
Click **Next**.

10 Provide the following details for each of the CP servers:

- Enter the virtual IP addresses or host names of the virtual IP address. The installer assumes these values to be identical as viewed from all the application cluster nodes.
- Enter the port that the CP server must listen on.
- Click **Next**.

11 If your server-based fencing configuration also uses disks as coordination points, perform the following steps:

- If you have not already checked the disks for SCSI-3 PR compliance, check the disks now, and click **OK** in the dialog box.
- If you do not want to use the default coordinator disk group name, enter a name for the new coordinator disk group you want to create.

- Select the disks to create the coordinator disk group.
 - Choose the fencing disk policy for the disk group.
- 12 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.
 - 13 Verify and confirm the I/O fencing configuration information.
 The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.
 - 14 Configure the CP agent on the SF Oracle RAC (application cluster), and click **Next**.
 - 15 Click **Next** to complete the process of configuring I/O fencing.
 On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.
 - 16 Select the checkbox to specify whether you want to send your installation information to Symantec.
 Click **Finish**. The installer prompts you for another task.

Setting up server-based I/O fencing using installsfrac program

See [“Setting up server-based I/O fencing using installsfrac program”](#) on page 157.

Setting up server-based I/O fencing manually

Tasks that are involved in setting up server-based I/O fencing manually include:

Table 9-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the SF Oracle RAC cluster	See “Preparing the CP servers manually for use by the SF Oracle RAC cluster” on page 158.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See “Configuring server-based fencing on the SF Oracle RAC cluster manually” on page 161.
Modifying SF Oracle RAC configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 151.

Table 9-2 Tasks to set up server-based I/O fencing manually (*continued*)

Task	Reference
Configuring Coordination Point agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 167.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 168.

Preparing the CP servers manually for use by the SF Oracle RAC cluster

Use this procedure to manually prepare the CP server for use by the SF Oracle RAC cluster or clusters.

[Table 9-3](#) displays the sample values used in this procedure.

Table 9-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1
Node #1 - SF Oracle RAC cluster	galaxy
Node #2 - SF Oracle RAC cluster	nebula
Cluster name	rac_cluster101
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SF Oracle RAC cluster

- 1 Determine the cluster name and uuid on the SF Oracle RAC cluster.

For example, issue the following commands on one of the SF Oracle RAC cluster nodes (galaxy):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster rac_cluster101

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2-bb31-00306eea460a}
```

- 2 Use the `cpsadm` command to check whether the SF Oracle RAC cluster and nodes are present in the CP server.

For example:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes
```

ClusName	UUID	Hostname(Node ID)	Registered
rac_cluster101	{f0735332-1dd1-11b2-bb31-00306eea460a}	galaxy(0)	0
rac_cluster101	{f0735332-1dd1-11b2-bb31-00306eea460a}	nebula(1)	0

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

3 Add the SF Oracle RAC cluster and nodes to each CP server.

For example, issue the following command on the CP server (mycps1.symantecexample.com) to add the cluster:

```
# cpsadm -s mycps1.symantecexample.com -a add_clus\  
-c rac_cluster101 -u {f0735332-1dd1-11b2}
```

```
Cluster rac_cluster101 added successfully
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the first node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c rac_cluster101 -u {f0735332-1dd1-11b2} -h galaxy -n0
```

```
Node 0 (galaxy) successfully added
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the second node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c rac_cluster101 -u {f0735332-1dd1-11b2} -h nebula -n1
```

```
Node 1 (nebula) successfully added
```

4 If security is to be enabled, check whether the CPSADM@VCS_SERVICES@cluster_uuid users are created in the CP server.

If the output below does not show the users, then add them as described in the next step.

```
# cpsadm -s mycps1.symantecexample.com -a list_users
```

Username/Domain	Type	Cluster Name / UUID	Role
CPSADM@VCS_SERVICES@f0735332-1dd1-11b2/vx		rac_cluster101/{f0735332-1dd1-11b2}	Operator

If security is to be disabled, then add the user name "cpsclient@hostname" to the server instead of the CPSADM@VCS_SERVICES@cluster_uuid (for example, cpsclient@galaxy).

The CP server can only run in either secure mode or non-secure mode, both connections are not accepted at the same time.

5 Add the users to the CP server.

Issue the following commands on the CP server
(mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
CPSADM@VCS_SERVICES@cluster_uuid\  
-f cps_operator -g vx
```

```
User CPSADM@VCS_SERVICES@cluster_uuid  
successfully added
```

6 Authorize the CP server user to administer the SF Oracle RAC cluster. You must perform this task for the CP server users corresponding to each node in the SF Oracle RAC cluster.

For example, issue the following command on the CP server
(mycps1.symantecexample.com) for SF Oracle RAC cluster rac_cluster101
with two nodes galaxy and nebula:

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c rac_cluster101\  
-u {f0735332-1dd1-11b2}\  
-e CPSADM@VCS_SERVICES@cluster_uuid\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
CPSADM@VCS_SERVICES@cluster_uuid privileges.
```

Configuring server-based fencing on the SF Oracle RAC cluster manually

The configuration process for the client or SF Oracle RAC cluster to use CP server as a coordination point requires editing the `/etc/vxfenmode` file.

You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

Note: Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxfencoorddg). You must specify this disk group in the `/etc/vxfenmode` file.

See [“Setting up coordinator disk groups”](#) on page 149.

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

To configure server-based fencing on the SF Oracle RAC cluster manually

- 1 Use a text editor to edit the following file on each node in the cluster:

```
/etc/sysconfig/vxfen
```

You must change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

- 2 Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

If your server-based fencing configuration uses a single highly available CP server as its only coordination point, make sure to add the `single_cp=1` entry in the `/etc/vxfenmode` file.

The following sample file output displays what the `/etc/vxfenmode` file contains:

See [“Sample vxfenmode file output for server-based fencing”](#) on page 162.

- 3 After editing the `/etc/vxfenmode` file, run the `vxfen init` script to start fencing.

For example:

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

- 4 For CP servers in secure mode, make sure that the security is enabled on the cluster and the credentials for the CPSADM are present in the `/var/VRTSvcs/vcsauth/data/CPSADM` directory.

Sample vxfenmode file output for server-based fencing

The following is a sample `vxfenmode` file for server-based fencing:

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3          - use scsi3 persistent reservation disks
```

```
# customized - use script based customized fencing
# disabled - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps - use a coordination point server with optional script
# controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp

# security when enabled uses secure communication to the cp server
# using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
# communication
# 1 - use Veritas Authentication Service for cp server
# communication
security=1

#
# Specify 3 or more odd number of coordination points in this file,
# one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks. Please ensure that the CP
# server coordination points are numbered sequentially and in the
# same order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
```

```

#
# cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames over
# different subnets, all of the IPs/names can be specified in a
# comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,...,
#   [<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#   is the serial number of the CPS as a coordination point; must
#   start with 1.
# <vip>
#   is the virtual IP address of the CPS, must be specified in
#   square brackets ("[]").
# <vhn>
#   is the virtual hostname of the CPS, must be specified in square
#   brackets ("[]").
# <port>
#   is the port number bound to a particular <vip/vhn> of the CPS.
#   It is optional to specify a <port>. However, if specified, it
#   must follow a colon (":") after <vip/vhn>. If not specified, the
#   colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>, a
# default port can be specified as follows:
#
# port=<default_port>
#
# Where <default_port> is applicable to all the <vip/vhn>s for
# which a <port> is not specified. In other words, specifying <port>
# with a <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[mycps1.company.com]
# cps2=[192.168.0.25]
# cps3=[mycps2.company.com]:59999
    
```

```

#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [mycps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#   [192.168.0.23]
#   [mycps1.company.com]
#   [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be used
#   for all remaining <vip/vhn>s:
#   [192.168.0.23]
#   [mycps1.company.com]
#   [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#   vx fendg=<coordinator disk group name>
# Example:
#   vx fendg=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
#   cps1=
#   cps2=
#   cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
#   coordinator disks
#   cps1=
#   vx fendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
#   vx fendg=
# Note: The disk group specified in case should have three disks
#

```

Table 9-4 defines the vxfenmode parameters that must be edited.

Table 9-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to "customized".
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".
scsi3_disk_policy	Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw". Note: The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that secure mode is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1.
cps1, cps2, or vxfendg	Coordination point parameters. Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server. <i>cps<number>=[virtual_ip_address/virtual_host_name]:port</i> Where <i>port</i> is optional. The default port value is 14250. If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example: <i>cps1=[192.168.0.23], [192.168.0.24]:58888, [mycps1.company.com]</i> Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxencoordg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file).

Table 9-4 vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
port	Default port for the CP server to listen on If you have not specified port numbers for individual virtual IP addresses or host names, the default port number value that the CP server uses for those individual virtual IP addresses or host names is 14250. You can change this default port value using the port parameter.
single_cp	Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point. Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.

Configuring CoordPoint agent to monitor coordination points

The following procedure describes how to manually configure the CoordPoint agent to monitor coordination points.

The CoordPoint agent can monitor CP servers and SCSI-3 disks.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on the agent.

To configure CoordPoint agent to monitor coordination points

- 1 Ensure that your SF Oracle RAC cluster has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

```
# haconf -makerw
# hagr -add vxfen
# hagr -modify vxfen SystemList galaxy 0 nebula 1
# hagr -modify vxfen AutoFailOver 0
# hagr -modify vxfen Parallel 1
# hagr -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SF Oracle RAC cluster using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state coordpoint

# Resource      Attribute    System    Value
coordpoint      State        galaxy    ONLINE
coordpoint      State        nebula    ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed CoordPoint agent monitoring information; including information about whether the CoordPoint agent is able to access all the coordination points, information to check on which coordination points the CoordPoint agent is reporting missing keys, etc.

To view all such information in the engine log, change the dbg level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

```
/var/VRTSvcS/log/engine_A.log
```

Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

To verify the server-based I/O fencing configuration

- 1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

Note: For troubleshooting any server-based I/O fencing configuration issues, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

- 2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

Performing post-installation and configuration tasks

This chapter includes the following topics:

- [Performing a postcheck on a node](#)
- [Verifying SF Oracle RAC installation using VCS configuration file](#)
- [Verifying LLT, GAB, and cluster operation](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [About configuring authentication for SFDB tools](#)
- [Configuring Veritas Volume Replicator](#)
- [Running SORT Data Collector to collect configuration information](#)

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See [“About using the postcheck option”](#) on page 591.

To run the postcheck command on a node

- 1 Run the installer with the `-postcheck` option.

```
# ./installer -postcheck system_name
```

- 2 Review the output for installation-related information.

Verifying SF Oracle RAC installation using VCS configuration file

The configuration file, `main.cf`, is created on each node at `/etc/VRTSvcs/conf/config/`. Review the `main.cf` configuration file after the SF Oracle RAC installation and before the Oracle installation.

Verify the following information in the `main.cf` file:

- The cluster definition within the `main.cf` includes the cluster information that was provided during the configuration. The cluster information includes the cluster name, cluster address, and the names of cluster users and administrators.
- The `UseFence = SCSI3` attribute is not automatically present; you must manually add it after the installation.
- If you configured the cluster in secure mode, the “`SecureClus = 1`” cluster attribute is set.

For more information on the configuration file:

See [“About VCS configuration file”](#) on page 631.

Verifying LLT, GAB, and cluster operation

Verify the operation of LLT, GAB, and the cluster using the VCS commands.

To verify LLT, GAB, and cluster operation

- 1 Log in to any node in the cluster as superuser.
- 2 Make sure that the `PATH` environment variable is set to run the VCS commands.
- 3 Verify LLT operation.
See [“Verifying LLT”](#) on page 172.
- 4 Verify GAB operation.
See [“Verifying GAB”](#) on page 174.

- 5 Verify the cluster operation.

See “[Verifying the cluster](#)” on page 175.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State      Links
*0 galaxy      OPEN      2
 1 nebula      OPEN      2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State      Links
* 0 galaxy      OPEN      2
 1 nebula      OPEN      2
 2 saturn      OPEN      1
```

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

LLT node information:

Node	State	Links
0 galaxy	OPEN	2
*1 nebula	OPEN	2

- To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv active
```

The output on galaxy resembles:

Node	State	Link	Status	Address
*0 galaxy	OPEN			
		eth1	UP	08:00:20:93:0E:34
		eth2	UP	08:00:20:93:0E:38
1 nebula	OPEN			
		eth1	UP	08:00:20:8F:D1:F2
		eth2	DOWN	

The command reports the status on the two active nodes in the cluster, galaxy and nebula.

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

- A state of OPEN
- A status for each link of UP
- 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 ... 60 61 62 63
      connects:   0 1
7      gab        0x7
      opens:      0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects:   0 1
31     gab        0x1F
      opens:      0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects:   0 1

```

Verifying GAB

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information. The output displays the nodes that have membership with the modules you installed and configured. You can use GAB port membership as a method of determining if a specific component of the SF Oracle RAC stack communicates with its peers.

[Table 10-1](#) lists the different ports that the software configures for different functions.

Table 10-1 GAB port description

Port	Function
a	GAB
b	I/O fencing
d	Oracle Disk Manager (ODM)
f	Cluster File System (CFS)
h	Veritas Cluster Server (VCS: High Availability Daemon)
o	VCSMM driver
u	Cluster Volume Manager (CVM) (to ship commands from slave node to master node) Port u in the <code>gabconfig</code> output is visible with CVM protocol version ≥ 100 .
v	Cluster Volume Manager (CVM)
w	vxconfigd (module for CVM)
y	Cluster Volume Manager (CVM) I/O shipping

For more information on GAB, refer to the *Veritas Cluster Server Administrator's Guide*.

To verify GAB

- ◆ To verify the GAB operation, type the following command on each node:

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

For example, the command returns the following output:

```
GAB Port Memberships
=====
Port a gen  ada401 membership 01
Port b gen  ada40d membership 01
Port d gen  ada409 membership 01
Port f gen  ada41c membership 01
Port h gen  ada40f membership 01
Port o gen  ada406 membership 01
Port u gen  ada41a membership 01
Port v gen  ada416 membership 01
Port w gen  ada418 membership 01
Port y gen  ada42a membership 01
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for a description of system states and the transitions between them.

To verify the cluster

- 1 To verify the status of the cluster, type the following command:

```
# 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  cvm                galaxy              Y        N              ONLINE
B  cvm                nebula              Y        N              ONLINE
```

- 2 Review the command output for the following information:
 - The system state
If the value of the system state is **RUNNING**, the cluster is successfully started.
 - The cvm group state
In the sample output, the group state lists the cvm group, which is **ONLINE** on both the nodes galaxy and nebula.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for information about the system attributes for VCS.

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	CPUThresholdLevel	Critical 90 Warning 80 Note 70 Info 60
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	
galaxy	ConfigChecksum	
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcs/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Thu Sep 22 07:14:23 CDT 2011
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	6.0.00.0
galaxy	FencingWeight	0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 7 Swap 0
galaxy	LLTNodeId	0

galaxy	LicenseType	PERMANENT_SITE
galaxy	Limits	
galaxy	LinkHbStatus	eth1 UP eth2 UP
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	PhysicalServer	
galaxy	ShutdownTimeout	600
galaxy	SourceFile	./main.cf
galaxy	SwapThresholdLevel	Critical 90 Warning 80 Note 70 Info 60
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	SystemRecipients	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	NONE
galaxy	VCSMode	VCS_RAC

About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Oracle Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.


See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 180.

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

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

[Figure 10-1](#) depicts the SF Oracle RAC cluster communication with the LDAP servers when clusters run in secure mode.

Figure 10-1 Client communication with LDAP servers



The LDAP schema and syntax for LDAP commands (such as `ldapadd`, `ldapmodify`, and `ldapsearch`) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is `posixAccount`)
 - UserObject Attribute (the default is `uid`)
 - User Group Attribute (the default is `gidNumber`)
 - Group Object Class (the default is `posixGroup`)
 - GroupObject Attribute (the default is `cn`)
 - Group GID Attribute (the default is `gidNumber`)
 - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, `UserBaseDN=ou=people,dc=comp,dc=com`)
- Distinguished name for the group container (for example, `GroupBaseDN=ou=group,dc=comp,dc=com`)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.6.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com" \
--server_url "ldap://my_openldap_host.symantecexample.com" \
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com" \
--user_attribute "cn" --user_object_class "account" \
--user_gid_attribute "gidNumber" \
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com" \
--group_attribute "cn" --group_object_class "posixGroup" \
--group_gid_attribute "member" \
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com" \
--admin_user_password "password"--auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the SF Oracle RAC nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, `vcsadmin1`.

```
galaxy# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:14149
```

Enter password for vcsadmin1: #####

```
authenticate
-----
-----
```

```
Authenticated User vcsadmin1
-----
```

3 Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

4 Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster rac_cluster101 (
    SecureClus = 1
    Administrators = {
        "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" }
    AdministratorGroups = {
        "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com " }
    )
...
...
```

5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=myenterprise.symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute      Value
galaxy       Attribute      RUNNING
nebula       Attribute      RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SF Oracle RAC node using the VCS Cluster Manager (Java Console).

7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -d \  
-s domain_controller_name_or_ipaddress \  
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-d -s 192.168.20.32 -u Administrator -g "Domain Admins"
```

Search User provided is invalid or Authentication is required to proceed further.
 Please provide authentication information for LDAP server.

Username/Common Name: **symantecdomain\administrator**
 Password:

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d windows_domain_name
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d symantecdomain.com
```

Attribute list file not provided, using default AttributeList.txt.
 CLI file name not provided, using default CLI.txt.

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x
```


- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL :          ldap://192.168.20.32:389
SSL Enabled :          No
User Base DN :          CN=people,DC=symantecdomain,DC=com
User Object Class :    account
User Attribute :       cn
User GID Attribute :   gidNumber
Group Base DN :        CN=group,DC=symantecdomain,DC=com
Group Object Class :   group
Group Attribute :      cn
Group GID Attribute :  cn
Auth Type :            FLAT
Admin User :
Admin User Password :
Search Scope :         SUB
```

- 5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SF Oracle RAC node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

About configuring authentication for SFDB tools

To configure authentication for Storage Foundation for Databases (SFDB) tools, perform the following tasks:

Configure the vxdbd daemon to require authentication

See [“Configuring vxdbd for SFDB tools authentication”](#) on page 186.

Add a node to a cluster that is using authentication for SFDB tools

See [“Adding nodes to a cluster that is using authentication for SFDB tools”](#) on page 490.

Configuring vxdbd for SFDB tools authentication

To configure vxdbd, perform the following steps as the root user

- 1 Run the sfcae_auth_op command to set up the authentication services.

```
# /opt/VRTS/bin/sfae_auth_op -o setup
Setting up AT
Starting SFAE AT broker
Creating SFAE private domain
Backing up AT configuration
Creating principal for vxdbd
```

- 2 Stop the vxdbd daemon.

```
# /opt/VRTS/bin/vxdbdctrl stop
Stopping Veritas vxdbd
vxdbd stop succeeded
```

- 3 Enable authentication by setting the `AUTHENTICATION` key to `yes` in the `/etc/vx/vxdbed/admin.properties` configuration file.

- 4 Start the `vxdbd` daemon.

```
# /opt/VRTS/bin/vxdbdctrl start
Starting Veritas vxdbd
/opt/VRTSdbed/bin/vxdbd start SUCCESS
```

The `vxdbd` daemon is now configured to require authentication.

Configuring Veritas Volume Replicator

Perform this step only if you have not already configured VVR during the installation.

By default, the installer installs the required VVR configuration files irrespective of whether or not you choose to enable VVR. To configure VVR manually in SF Oracle RAC, simply start VVR using the `vxstart_vvr` command. The command starts the VVR daemons and configures the ports. You may change the default settings at any time.

For instructions on changing the default settings, see the *Veritas Volume Replicator Administrator's Guide*.

To configure VVR

- 1 Log into each node in the cluster as the root user.
- 2 Start VVR:

```
# vxstart_vvr start
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmind: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
VxVM VVR V-5-2-5942 Starting Communication daemon: [OK]
```

Running SORT Data Collector to collect configuration information

SORT Data Collector now supersedes the `VRTSexplorer` utility. Run the Data Collector with the `VxExplorer` option to gather information about the system.

Visit the SORT Website and download the UNIX Data Collector appropriate for your operating system.

<https://sort.symantec.com>

For more information:

<https://sort.symantec.com/public/help/wwhelp/wwhimpl/js/html/wwhelp.htm>

Upgrade of SF Oracle RAC

- [Chapter 11. About upgrading SF Oracle RAC 6.0](#)
- [Chapter 12. Performing a full upgrade to SF Oracle RAC 6.0](#)
- [Chapter 13. Performing a phased upgrade to SF Oracle RAC 6.0](#)
- [Chapter 14. Performing a rolling upgrade to SF Oracle RAC 6.0](#)
- [Chapter 15. Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0](#)
- [Chapter 16. Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC](#)
- [Chapter 17. Performing post-upgrade tasks](#)

About upgrading SF Oracle RAC 6.0

This chapter includes the following topics:

- [About types of upgrade](#)
- [Supported upgrade paths](#)

About types of upgrade

SF Oracle RAC supports various ways of upgrading your cluster to the latest version. Choose a method that best suits your environment and supports your planned upgrade path.

[Table 11-1](#) lists the supported types of upgrade.

Table 11-1 Types of upgrade

Type of upgrade	Method of upgrade	Procedures
Full upgrade	Veritas script-based installation programs ■ Interactive mode ■ Non-interactive mode using response files Veritas Web-based installation program	Complete the following steps: ■ Preparing to upgrade See “Preparing to perform a full upgrade to SF Oracle RAC 6.0” on page 194. ■ Upgrading to SF Oracle RAC 6.0 See the chapter <i>Performing a full upgrade to SF Oracle RAC 6.0</i>. ■ Completing post-upgrade tasks See the chapter <i>Performing post-upgrade tasks</i>.

Table 11-1 Types of upgrade (*continued*)

Type of upgrade	Method of upgrade	Procedures
Phased upgrade	Combination of manual steps and the Veritas script-based installation programs Note: SF Oracle RAC does not support phased upgrades using the Web installer.	Complete the steps in the chapter <i>Performing a phased upgrade to SF Oracle RAC 6.0</i> .
Rolling upgrade	Veritas script-based installation programs Veritas Web-based installation program	Complete the steps in the chapter <i>Performing a rolling upgrade to SF Oracle RAC 6.0</i> .

Supported upgrade paths

SF Oracle RAC software must be at the same version across all nodes in an SF Oracle RAC cluster after the upgrade, that is 6.0.

Review the following information before you upgrade:

- If you are running an Oracle RAC version that is not supported in this release, first upgrade Oracle RAC to the supported version before upgrading SF Oracle RAC. For instructions, see the product documentation of the version in use.

[Table 11-2](#) lists the supported upgrade paths for upgrades on RHEL and OEL.

Table 11-2 Supported upgrade paths on RHEL and OEL

From product version	From OS version	To SF Oracle RAC version	To OS version	Supported upgrade type
SF Oracle RAC 5.1 Patch 3	RHEL 5 Update 3 or OEL 5 Update 3 RHEL 5 Update 4 or OEL 5 Update 4	SF Oracle RAC 6.0	RHEL 5 Update 5 or OEL 5 Update 5 and later	Full or rolling upgrade

Table 11-2 Supported upgrade paths on RHEL and OEL (*continued*)

From product version	From OS version	To SF Oracle RAC version	To OS version	Supported upgrade type
SF Oracle RAC 5.1 PR1	RHEL 5 Update 3 or OEL 5 Update 3 RHEL 5 Update 4 or OEL 5 Update 4	SF Oracle RAC 6.0	RHEL 5 Update 5 or OEL 5 Update 5 and later	Full or rolling upgrade
SF Oracle RAC 5.1 SP1	RHEL 5 Update 3 or OEL 5 Update 3 RHEL 5 Update 4 or OEL 5 Update 4	SF Oracle RAC 6.0	RHEL 5 Update 5 or OEL 5 Update 5 and later	Full or rolling upgrade

[Table 11-3](#) lists the supported upgrade paths for upgrades on SLES.

Table 11-3 Supported upgrade paths on SLES

From product version	From OS version	To SF Oracle RAC version	To OS version	Supported upgrade type
SF Oracle RAC 5.1 Patch 3	SLES 10 SP2	SF Oracle RAC 6.0	SLES 10 SP4	Full or phased upgrade
SF Oracle RAC 5.1 PR1	SLES 10 SP2 SLES 10 SP3	SF Oracle RAC 6.0	SLES 10 SP4 SLES 11 SP1	Full or phased upgrade
SF Oracle RAC 5.1 SP1	SLES 10 SP2 SLES 10 SP3 SLES 11	SF Oracle RAC 6.0	SLES 10 SP4 SLES 11 SP1	Full or phased upgrade

Performing a full upgrade to SF Oracle RAC 6.0

This chapter includes the following topics:

- [About full upgrades](#)
- [Preparing to perform a full upgrade to SF Oracle RAC 6.0](#)
- [Pre-upgrade tasks for migrating the SFDB repository database](#)
- [Upgrading to SF Oracle RAC 6.0](#)

About full upgrades

A full upgrade involves upgrading all the nodes in the cluster at the same time. The cluster remains unavailable for the duration of the upgrade.

Note: You can not roll back the upgrade to a previous version after you upgrade to version 6.0.

You can perform the upgrade using one of the following Veritas script-based installation programs:

- **Common product installer** (`installer` or `webinstaller`)
The common product installer provides menu options for installing and configuring multiple Veritas products.
- **SF Oracle RAC installation programs** (`installsfprac`)
The SF Oracle RAC installation programs provide menu options for installing and configuring SF Oracle RAC. You can use either the script-based installer or the Web-based installer.

Note: If you obtained SF Oracle RAC from an electronic download site, you must use the product installer (`installsffrac`) instead of the common product installer (`installer`).

You can also perform a full upgrade using a response file. You can create a response file by using the response file template or by customizing a response file that is generated by the script-based installer.

For more information about response files:

See [“About response files”](#) on page 373.

Preparing to perform a full upgrade to SF Oracle RAC 6.0

Perform the preparatory steps in this section if you are performing a full upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

To prepare to upgrade SF Oracle RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

3 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

4 Stop all Oracle RAC resources.

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

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

- If the database instances are not managed by VCS, then run the following on one node:

```
$ srvctl stop database -d db_name
```

5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

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

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

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

6 Unmount the VxFS file system, which is not under VCS control.

```
# mount |grep vxfs

# fuser -m /mount_point

# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

7 Stop VCS on all nodes:

```
# hastop -all
```

Note: The installer displays the following warning:

VCS is not running before upgrade. Please make sure all the configurations are valid before upgrade.

If the configuration files are valid, you may ignore the message.

One way to check whether or not the configuration is valid is to check the main.cf file as follows:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

However, this method can not verify whether all configurations are valid. If SF Oracle RAC was running properly before the upgrade, the configurations are valid.

8 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 196.

9 If you plan to upgrade the operating system, stop all ports.

If you are running version 5.1 and later, stop the ports using the installer:

```
# ./installsfrc -stop
```

Pre-upgrade tasks for migrating the SFDB repository database

If you plan to continue using Database Storage Checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle RAC, you must migrate the SFDB repository database to 6.0.

If you are upgrading from 5.1 or 5.1SP1 to 6.0, no upgrade steps are required for the SFDB tools. The migration process is automatic.

Note: The Sfua_Base repository resource group will be removed from the main.cf file. It is not required as a separate service group for SF Oracle RAC 6.0.

Perform the following before upgrading SF Oracle RAC.

To prepare to migrate the repository database

- ◆ Resynchronize all existing snapshots before upgrading. As Oracle user, enter:

```
$ /opt/VRTS/bin/dbed_vmsnap -S $ORACLE_SID \  
-f SNAPPLAN -o resync
```

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 6.0.

Upgrading to SF Oracle RAC 6.0

This section provides instructions for the following upgrade scenarios:

SF Oracle RAC and major operating system upgrade	Perform the steps in the following procedure if you plan to perform a major upgrade of the operating system, for example from SLES 10 to SLES 11, along with SF Oracle RAC.

See [“To upgrade SF Oracle RAC and the operating system \(major OS upgrade\)”](#) on page 199.

SF Oracle RAC and minor operating system upgrade	Perform the steps in the following procedure if you plan to perform a minor upgrade of the operating system, for example from SLES 10 SP2 to SLES 10 SP4, along with SF Oracle RAC.

See [“To upgrade SF Oracle RAC and operating system \(minor OS upgrade\)”](#) on page 197.

To upgrade SF Oracle RAC and operating system (minor OS upgrade)

- 1 If you want to upgrade the operating system, perform the following steps:
 - Rename the /etc/llttab file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- Upgrade the operating system on all nodes in the cluster. For instructions, see the operating system documentation.

Note: If reboot is required, use `shutdown -r now` command to reboot the nodes.

- After the system restarts, restore the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 2 Upgrade to SF Oracle RAC 6.0 using the script-based installer or the Web-based installer.

See [“Upgrading SF Oracle RAC using the Veritas script-based installation program”](#) on page 199.

See [“Upgrading Veritas Storage Foundation for Oracle RAC using the Veritas Web-based installer”](#) on page 201.

You can also perform a silent upgrade:

See [“Upgrading SF Oracle RAC using a response file”](#) on page 203.

- 3 Manually mount the VxFS and CFS file systems that are not managed by VCS.

- 4 Relink the SF Oracle RAC libraries with Oracle.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 5 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online Oracle_group -any
```

- If the Oracle database is not managed by VCS:

```
# srvctl start database -d db_name
```

- 6 Start all applications that are not managed by VCS. Use native application commands to start the applications.

- 7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oradb_grpname AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

8 Complete other post-upgrade steps.

For instructions, see the chapter *Performing post-upgrade tasks* in this document.

9 If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

10 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

11 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

To upgrade SF Oracle RAC and the operating system (major OS upgrade)

1 Uninstall SF Oracle RAC.

For instructions, see the chapter *Uninstalling SF Oracle RAC*.

2 Upgrade the operating system.

For instructions, see the operating system documentation.

3 Install SF Oracle RAC 6.0.

For instructions, see the chapter *Installing and configuring SF Oracle RAC*.

4 Complete the post-installation and configuration tasks.

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

Upgrading SF Oracle RAC using the Veritas script-based installation program

Use the `installer` or the `installsfprac` Veritas script-based installation programs to upgrade SF Oracle RAC.

The installer performs the following tasks to upgrade SF Oracle RAC:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Oracle RAC processes before the upgrade.
- Uninstalls SF Oracle RAC.
- Installs the SF Oracle RAC 6.0 RPMs on the nodes.
- Starts SF Oracle RAC 6.0 on all the nodes.
- Displays the location of the log files, summary file, and response file.

To upgrade to SF Oracle RAC 6.0 using the `installsfrac` program

- 1 Start the installation program using one of the following ways:

SF Oracle RAC
installer

Navigate to the product directory on the installation media that contains the installation program.

The program is located in the
`storage_foundation_for_oracle_rac` directory.

Run the program:

```
# ./installsfrac galaxy nebula
```

Common product
installer

Navigate to the product directory on the installation media that contains the installation program.

Run the program:

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

From the opening Selection Menu, choose **G** for "**Upgrade a Product.**"

Select the option **Full Upgrade.**"

Select the option **Veritas Storage Foundation for Oracle RAC.**

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

- 2 Enter **3** to install all the SF Oracle RAC RPMs.

The installer verifies the systems for compatibility.

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

Review the messages displayed and make sure that you meet the requirements before proceeding with the upgrade.

- 3 Press **Enter** to continue with the upgrade.

Depending on the installation option selected, the installer displays the list of RPMs that will be installed.

The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's boot disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer y.

- 4 Enter the name of the backup boot disk group when prompted. Press **Enter** to accept the default.

You are prompted to start the split operation.

- 5 Enter **y** to continue with the split operation.

The split operation can take some time to complete.

Note: Verify the boot device from which the system is set to boot. Make sure that the system is set to start from the boot device with the required version of SF Oracle RAC.

- 6 Enter **y** to stop the SF Oracle RAC processes.

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

The installer stops the processes and uninstalls SF Oracle RAC. After the uninstallation, the installer installs SF Oracle RAC 6.0 and starts SF Oracle RAC 6.0 on all the nodes.

- 7 Complete the remaining tasks to finish the upgrade:

See [“Upgrading to SF Oracle RAC 6.0”](#) on page 197.

Upgrading Veritas Storage Foundation for Oracle RAC using the Veritas Web-based installer

This section describes upgrading SF Oracle RAC with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade SF Oracle RAC

- 1 Perform the required steps to save any data that you wish to preserve. For example, make configuration file backups.
- 2 If you are upgrading a high availability (HA) product, take all service groups offline. List all service groups:

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

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -all
```

- 3 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
- 5 The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the boot disk group. To create the backup, check the **Split mirrors on all the systems** box. Check the appropriate box to use the same name for the backup disk group on all systems--you can use the default name or choose a new one. Check the systems where you want to create the backup. When you are ready, click the **Next** button.
- 6 Click **Next** to complete the upgrade.
After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.
- 7 After the upgrade, if the product is not configured, the Web-based installer asks: "Do you want to configure this product?" If the product is already configured, it will not ask any questions.
- 8 If you want to upgrade VCS or SFHA 5.1 on the CP server systems to version SF Oracle RAC 6.0, make sure that you upgraded all application clusters to version SF Oracle RAC 6.0. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

- 9 Only perform this step if you have split the mirrored root disk to back it up. After a successful reboot, verify the upgrade and re-join the backup disk group into the upgraded boot disk group. If the upgrade fails, revert the upgrade boot disk group to the backup disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 255.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 255.
- 10 Complete the remaining tasks to finish the upgrade:

See [“Upgrading to SF Oracle RAC 6.0”](#) on page 197.

Upgrading SF Oracle RAC using a response file

You can upgrade from SF Oracle RAC version 5.1 PR1 using a response file.

Perform the steps in the following procedure to upgrade to SF Oracle RAC 6.0 using a response file.

To upgrade SF Oracle RAC using a response file

- 1 Upgrade the operating system, if required.

For instructions, see the operating system documentation.
- 2 Create a response file using one of the available options.

Note: Make sure that you replace the host names in the response file with the names of the systems that you plan to upgrade.

For information on various options available for creating a response file:

See [“About response files”](#) on page 373.

For response file variable definitions:

See [“Response file variables to upgrade SF Oracle RAC”](#) on page 204.

For a sample response file:

See [“Sample response file for upgrading SF Oracle RAC”](#) on page 205.

- 3 Navigate to the product directory on the installation media that contains the SF Oracle RAC installation program.

4 Start the installation:

```
# ./installsf rac -responsefile /tmp/response_file
```

Where /tmp/response_file is the full path name of the response file.

5 Complete the post-upgrade steps.

See the chapter "Performing post-upgrade tasks" in this document.

Response file variables to upgrade SF Oracle RAC

Table 12-1 lists the response file variables that you can define to upgrade SF Oracle RAC.

Table 12-1 Response file variables specific to upgrading SF Oracle RAC

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades SF Oracle RAC RPMs. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be upgraded. (Required)
CFG{prod}	Scalar	Defines the product to be upgraded. (Optional)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)

Table 12-1 Response file variables specific to upgrading SF Oracle RAC
(continued)

Variable	List or Scalar	Description
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)

Sample response file for upgrading SF Oracle RAC

The following sample response file upgrades SF Oracle RAC to version 6.0 on two nodes, galaxy and nebula.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_allowcomms}=1;
```

1;

Performing a phased upgrade to SF Oracle RAC 6.0

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing phased upgrade of SF Oracle RAC from version 5.1 Patch 3 and later releases](#)

About phased upgrade

The phased upgrade methodology involves upgrading half of the nodes in the cluster at a time.

For supported upgrade paths:

See [“Supported upgrade paths”](#) on page 191.

Note: You can perform a phased upgrade only if the Oracle RAC binaries are present on the local file system.

Caution: There is a potential for dependency problems between product components that no longer match when upgrading part of a cluster at a time. Follow the phased upgrade procedures carefully to avoid these problems.

Note: There will be some downtime involved. Review the procedures and carefully plan your downtime before proceeding with any steps. The sample procedures assume that Oracle RAC binaries are installed on local file systems for each node in the cluster.

The examples in the procedures assume a four-node SF Oracle RAC cluster with the nodes *galaxy* and *nebula* constituting the first half of the cluster and the nodes *jupiter* and *mercury* constituting the second half of the cluster.

Performing phased upgrade of SF Oracle RAC from version 5.1 Patch 3 and later releases

Table 13-1 illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

Table 13-1 Summary of phased upgrade

First half of the cluster	Second half of the cluster
SF Oracle RAC cluster before the upgrade:	
First half of the cluster galaxy nebula Second half of the cluster jupiter mercury	
STEP 1: Perform the following pre-upgrade steps: - Switch failover applications. - Stop all parallel applications. See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 209. STEP 2: Upgrade SF Oracle RAC. See “Step 2: Upgrading the first half of the cluster” on page 211.	The second half of the cluster is up and running. jupiter mercury

Table 13-1 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
The first half of the cluster is not running.  galaxy nebula	STEP 3: Perform the following pre-upgrade steps: ■ Stop all parallel and failover applications. ■ Stop SF Oracle RAC. See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 213. The downtime starts now.
STEP 4: Perform the following post-upgrade steps: ■ Start SF Oracle RAC. ■ Start all applications. See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 214. The downtime ends here.	The second half of the cluster is not running.  jupiter mercury
The first half of the cluster is up and running.  galaxy nebula	STEP 5: Upgrade SF Oracle RAC. See “Step 5: Upgrading the second half of the cluster” on page 215. STEP 6: Perform the following post-upgrade steps: ■ Start SF Oracle RAC. ■ Start all applications. See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 216.
The phased upgrade is complete and both the first and the second half of the cluster are running.  galaxy nebula jupiter mercury	

Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

- 2 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.
- 3 Stop the applications configured under VCS. Stop the Oracle database:

- If the Oracle RAC instance is managed by VCS:

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

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

For Oracle RAC 11.2.0.2:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 4 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

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

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```

- 5 Unmount the VxFS and CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs

# fuser -cu /mount_point
```

- Unmount the non-system VxFS or CFS file system:

```
# umount /mount_point
```

- 6 Stop any open volumes that are not managed by VCS.

- 7 Stop the parallel service groups and switch over failover service groups:

```
# hastop -local -evacuate
```

- 8 If you plan to upgrade the operating system, stop all ports.

```
# /opt/VRTS/install/installsfrac -stop galaxy nebula
```

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 6 Upgrade SF Oracle RAC. Navigate to the product directory on the installation media. When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

For SLES 11 (x86_64)

```
# cd /dvd_mount/sles11_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
```

For SLES 10 (x86_64)

```
# cd /dvd_mount/sles10_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
# ./installsfrac -upgrade galaxy nebula
```

Note: After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop all applications that are configured under VCS. Stop the Oracle database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys jupiter
# hagrps -offline oracle_group -sys mercury
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

For Oracle RAC 11.2.0.2:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 3 Unmount the VxFS or CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs

# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

4 Stop any open volumes that are not managed by VCS.

5 Stop VCS:

```
# hastop -local
```

6 If you plan to upgrade the operating system, stop all ports.

```
# /opt/VRTS/install/installsfrac -stop jupiter mercury
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
```

```
# /etc/init.d/gab start
```

```
# gabconfig -x
```

2 On the first half of the cluster, start SF Oracle RAC:

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

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

Verify that the GAB ports a, b, d, o and h appear in the `gabconfig -a` command output.

3 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.

4 Relink the SF Oracle RAC libraries with Oracle.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 5 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

Note: The downtime ends here.

- 6 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 On the second half of the cluster, upgrade SF Oracle RAC. Navigate to the product directory on the installation media.

When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

For SLES 11 (x86_64)

```
# cd /dvd_mount/sles11_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
```

For SLES 10 (x86_64)

```
# cd /dvd_mount/sles10_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
```

```
# ./installsfrac -upgrade jupiter mercury
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 Relink the SF Oracle RAC libraries with Oracle.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 3 On the second half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install
# ./installsfrac -start jupiter mercury
```

- 4 Upgrade VxVM disk group version.

See [“Upgrading CVM protocol version and VxVM disk group version ”](#) on page 257.

- 5 Upgrade disk layout version.

See [“Upgrading disk layout versions”](#) on page 256.

- 6 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys jupiter
# hagrps -online oracle_group -sys mercury
```

If the Oracle database is not managed by VCS:

For Oracle RAC 11.2.0.2:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

- 7
 - If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

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

- 8 Start all applications that are not managed by VCS. Use native application commands to start the applications.

- 9 Set or change the product license level, if required.

See [“Setting or changing the product license level”](#) on page 256.

- 10 Migrate the SFDB repository database.

See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 257.

11 Verify the cluster.

See [“Verifying the cluster”](#) on page 260.

12 If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Performing a rolling upgrade to SF Oracle RAC 6.0

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Preparing to perform a rolling upgrade to SF Oracle RAC 6.0](#)
- [Performing a rolling upgrade using the installer](#)

About rolling upgrades

The rolling upgrade minimizes downtime for highly available clusters to the amount of time that it takes to perform a service group failover. The rolling upgrade has two main phases where the installer upgrades kernel RPMs in phase 1 and VCS and VCS agent RPMs in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.


The following is an overview of the flow for a rolling upgrade:

1. The installer performs prechecks on the cluster.
2. The installer moves service groups to free nodes for the first phase of the upgrade as is needed. Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over.

3. The installer performs the second phase of the upgrade on all of the nodes in the cluster. The second phase of the upgrade includes downtime of the Veritas Cluster Server (VCS) engine HAD, but does not include application downtime.

[Figure 14-1](#) illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 14-1 Example of the installer performing a rolling upgrade



The following limitations apply to rolling upgrades:

- Rolling upgrades are not compatible with phased upgrades.
- Do not mix rolling upgrades and phased upgrades.

Preparing to perform a rolling upgrade to SF Oracle RAC 6.0

Perform the preparatory steps in this section if you are performing a rolling upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

Note: Perform the steps on the first subcluster.

To prepare to upgrade SF Oracle RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

- 3 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -sys node_name
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 4 Stop the Oracle RAC resources on each node.
 - If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrpf -offline grp_name -sys node_name
```

- If the database instances are not managed by VCS, then run the following on one node:
 For Oracle RAC 11.2.0.2:

```
$ srvctl stop instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \  
-i instance_name
```

- 5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw  
# hagrpf -modify oracle_group AutoStart 0  
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```

- 6 Switch over all failover service groups to the other nodes in the cluster:

```
# hagrpf -switch grp_name -to sys_name
```

- 7 Take all the VCS service groups offline:

```
# hagrpf -offline grp_name -sys sys_name
```

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

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

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 9 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 196.

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Veritas Storage Foundation for Oracle RAC to the latest release with minimal application downtime.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To perform a rolling upgrade

- 1 If you want to upgrade the operating system, perform the following steps:
 - Change to the `/opt/VRTS/install` directory on the node where you want to upgrade the operating system:

```
# cd /opt/VRTS/install
```
 - Stop SF Oracle RAC:


```
# ./installsfrac -stop
```

- Upgrade the operating system. For instructions, see the operating system documentation.
- Reboot the nodes:

```
# shutdown -r now
```

- 2 Complete the preparatory steps on the first sub-cluster.

See [“Preparing to perform a rolling upgrade to SF Oracle RAC 6.0”](#) on page 222.

- 3 Log in as superuser and mount the SF Oracle RAC SF Oracle RAC 6.0 installation media.
- 4 From root, start the installer.

```
# ./installer
```

- 5 From the menu select **G) Upgrade a product; then select 2)Rolling Upgrade;.**
- 6 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 7 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 8 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 9 Review the EULA, and type **y** if you agree to its terms.
- 10 If the upgrade is from the 5.1 SP1 release or later and the boot disk is encapsulated and mirrored, you can create a backup boot disk.

If you choose to create a backup boot disk, type **y**. Provide a backup name for the boot disk group or accept the default name. The installer then creates a backup copy of the boot disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 255.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 255.

The installer lists the RPMs to upgrade on the selected node or nodes.

- 11 The installer prompts you to stop the applicable processes. Type **y** to continue.

The installer fails over failover service groups to the node or nodes that are not upgraded at this time. The downtime is the time that it normally takes for the service group's failover. The installer stops parallel service groups on the nodes that are to be upgraded.

- 12 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs.

It performs the configuration for the upgrade and re-starts processes.

In case of failure in the startup of some of the processes, you may need to reboot the nodes and manually check the cluster's status.

Note: The Oracle service group is offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

- 13 Relink the SF Oracle RAC libraries with Oracle by choosing the option **Relink Oracle Database Binary** from the program menu.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 14 If the boot disk is encapsulated, the installer strongly recommends a reboot of the nodes. Reboot the nodes as prompted by the installer.

Note: Before you reboot the nodes, ensure that the boot device is set to the disk containing the upgraded version of the product.

```
# eeprom
```

- 15 Manually mount the VxFS and CFS file systems that are not managed by VCS.

- 16 Bring the Oracle database service group online.

- If VCS manages the Oracle database:

```
# hagrps -online oracle_group -sys node_name
```

- If VCS does not manage the Oracle database:

```
# srvctl start database -d db_name
```

- 17 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 18 Complete the preparatory steps on the nodes that you have not yet upgraded. See [“Preparing to perform a rolling upgrade to SF Oracle RAC 6.0”](#) on page 222.
- 19 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

If the installer reboots nodes, restart the installer.

The installer repeats step 7 through step 17.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

This completes phase 1 of the upgrade.
- 20
 - If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```
 - If VCS does not manage the Oracle database, change the management policy for the database to automatic:

```
$ srvctl modify database -d db-name -y AUTOMATIC
```
- 21 Migrate the SFDB repository database.

See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 257.
- 22 Phase 2 of the upgrade begins here. This phase includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.
- 23 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.
- 24 The installer stops Veritas Cluster Server (VCS) processes. Type **y** to continue.

The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.
- 25 Type **y** or **n** to help Symantec improve the automated installation.

- 26 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.
- 27 To upgrade VCS or Storage Foundation High Availability (SFHA) on the Coordination Point (CP) server systems to version 6.0, upgrade all the application clusters to 6.0. You then upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of SF Oracle RAC using the Web-based installer

This section describes using the Veritas Web-based installer to perform a rolling upgrade. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

See [“About rolling upgrades”](#) on page 219.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To start the rolling upgrade—phase 1

- 1 To do an optional OS upgrade, stop the application or mount point manually if it is outside of VCS.

```
# hastop -local
```

```
# installsfrac -stop first half of nodes
```

Perform the OS upgrade. Refer to the documentation for your OS. A reboot is required.

- 2 Complete the preparatory steps on the first sub-cluster.

See [“Preparing to perform a rolling upgrade to SF Oracle RAC 6.0”](#) on page 222.
- 3 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 4 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 86.

- 5 In the Task pull-down menu, select G) Upgrade a product; then select 2) Rolling Upgrade;.

Click the **Next** button to proceed.

- 6 Review the systems that the installer has chosen to start the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

Click **Yes** to proceed.

The installer validates systems. If it throws an error, address the error and return to the installer.

- 7 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.

- 8 If the upgrade is from the 5.1 SP1 release or later and the boot disk is encapsulated and mirrored, you can create a backup boot disk.

If you choose to create a backup boot disk, type **y**. Provide a backup name for the boot disk group or accept the default name. The installer then creates a backup copy of the boot disk group.

If you choose to create a backup boot disk, provide the name of the backup disk group or accept the default name.

Additionally, select the following options:

- Split mirrors on all the systems
- Use the same disk group name on all the mirrored systems

Select the systems where you want to create a backup of the boot disk.

Click **Next**. The installer then creates a backup copy of the boot disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 255.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 255.

The installer lists the RPMs to upgrade on the selected node or nodes.

- 9 The installer stops all processes. Click **Next** to proceed.
- 10 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. The installer starts all the relevant processes and brings all the service groups online.

If the installer reboots nodes, restart the installer.

- 11 Relink the SF Oracle RAC libraries with Oracle by choosing the option **Relink Oracle Database Binary** from the program menu.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 12 If the boot disk is encapsulated, the installer strongly recommends a reboot of the nodes. Reboot the nodes as prompted by the installer.

Note: Before you reboot the nodes, ensure that the boot device is set to the disk containing the upgraded version of the product.

`eeeprom`

- 13 Complete the preparatory steps on the nodes that you have not yet upgraded. See [“Preparing to perform a rolling upgrade to SF Oracle RAC 6.0”](#) on page 222.

- 14 When prompted, perform step 5 through step 10 on the nodes that you have not yet upgraded.

- 15 Relink the SF Oracle RAC libraries with Oracle on the nodes that you have not yet upgraded by choosing the option **Relink Oracle Database Binary** from the program menu.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

- 16 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade. You may need to restart the Web-based installer to perform phase 2. See [“Starting the Veritas Web-based installer”](#) on page 86.

To upgrade the non-kernel components—phase 2

- 1 In the Task pull-down menu, make sure that **Rolling Upgrade** is selected. Click the **Next** button to proceed.
- 2 The installer detects the information of cluster and the state of rolling upgrade. The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.
- 3 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 4 The installer validates systems. If it throws an error, address the error and return to the installer. Click **Next** to proceed.
- 5 The installer stops all processes. Click **Next** to proceed.

- 6 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. The installer starts all the relevant processes and brings all the service groups online.
- 7 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.

The upgrade is complete.

See [“Verifying the cluster”](#) on page 260.

Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0

This chapter includes the following topics:

- [About upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0](#)
- [Upgrading from Storage Foundation Cluster File System for Oracle RAC to SF Oracle RAC 6.0](#)
- [Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0](#)

About upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0

You can upgrade from earlier versions of the following Symantec high availability products to SF Oracle RAC 6.0:

- Storage Foundation High Availability
- Storage Foundation Cluster File System
- Storage Foundation Cluster File System for Oracle RAC

Note: The installer does not support direct upgrade from earlier versions of the products, for example from Storage Foundation Cluster File System 5.0, to SF Oracle RAC 6.0. You must first upgrade to version 6.0 of the installed product, then install SF Oracle RAC 6.0.

If you are running SFCFS RAC versions, you need to uninstall the version and install SF Oracle RAC 6.0.

Upgrading from Storage Foundation Cluster File System for Oracle RAC to SF Oracle RAC 6.0

This release of the product does not support a direct upgrade from SFCFS RAC versions to SF Oracle RAC 6.0. You must uninstall the SFCFS RAC version and install SF Oracle RAC 6.0.

Perform the steps in the following procedure to upgrade SFCFS RAC versions to SF Oracle RAC 6.0.

To upgrade SFCFS RAC versions to SF Oracle RAC 6.0

- 1 Uninstall the current version of SFCFS RAC.

For instructions, see the installation guide of the product.

- 2 Install and configure SF Oracle RAC 6.0.

See *Section 2: Installation and Configuration of SF Oracle RAC* in this document.

- 3 Install Oracle RAC.

See *Section 4: Installation and Upgrade of Oracle RAC* in this document.

Upgrading from Storage Foundation High Availability products to SF Oracle RAC 6.0

Perform the steps in the following procedure to upgrade Storage Foundation High Availability products to SF Oracle RAC 6.0.

To upgrade Storage Foundation High Availability products to SF Oracle RAC 6.0

- 1 If you are running an earlier version of the product, upgrade to version 6.0 of the product.

For instructions, see the installation guide of the product.

- 2 Install SF Oracle RAC 6.0.

See “[About installing SF Oracle RAC](#)” on page 80.

- 3 When you are prompted to configure SF Oracle RAC, enter **y** at the prompt:

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

The program menu is displayed.

- 4 Select the option **Configure SF Oracle RAC sub-components**.

If VCS is configured in your environment, the installation program detects the VCS configuration.

Note: Do not reconfigure VCS when you are prompted by the installation program:

```
Do you want to re-configure VCS? [y,n,q] (n)
```

The installation program creates the VCS configuration file and the `/etc/vcsmmtab` file and starts all processes.

- 5 Verify that all GAB ports are up and running.

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

```
Port a gen ada401 membership 01
```

```
Port b gen ada40d membership 01
```

```
Port d gen ada409 membership 01
```

```
Port f gen ada41c membership 01
```

```
Port h gen ada40f membership 01
```

```
Port o gen ada406 membership 01
```

```
Port u gen ada41a membership 01
```

```
Port v gen ada416 membership 01
```

```
Port w gen ada418 membership 01
```

```
Port y gen ada42a membership 01
```

- 6 Restart ODM on all nodes. This will start ODM in clustered mode.

```
# /etc/init.d/vxodm stop
# /etc/init.d/vxodm start
```

- 7 If fencing is not configured, configure fencing.

For instructions, see the chapter *Configuring SF Oracle RAC clusters for data integrity* in this document.

- 8
 - If you have neither single-instance Oracle nor Oracle RAC running in your environment, install Oracle RAC.
For instructions, see the chapter *Installing and configuring Oracle RAC* in this document.
 - If you have a single-instance Oracle database in your environment, migrate the database to Oracle RAC.
For instructions, see the chapter *Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC* in this document.
 - If you have an Oracle RAC database in your environment, upgrade the database to a higher Oracle RAC version, if necessary.
For instructions, see the chapter *Upgrading Oracle RAC* in this document.

- 9 If you are upgrading from SFCFS RAC to SF Oracle RAC, manually set the `LLT PEERINACT` value to 16 seconds:

```
# lltconfig -T peerinact:1600
```

Verify that the value is set to 16 seconds:

```
# lltconfig -T query
```

Remove the following line from the `/etc/llttab` file:

```
set-timer peerinact:15000
```

Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC

This chapter includes the following topics:

- [Migration overview](#)
- [Migration requirements](#)
- [Before you migrate](#)
- [Migrating to SF Oracle RAC 6.0](#)
- [Sample configuration files](#)


Migration overview

This chapter provides instructions for migrating from Veritas Storage Foundation for Oracle High Availability to Veritas Storage Foundation for Oracle RAC (SF Oracle RAC).

Note: For Oracle RAC 11g Release 2: The instructions in this chapter support migration to administrator-managed databases only.

[Figure 16-1](#) illustrates the migration from Storage Foundation for Oracle HA to SF Oracle RAC.

Figure 16-1 Migration from Storage Foundation for Oracle HA to SF Oracle RAC



Sample configuration before and after migration

Table 16-1 describes a sample existing configuration in a Storage Foundation for Oracle HA database environment.

Table 16-1 Configuration in a Storage Foundation for Oracle HA database environment

Host name	Instance name	Database name	Oracle home	Database file storage
galaxy	vrts	vrts	/u01/app/oracle/ product/ <i>release</i> /dbhome_1 where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version	Shared storage using VxFS
nebula	vrts	vrts	/u01/app/oracle/ product/ <i>release</i> /dbhome_1 where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version	Shared storage using VxFS

Table 16-2 describes a sample configuration in an SF Oracle RAC database environment.

Table 16-2 Configuration in an SF Oracle RAC environment

Host name	Instance name	Database name	Oracle home	Database file storage	OCR and voting disk
galaxy	vrts1	vrts	/u01/app/oracle/ product/ <i>release</i> /dbhome_1 where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version	Shared storage using CFS	Resides on CFS
nebula	vrts2	vrts	/u01/app/oracle/ product/ <i>release</i> /dbhome_1 where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version	Shared storage using CFS	Resides on CFS

Migration requirements

Make sure that you meet the following requirements before migrating to SF Oracle RAC:

- Storage Foundation for Oracle HA version 6.0 is installed on the systems.
- Oracle 10g Release 2 or Oracle 11g Release 2 is installed on the systems.

Note: In the case of Oracle RAC 11g Release 1, only the database is supported.

Before you migrate

Before you migrate to SF Oracle RAC, complete the following tasks:

- Set up the single instance Storage Foundation for Oracle database with the following configuration:
 - Storage Foundation for Oracle HA installed on all nodes
 - Oracle binaries installed on each node
 - Oracle database created on shared storage using Veritas File System (VxFS)
- Back up the existing database before the migration.

Migrating to SF Oracle RAC 6.0

Perform the following steps to migrate from Storage Foundation for Oracle HA to SF Oracle RAC.

1. Migrate Storage Foundation for Oracle HA to SF Oracle RAC.
 See [“Migrating Storage Foundation for Oracle HA to SF Oracle RAC”](#) on page 239.
2. Migrate single instance Oracle database to Oracle RAC database.
 See [“Migrating a single instance Oracle database to Oracle RAC database”](#) on page 240.
3. Complete the post-migration tasks.
 See [“Completing post-migration tasks”](#) on page 243.

Migrating Storage Foundation for Oracle HA to SF Oracle RAC

Perform the following steps to migrate from Storage Foundation for Oracle HA to SF Oracle RAC.

To migrate from Storage Foundation for Oracle HA to SF Oracle RAC

- 1 Log in as a superuser.
- 2 Back up the existing Storage Foundation for Oracle database HA resource configuration:

```
# cp -rp /etc/VRTSvcs/conf/config config.old
```

- 3 Take the database service groups offline:

```
# hagr -offline group_name -any
```

- 4 Unmount the VxFS mount points that are not managed by VCS:

```
# umount mount_point
```

- 5 Stop all the other VCS service groups.

To view the current state of the service groups:

```
# hagr -state
```

To stop each group:

```
# hagr -offline servicegroup -sys node_name
```

6 Freeze the VCS service groups:

```
# haconf -makerw
# hagrps -freeze servicegroup -persistent
# haconf -dump -makero
```

7 Stop VCS on all nodes:

```
# hastop -all -force
```

8 Install and configure SF Oracle RAC.

For information on installing and configuring SF Oracle RAC, see *Section 2: Installation and configuration of SF Oracle RAC* in this guide.

9 Unfreeze the VCS service groups:

```
# haconf -makerw
# hagrps -unfreeze servicegroup -persistent
# haconf -dump -makero
```

Migrating a single instance Oracle database to Oracle RAC database

Complete the steps in the following procedure for migrating a single instance Oracle database to the Oracle RAC database.

To migrate a single instance Oracle database to Oracle RAC database

1 Install Oracle Clusterware.

For information on installing Oracle Clusterware, see *Section 4: Installation and upgrade of Oracle RAC* in this guide.

2 Import the single instance Storage Foundation for Oracle database HA storage disk group in shared mode:

```
# vxdbg -s import oradatadg
```

where oradatadg is a disk group in the Storage Foundation for Oracle database HA environment.

In an Oracle RAC environment, all instances concurrently access a single database. Thus, all datafiles, control files, SPFILE, redo log files and archive log files must reside on shared storage.

3 Mount the file system in cluster mode on all nodes:

```
# mount -t vxfs -o cluster \
/dev/vx/dsk/oradatadg/oradatavol1 /oradata
```

where oradata is a file system in the Storage Foundation for Oracle database HA environment.

4 Relink Oracle binaries with Veritas libraries.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 253.

5 Start the database as an Oracle user from any one of the nodes in the cluster using SQLPLUS:

```
$ export ORACLE_SID=vrts
$ export ORACLE_HOME=/u02/app/oracle/product/11.2.0/dbhome_1
$ /u02/app/oracle/product/11.2.0/dbhome_1/bin/sqlplus "/as sysdba"
```

where vrts is a database in the Storage Foundation for Oracle database HA environment.

6 Add redo logs. Each instance requires its own redo thread. The following example assumes a two node configuration, with galaxy and nebula as the nodes in the cluster. The existing redo thread will be used by galaxy.

To add a new redo thread, log on to nebula as Oracle user and run the following commands:

```
SQL> alter database add logfile thread 2
      group 4 ('/oradata/vrts/redo04.log') size 50M REUSE;
Database altered
```

```
SQL> alter database add logfile thread 2
      group 5 ('/oradata/vrts/redo05.log') size 50M REUSE;
Database altered.
```

```
SQL> alter database add logfile thread 2
      group 6 ('/oradata/vrts/redo06.log') size 50M REUSE;
Database altered
```

7 Enable redo log thread. While enabling, it may be designated as a public thread or a private thread:

```
SQL> ALTER DATABASE ENABLE PUBLIC THREAD 2;
OR
SQL> ALTER DATABASE ENABLE PRIVATE THREAD 2;
```

8 Add UNDO tablespaces for each additional instance:

```
SQL> create undo tablespace UNDOTBS2 datafile\
'/oradata/vrts/undotbs02.dbf' size 500M autoextend on;
```

9 Create the cluster views needed for Oracle RAC:

```
SQL> @?/rdbms/admin/catclust
```

10 If you are using an SPFILE, create a PFILE from it:

```
SQL> create pfile='/tmp/initORA.ora' from spfile
```

11 Edit the Oracle initialization parameter file /tmp/initORA.ora to include cluster parameters:

```
*.cluster_database_instances=2
*.cluster_database=TRUE
vrts1.instance_name=vrts1
vrts2.instance_name=vrts2
vrts1.instance_number=1
vrts2.instance_number=2
vrts1.thread=1
vrts2.thread=2
vrts1.undo_tablespace='UNDOTBS1'
vrts2.undo_tablespace='UNDOTBS2'
vrts1.local_listener= 'LISTENER_GALAXY'
vrts2.local_listener= 'LISTENER_NEBULA'
vrts1.remote_listener= 'LISTENERS_VRTS'
vrts2.remote_listener= 'LISTENERS_VRTS'
```

12 Start the Oracle initialization parameter file:

```
SQL> connect / as sysdba
SQL> shutdown immediate
SQL> startup pfile='/tmp/initORA.ora'
```

13 If the database starts successfully using the Oracle initialization parameter file, create an spfile in a shared location:

```
SQL> create spfile='sharedlocation/spfiledbrac.ora'\
from pfile='/tmp/initORA.ora'
```

- 14 On each node create a link in \$ORACLE_HOME/dbs to the shared spfile directory.

For galaxy:

```
cd $ORACLE_HOME/dbs
ln -s $sharedlocation/spfiledbrac.ora spfiledbrac1.ora
```

For nebula:

```
cd $ORACLE_HOME/dbs
ln -s $sharedlocation/spfiledbrac.ora spfiledbrac2.ora
```

- 15 Add the database and instances to the cluster registry:

```
# srvctl add database -d vrts -o $ORACLE_HOME
# srvctl add instance -d vrts -i vrts1 -n galaxy
# srvctl add instance -d vrts -i vrts2 -n nebula
```

- 16 Stop and start the database:

```
# srvctl stop database -d vrts
# srvctl start database -d vrts
```

Completing post-migration tasks

Perform the steps in the following procedure to complete the migration process. The examples in the procedures assume a two-node SF Oracle RAC cluster with the nodes galaxy and nebula.

To complete the post-migration tasks:

- 1 For Oracle RAC 11g Release 2, see the Oracle documentation to update the `tnsnames.ora` file.
- 2 Optional: Add services to a database and assign them to instances. You can create services manually or by using the Database Configuration Assistant (DBCA) utility.

In Real Application Clusters (RAC), a service can span one or more instances and facilitate real workload balancing based on real transaction performance. You can control the instances in your cluster that are allocated to different services at different times.

To create the services manually:

```
# srvctl add service -d vrts -s ORA_TAF -r vrts1, vrts2
```

- 3 Configure the CSSD, PrivNIC or MultiPrivNIC and Oracle resource agents in the VCS configuration file.

For more information see *Section 4: Installation and upgrade of Oracle RAC* in this guide.

- 4 Depending on your deployment needs you may perform the following optional tasks:
 - Update the application (database client) configuration to connect to Oracle RAC database (if required) and to take advantage of Oracle RAC features like load balancing and scalability.
 - Update the VCS configuration file for any additional resources, which were present in the VCS configuration prior to the migration process.

Migrating Storage Foundation for Databases (SFDB) objects from single instance Oracle to Oracle RAC

After migrating from Storage Foundation HA to Storage Foundation for Oracle RAC, the Storage Foundation for Databases (SFDB) objects created in a single instance Oracle environment will not be available in the Oracle RAC environment.

To re-create the SFDB repository

- 1 Run the following command:

```
$ /opt/VRTS/bin/dbed_update -s $ORACLE_SID -h $ORACLE_HOME
```

- 2 Re-create the SFDB objects that existed in the Storage Foundation HA environment. They will now be supported for SFRAC.

For information on configuring SFDB objects, see the *Storage Foundation: Storage and Availability Management for Oracle Databases* guide.

Sample configuration files

This section illustrates sample configurations for the following files:

VCS configuration file

See [“VCS configuration file for Storage Foundation for Oracle HA”](#) on page 245.

See [“VCS configuration file for SF Oracle RAC”](#) on page 248.

Oracle initialization parameter file	See “Oracle initialization parameter file for Storage Foundation for Oracle HA” on page 246. See “Oracle initialization parameter file for SF Oracle RAC” on page 251.
tnsnames.ora	See “tnsnames.ora for Storage Foundation for Oracle HA” on page 247. See “tnsnames.ora file for SF Oracle RAC” on page 252.
listener.ora	See “listener.ora for Storage Foundation for Oracle HA” on page 247. See “listener.ora file for SF Oracle RAC” on page 252.

VCS configuration file for Storage Foundation for Oracle HA

The sample VCS configuration file for Storage Foundation for Oracle HA is as follows:

```
include "types.cf"
include "OracleTypes.cf"

cluster sfha_clus (
  UserNames = { admin = bQRjQLqNRmRRpZRlQO }
  Administrators = { admin }
)

system galaxy (
)

system nebula (
)

group sfora (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoStartList = { galaxy, nebula }
)

DiskGroup oradatadg (
  Critical = 0
  DiskGroup = oradatadg
)
```

```

)

Mount oradatamnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/datavol"
    FSType = vxfs
    FsckOpt = "-n"
)

Oracle oradb (
    Critical = 0
    Sid = vrts
    Owner = oracle
    Home = "/oracle/dbhome"
)

oradatamnt requires oradatadg
oradb requires oradatamnt

```

Oracle initialization parameter file for Storage Foundation for Oracle HA

The sample Oracle initialization parameter file for Storage Foundation for Oracle HA is as follows:

```

vrts.__db_cache_size=1375731712
vrts.__java_pool_size=16777216
vrts.__large_pool_size=16777216
vrts.__oracle_base='/u02/app/oracle/product/11.2.0/dbhome_1\'
#ORACLE_BASE set from environment
vrts.__pga_aggregate_target=1224736768
vrts.__sga_target=1811939328
vrts.__shared_io_pool_size=0
vrts.__shared_pool_size=385875968
vrts.__streams_pool_size=0
*.audit_file_dest='/u02/app/oracle/product/11.2.0/dbhome_1\'
/admin/vrts/adump'
*.audit_trail='none'
*.compatible='11.2.0.0.0'
*.control_files='/oradata/vrts/control01.ctl',
'/oradata/vrts/control02.ctl',
'/oradata/vrts/control03.ctl'
*.db_block_size=8192

```

```
*.db_domain=''
*.db_name='vrts'
*.diagnostic_dest='/u02/app/oracle/product/11.2.0/dbhome_1'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=dbracXDB)'
*.log_archive_dest_1='LOCATION=/oradata/archive'
*.log_archive_format='%t_%s_%r.dbf'
*.memory_target=3036676096
*.open_cursors=300
*.processes=150
*.remote_login_passwordfile='EXCLUSIVE'
*.undo_tablespace='UNDOTBS1'
```

tnsnames.ora for Storage Foundation for Oracle HA

The sample tnsnames.ora for Storage Foundation for Oracle HA is as follows:

```
Host : galaxy
VRTS =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = vrts)
    )
  )

Host : Nebula
VRTS =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = nebula) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = vrts)
    )
  )
```

listener.ora for Storage Foundation for Oracle HA

The sample listener.ora for Storage Foundation for Oracle HA is as follows:

Host : galaxy

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )
```

Host : Nebula

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = TCP) (HOST = nebula) (PORT = 1521))
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )
```

VCS configuration file for SF Oracle RAC

The sample VCS configuration file for SF Oracle RAC is as follows:

```
include "OracleASMTTypes.cf"
include "types.cf"
include "CFSTypes.cf"
include "CRSResource.cf"
include "CVMTTypes.cf"
include "Db2udbTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SybaseTypes.cf"

cluster rac_cluster101 (
  UserNames = { admin = enoGniNkoJooMwoInl }
  Administrators = { admin }
  UseFence = SCSI3
  HacliUserLevel = COMMANDROOT
)
```



```

system galaxy (
)

system nebula (
)

group ora_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy , nebula }
)

Oracle ora_res (
    Critical = 0
    Sid @galaxy = "vrts1"
    Sid @nebula = "vrts2"
    Owner = oracle
    Home = "/u02/app/product/orabase/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol1"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol1 }
    CVMActivation = sw
)

requires group cvm online local firm
ora_res requires oradata_mnt
oradata_mnt requires oradata_voldg

```

```

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"
    OnlineRetryLimit = 20
    OnlineWaitLimit = 5
    Critical = 0
)

CFSfsckd vxfsckd (
)

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

CFSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CVMVxconfigd cvm_vxconfigd (
    CVMVxconfigdArgs = { syslog }
)

```

```
PrivNIC ora_priv (
    Critical = 0
    Device = { eth1 = 0, eth2 = 1 }
    Address @galaxy = "192.168.1.128"
    Address @nebula = "192.168.1.131"
    NetMask = "255.255.255.0"
)
```

```
cssd requires ocrvote_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
```

Oracle initialization parameter file for SF Oracle RAC

The sample Oracle initialization parameter file for SF Oracle RAC is as follows:

```
vrts2.__db_cache_size=331350016
vrts1.__db_cache_size=331350016
vrts2.__java_pool_size=4194304
vrts1.__java_pool_size=4194304
vrts2.__large_pool_size=4194304
vrts1.__large_pool_size=4194304
vrts2.__pga_aggregate_target=385875968
vrts1.__pga_aggregate_target=385875968
vrts2.__sga_target=578813952
vrts1.__sga_target=578813952
vrts2.__shared_io_pool_size=0
vrts1.__shared_io_pool_size=0
vrts2.__shared_pool_size=226492416
vrts1.__shared_pool_size=226492416
vrts2.__streams_pool_size=0
vrts1.__streams_pool_size=0
*.audit_file_dest='/u02/app/product/orabase/admin/vrts/adump'
*.audit_trail='db'
*.cluster_database=true
*.compatible='11.2.0.0.0'
*.control_files='/oradata/vrts/control01.ctl','/oradata/vrts/control02.ctl'
*.db_block_size=8192
```

```
*.db_domain=''
*.db_name='vrts'
*.diagnostic_dest='/u02/app/product/orabase'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=vrtsXDB)'
vrts1.instance_number=1
vrts2.instance_number=2
*.memory_target=963641344
*.open_cursors=300
*.processes=150
*.remote_listener='cert-s
can-vip01:1521'
*.remote_login_passwordfile='exclusive'
vrts2.thread=2
vrts1.thread=1
vrts1.undo_tablespace='UNDOTBS1'
vrts2.undo_tablespace='UNDOTBS2'
```

tnsnames.ora file for SF Oracle RAC

The sample tnsnames.ora file for SF Oracle RAC is as follows:

```
VRTS =
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP) (HOST = scan-vip01) (PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = vrts)
  )
)
```

listener.ora file for SF Oracle RAC

The sample listener.ora file for SF Oracle RAC is as follows:

```
LIST1=(DESCRIPTION=(ADDRESS_LIST=\
  (ADDRESS=(PROTOCOL=IPC) (KEY=LIST1))))
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LIST1=ON
```

Performing post-upgrade tasks

This chapter includes the following topics:

- [Relinking Oracle RAC libraries with the SF Oracle RAC libraries](#)
- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)
- [Setting or changing the product license level](#)
- [Upgrading disk layout versions](#)
- [Upgrading CVM protocol version and VxVM disk group version](#)
- [Post upgrade tasks for migrating the SFDB repository database](#)
- [Verifying the cluster](#)

Relinking Oracle RAC libraries with the SF Oracle RAC libraries

You must relink the Oracle RAC libraries with the SF Oracle RAC libraries after upgrading SF Oracle RAC.

To relink Oracle RAC 11g using the installer

1 Run the `installsffrac` installer:

```
# cd /opt/VRTS/install
# ./installsffrac -configure galaxy nebula
```

2 Enter **5** to select the option **Post Oracle Installation Tasks**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 5
```

3 Select the option **Relink Oracle Database Binary**.

```
1) Configure CSSD agent
2) Relink Oracle Database Binary
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

4 Provide the following information:

```
Enter Oracle UNIX user name: [b] (oracle)
Enter Oracle UNIX group name: [b] (oinstall)
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
Enter absolute path of Oracle Database Home directory: [b]
```

The installer detects the Oracle version.

5 Enter **y** to proceed with relinking.

```
Do you want to continue? [y,n,q] (y)
```

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

See [“Performing a rolling upgrade using the installer”](#) on page 224.

To re-join the backup boot disk group

- ◆ Re-join the *backup_bootdg* disk group to the boot disk group.

```
# /etc/vx/bin/vxrootadm -Y join backup_bootdg
```

where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

See [“Performing a rolling upgrade using the installer”](#) on page 224.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vx dg` command to find the boot disk group where you are currently booted.

```
# vx dg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Setting or changing the product license level

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

After you upgrade, perform one of the following steps:

- Obtain a valid license key and run the `vxlicinst` command to add it to your system.
- Use the `vxkeyless` command to update the license keys to the keyless license model.

For more information and instructions, see the chapter *Licensing SF Oracle RAC*.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 7, 8, and 9. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

See the `vxupgrade(1M)` manual page.

Support for disk layout Version 4 has been removed. You must upgrade any existing file systems with disk layout Version 4 to disk layout Version 7 or later using the `vxfsconvert` command.

See the `vxfsconvert(1M)` manual page.

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/voll | grep -i version
```

For more information about disk layout versions, see the *Veritas Storage Foundation Administrator's Guide*.

Upgrading CVM protocol version and VxVM disk group version

The default Cluster Volume Manager protocol version is 110.

Run the following command to verify the CVM protocol version:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the protocol version is not 110, run the following command to upgrade the version:

```
# /opt/VRTS/bin/vxdctl upgrade
```

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 170.

Check the existing disk group version:

```
# vxdg list dg_name | grep -i version
```

If the disk group version is not 170, run the following command on the master node to upgrade the version:

```
# vxdg -T 170 upgrade dg_name
```

Post upgrade tasks for migrating the SFDB repository database

Database Storage Checkpoints that have been created by using the SFDB tools before upgrade are visible using the `vxsfdm` CLI, and you can mount these Database Storage Checkpoints and roll back to them, if required. However, creating clones by using migrated Database Storage Checkpoints is not supported.

If you want to continue using previously created FlashSnap snapplans to take snapshots, you must validate them by using the `-o validate` option of the `vxsfdm` command.

To continue using the Database Storage Checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle RAC, you must perform one of the following procedures after upgrading SF Oracle RAC to 6.0:

- Rename startup script after upgrading from 5.0x and before migrating the SFDB repository
 See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 260.
- Migrate from a 5.0x SFDB repository database to 6.0
 See [“Migrating from a 5.0 repository database to 6.0”](#) on page 258.

Migrating from a 5.0 repository database to 6.0

Perform the following on one node only.

To migrate from a 5.0 repository database to 6.0

- 1 Rename the startup script NO_S*vxdbms3 to S*vxdbms3.
 See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 260.
- 2 As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to do this on one node.


```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
```
- 3 On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.
- 4 By default, the repository is created on the file system which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:
 - Repository path has to be a directory writable by Oracle user.
 - The repository must be accessible by all nodes. You can put it in a CFS file system, or put it in a resource group under VCS control so it can be failed over together with the Oracle database.
 - The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

Create an alternate repository path.
- 5 If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform the previous steps on the secondary host.

- 6 On the primary host, edit your snapplans to remove the "SNAPSHOT_DG=SNAP_*" parameter and add "SNAPSHOT_DG_PREFIX=SNAP_*". The parameter can be any PREFIX value and not necessarily "SNAP_*".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

- 7 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/vxsfadm -s flashsnap \  
-a oracle -c SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

To begin using the Storage Foundation for Databases (SFDB) tools:

See *Storage Foundation: Storage and Availability Management for Oracle Databases*

After upgrading from 5.0.x and before migrating SFDB

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 6.0 the S*vxdbs3 startup script is renamed to NO_S*vxdbs3. The S*vxdbs3 startup script is required by sfua_rept_migrate. Thus when sfua_rept_migrate is run, it is unable to find the S*vxdbs3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbs3 not found  
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.  
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

To prevent S*vxdbs3 startup script error

- ◆ Rename the startup script NO_S*vxdbs3 to S*vxdbs3.

Verifying the cluster

Run the SF Oracle RAC script-based installer or the Web-based installer to verify the functional health of the cluster.

The installer performs the following checks on each node:

- Verifies that SF Oracle RAC is installed.
- Verifies that the Oracle process daemon (oproc) is not running.
- Verifies that the Veritas libraries (ODM and VCSMM) are linked with the Oracle RAC libraries.
- Verifies whether the tunables for LLT/VCSMM tunables are set correctly.
- Verifies that the CSSD resource is configured under the CVM service group.
- Verifies whether the specified Oracle user has read permissions on the /etc/lltab and /etc/llhosts files.
- Verifies that both LLT and Oracle Clusterware/Grid Infrastructure have identical node information.

- Verifies that the Oracle RAC database version is the same on all nodes in the cluster. The check fails if the version or patchset information varies across nodes.

For information on resolving issues that you may encounter during the checks, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter: Troubleshooting SF Oracle RAC, Section: Troubleshooting installation and configuration check failures.

The installer skips the following checks when LLT is configured over UDP:

- LLT links' full duplex setting
- LLT link jumbo frame setting (MTU)
- LLT links' cross connection
- LLT links' speed and auto negotiation settings

To verify the cluster using the script-based installer

- 1 Start the SF Oracle RAC installer:

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

- 2 Select the option **SF Oracle RAC Installation and Configuration Checks**.

```
1)  Configure SF Oracle RAC sub-components  
2)  SF Oracle RAC Installation and Configuration Checks  
3)  Prepare to Install Oracle  
4)  Install Oracle Clusterware/Grid Infrastructure and Database  
5)  Post Oracle Installation Tasks  
6)  Exit SF Oracle RAC Configuration  
Choose option: [1-6,q] (1) 2
```

- 3 Enter **y** to confirm that Oracle RAC is installed on the nodes.

Note: The installer tries to discover the location of the Oracle Clusterware/Grid Infrastructure home directory from the Oracle inventory. If the installer discovers the information, it displays the information.

- 4 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

Note: The installer tries to discover the location of the Oracle Clusterware/Grid Infrastructure home directory from the Oracle inventory. If the installer discovers the information, it displays the information.

The installer verifies the path information and re-prompts if the path is incorrect.

- 5 Enter the full path of the Oracle RAC database home directory.

Note: The installer tries to discover the location of the Oracle RAC database home directory from the Oracle inventory. If the installer discovers the information, it displays the information.

The installer verifies the path information.

- 6 Enter the Oracle user name.

The installer starts the installation and configuration checks.

To verify the cluster using the Web-based installer

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 86.

- 2 Select **Configure a Product** from the Task drop-down list.

- 3 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list. Click **Next**.

- 4 Enter the system names in the System Names text box and click **Next**.

The installer performs system verification checks. Click **Next** after the verification completes.

- 5 Select **SF Oracle RAC Installation and Configuration Checks** from the Select a Task drop-down list. Click **Next**.

- 6 Click **Yes** to confirm that Oracle RAC is installed on the nodes and to enter the Oracle software installation information.
- 7 Provide the following Oracle RAC installation information:

Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory	Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory. Note: The installer tries to discover the location of the Oracle Clusterware/Grid Infrastructure home directory from the Oracle inventory. If the installer discovers the information, it displays the information. The installer verifies the path information and re-prompts if the path is incorrect.
Enter absolute path of Oracle Database Home directory	Enter the full path of the Oracle RAC database home directory. Note: The installer tries to discover the location of the Oracle RAC database home directory from the Oracle inventory. If the installer discovers the information, it displays the information. The installer verifies the path information.
Enter the Oracle user name	Enter the Oracle user name.

Click **Next**.

The installer starts the installation and configuration checks.

Installation and upgrade of Oracle RAC

- [Chapter 18. Before installing Oracle RAC](#)
- [Chapter 19. Installing Oracle RAC](#)
- [Chapter 20. Performing Oracle RAC post-installation tasks](#)
- [Chapter 21. Upgrading Oracle RAC](#)

Before installing Oracle RAC

This chapter includes the following topics:

- [Important preinstallation information for Oracle RAC](#)
- [About preparing to install Oracle RAC](#)
- [Preparing to install Oracle RAC using the SF Oracle RAC installer or manually](#)

Important preinstallation information for Oracle RAC

Before you install Oracle RAC, ensure that you meet the following requirements:

- Review the Oracle documentation for additional requirements pertaining to your version of Oracle.
- Keep the Oracle worksheets handy with the values appropriate for your installation setup.

See [“Oracle RAC worksheet”](#) on page 613.

Note: The manual procedures use variables, which are indicated by *italicized* text. Ensure that you replace these variables with actual values.

About preparing to install Oracle RAC

Use one of the following ways to perform the pre-installation tasks:

SF Oracle RAC installer	The SF Oracle RAC installer provides a menu-driven command line interface to step you through the pre-installation tasks. Note: Some of the pre-installation steps are not supported by the SF Oracle RAC installer and must be done manually as described in the manual procedures.
Manual	You need to perform the pre-installation tasks manually as described in the manual procedures.
Response file	You can pre-configure the systems for Oracle RAC installation using an SF Oracle RAC response file. The SF Oracle RAC response file in tandem with the Oracle RAC response files simplify the process of automating and standardizing Oracle RAC installations. Note: You can use the response file to automate only those pre-configuration tasks that are supported by the SF Oracle RAC installer. For instructions, see the chapter "Installing Oracle RAC using a response file" in this document.

The examples in this chapter assume a two-node cluster comprising the nodes galaxy and nebula.

Before installing Oracle RAC, review the Oracle installation manuals and the appropriate Oracle support Web sites. Some of the pre-installation tasks, wherever indicated in the document, must be done in accordance with the instructions in the Oracle installation manuals. The instructions for these tasks are not provided in this document.

Note: The instructions in this chapter use variables and sample values wherever required. Replace these variables and sample values with values that conform to your installation requirements.

Before you start the preparatory tasks, you may want to update the sample worksheet with the correct installation values and keep them handy during the process.

Preparing to install Oracle RAC using the SF Oracle RAC installer or manually

This section provides instructions for both manual as well as SF Oracle RAC installer-based procedures.

To prepare to install Oracle RAC

- 1 Identify the public virtual IP addresses for use by Oracle.
See [“Identifying the public virtual IP addresses for use by Oracle”](#) on page 268.
- 2 Set the kernel parameters.
See [“Setting the kernel parameters”](#) on page 268.
- 3 Verify that RPMs and patches required by Oracle are installed.
See [“Verifying that RPMs and patches required by Oracle are installed”](#) on page 269.
- 4 Verify that the user `nobody` exists.
See [“Verifying the user `nobody` exists”](#) on page 269.
- 5 Launch the SF Oracle RAC installer.
See [“Launching the SF Oracle RAC installer”](#) on page 269.
- 6 Create Oracle user and groups.
See [“Creating users and groups for Oracle RAC”](#) on page 271.
- 7 Create the storage for OCR and voting disk.
See [“Creating storage for OCR and voting disk ”](#) on page 275.
- 8 Configure the private network for Oracle RAC.
See [“Configuring private IP addresses for Oracle RAC”](#) on page 292.
- 9 For Oracle RAC 11g Release 2: Verify that multicast is functional on all private network interfaces.
See [“Verifying that multicast is functional on all private network interfaces”](#) on page 311.
- 10 Create the Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually.
See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 312.
- 11 Set up Oracle user equivalence on all nodes.
See [“Setting up user equivalence”](#) on page 320.
- 12 Verify the cluster configuration before the installation to ensure that there are no issues that may prevent a successful installation.
See [“Verifying the systems for Oracle RAC installation”](#) on page 320.

Identifying the public virtual IP addresses for use by Oracle

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

The IP address and the corresponding host name should be registered in the domain name service (DNS). Alternatively, an entry for the virtual IP address and virtual public name can be placed in the `/etc/hosts` file as shown in the following example:

```
10.182.79.239 galaxy-vip
10.182.79.240 nebula-vip
```

The `/etc/hosts` file on each node of the cluster should have these entries.

Oracle recommends that the public node name for the virtual IP address be in the following format *hostname-vip*. For example, *galaxy-vip*.

Note: The public node name (in other words, the alias for the virtual IP address) for the nodes must be different from the host's current fully qualified domain name (FQDN).

For Oracle RAC 11g Release 2: Additionally, you need a Single Client Access Name (SCAN) registered in Enterprise DNS that resolves to three IP addresses (recommended) using a round robin algorithm or at least one IP address. The IP addresses must be on the same subnet as your public network in the cluster. SCAN provides a single name for clients to access an Oracle database running in a cluster.

Note: The virtual IP addresses that are used for SCAN IP resolution must be on the same subnet. Oracle RAC does not support their configuration on different subnets.

Setting the kernel parameters

Set the kernel parameter values to meet Oracle RAC deployment requirements. The Oracle Universal Installer (OUI) verifies the settings at the time of installation to ensure that they satisfy the minimum requirements. The Oracle RAC installation fails if the kernel parameters are not configured properly. The settings can also be tuned to optimize system performance.

For instructions and guidelines, see the Oracle Metalink document: 169706.1

Verifying that RPMs and patches required by Oracle are installed

Oracle may require certain RPMs and patches to be installed before the installation of Oracle RAC.

For the list of RPMs and patches and related instructions, see the Oracle documentation.

Verifying the user `nobody` exists

To verify the user "nobody" exists on each system in the cluster:

```
# id nobody
uid=65001(nobody) gid=65001(nobody) groups=65001(nobody)
```

If the user does not exist, create the user:

```
# groupadd -g 65001 nobody
# useradd -g 65001 -u 65001 nobody
```

Note: Make sure that the user ID and group ID are the same across the nodes in your cluster.

Launching the SF Oracle RAC installer

You can use one of the following ways to launch the SF Oracle RAC installer:

- Veritas script-based installation program
See [“To launch the SF Oracle RAC script-based installer”](#) on page 270.
- Veritas Web-based installation program
See [“To launch the SF Oracle RAC Web-based installer”](#) on page 270.

To launch the SF Oracle RAC script-based installer

- 1 Log in as the root user on any one node and start the SF Oracle RAC script-based installer:

```
# cd /opt/VRTS/install

# ./installsfrc -configure galaxy nebula
```

The program displays the Symantec copyright information as well as the location of the installation logs.

- 2 Review the installer instructions and press **Enter** to proceed. From the configuration program menu, select **Prepare to Install Oracle**.
- 3 Select the pre-installation task you want to perform.
 - Enter **1** to select the option **Create Oracle User and Group**:
See [“Creating users and groups for Oracle RAC”](#) on page 271.
 - Enter **2** to select the option **Create Storage for OCR and voting disk**:
See [“Creating storage for OCR and voting disk ”](#) on page 275.
 - Enter **3** to select the option **Oracle Network Configuration**:
See [“Configuring private IP addresses for Oracle RAC”](#) on page 292.

To launch the SF Oracle RAC Web-based installer

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 2 Select **Configure a Product** from the Task drop-down list.
- 3 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list.
- 4 Enter the system names in the System Names text box and click **Validate**.
The installer performs system verification checks. Click **Next** after the verification completes.
- 5 Select **Prepare to install Oracle** from the Select a Task drop-down list. Click **Next**.
- 6 Select the pre-installation task you want to perform.
 - Select the option **Create Oracle User and Group** to create Oracle user and group.
See [“Creating users and groups for Oracle RAC”](#) on page 271.
 - Select the option **Create Storage for OCR and voting disk** to create the storage for OCR and voting disk.

See [“Creating storage for OCR and voting disk ”](#) on page 275.

- Select the option **Oracle Network Configuration** to configure the private IP addresses for Oracle RAC.

See [“Configuring private IP addresses for Oracle RAC”](#) on page 292.

Creating users and groups for Oracle RAC

Depending on the Oracle RAC version, create the required users and groups.

Create the following groups and users for Oracle RAC 11g Release 2:

- Oracle Inventory group
- dba group
- Oracle grid user
- Oracle user

Create additional users and groups as required by Oracle. Before creating Oracle users and groups, see the Oracle documentation.

You must assign Oracle Inventory as the primary group and dba as the secondary group. Oracle requires secondary groups for identifying operating system accounts that have database administrative (SYSDBA) privileges and for those accounts that have limited sets of database administrative (SYSOPER) privileges. Create the groups on all systems and assign the Oracle user to these groups.

Use one of the following ways to create the Oracle user and groups:

Using the SF Oracle RAC script-based installer	See “Creating the Oracle user and groups using the SF Oracle RAC script-based installer” on page 271.
Using the SF Oracle RAC Web-based installer	See “Creating the Oracle user and groups using the SF Oracle RAC Web-based installer” on page 273.
Manual	See “Creating the Oracle user and groups manually” on page 275.

Creating the Oracle user and groups using the SF Oracle RAC script-based installer

This procedure provides instructions for creating the Oracle user and groups using the SF Oracle RAC installer.

To create the Oracle user and groups on all nodes in the cluster

- 1 From the SF Oracle RAC installer menu, enter **1** to select the option **Create Oracle User and Group**.
- 2 Provide the following information for creating the Oracle user and groups: user name, group name, user ID, group ID, and the full path of the Oracle user home directory.

The user ID and group ID must not be in use on any node in the cluster. The installer suggests unused values, which you may use or change as required. The configuration program assigns the same values on all the nodes.

Note: If you are configuring GCO, then the user IDs and group IDs of all nodes on both the primary and secondary clusters must be the same. While configuring the user ID and group ID values on the secondary site, make sure that they are identical to the values used at the primary site.

```
Enter Oracle UNIX user name: [b] oracle
Enter Oracle UNIX group name: [b] oinstall
Enter Oracle user's ID (numerical): [b] (1165)
Enter Oracle group's ID (numerical): [b] (1165)
Enter absolute path of Oracle user's Home directory: [b] /home/oracle
Is this information correct? [y,n,q] (y)
```

Review and confirm the information provided.

- 3 For Oracle RAC 11g Release 2: Repeat the above steps to create the grid user.

```
Enter Oracle UNIX user name: [b] oracle
Enter Oracle UNIX group name: [b] (oinstall)
Enter Oracle user's ID (numerical): [b] (1168)
Enter Oracle group's ID (numerical): [b] (1000)
Enter absolute path of Oracle user's Home directory: [b] /home/grid
```


4 Enter the information for the secondary group.

```
Do you want to create a secondary group
for Oracle user? [y,n,q,b,?] (n) y
Enter Oracle UNIX group name: [b] (dba, oper etc.)
Enter Oracle group's ID (numerical): [b] (1996)

      Creating secondary group dba for
      Oracle user oracle on galaxy ..... Done
      Creating secondary group dba for
      Oracle user oracle on nebula ..... Done
Do you want to create another secondary group
for Oracle user? [y,n,q,b,?] (n)
```

5 Create a password for the oracle user on each node.

```
# passwd oracle
```

6 Set up passwordless SSH to install Oracle RAC binaries.

For instructions, see the appendix *Setting up inter-system communication* in this document.

Creating the Oracle user and groups using the SF Oracle RAC Web-based installer

This procedure provides instructions for creating the Oracle user and groups using the SF Oracle RAC Web-based installer.

To create the Oracle user and groups on all nodes in the cluster

- 1 From the SF Oracle RAC installer menu, select the option **Create Oracle User and Group**.
- 2 Provide the following information for creating the Oracle user and groups: user name, group name, user ID, group ID, and the full path of the Oracle user home directory.

The user ID and group ID must not be in use on any node in the cluster. The installer suggests unused values, which you may use or change as required. The configuration program assigns the same values on all the nodes.

Note: If you are configuring GCO, then the user IDs and group IDs of all nodes on both the primary and secondary clusters must be the same. While configuring the user ID and group ID values on the secondary site, make sure that they are identical to the values used at the primary site.

```
Enter Oracle UNIX user name:
Enter Oracle UNIX group name:
Enter Oracle user's ID (numerical) (1165):
Enter Oracle group's ID (numerical) (1165):
Enter absolute path of Oracle user's Home directory:
```

- 3 Click **Yes** at the confirmation prompts.
- 4 Click **Yes** to enter the information for the secondary group.

```
Enter Oracle UNIX secondary group name:
Enter Oracle group's ID (numerical) (1996):
```

- 5 Click **Ok** to confirm.
- 6 Click **Yes** if you want to create another secondary group, else click **No**.
- 7 For Oracle RAC 11g Release 2: Repeat the above steps to create the grid user.

```
Enter Oracle UNIX user name:
Enter Oracle UNIX group name (oinstall):
Enter Oracle user's ID (numerical) (1168):
Enter Oracle group's ID (numerical) (1000):
Enter absolute path of Oracle user's Home directory:
```

Creating the Oracle user and groups manually

Depending on the Oracle RAC version, create the necessary Oracle groups and users. Be sure to assign the same group ID, user ID, and home directory for the user on each system.

Note: When you create the user and group, make sure that you specify a user and group ID that is not in use.

To create the operating system Oracle user and group on each system

- 1 Create the primary and secondary group on each system.

Primary group:

```
# groupadd -g grp_id grp_name
```

Secondary group:

```
# groupadd -g grp_id_sec grp_name_sec
```

- 2 Create the Oracle user and the user home directory on each system:

```
# useradd -g grp_name -u user_id \
-G grp_name_sec -m -d user_home user_name
```

Creating storage for OCR and voting disk

Create appropriate storage for Oracle Cluster Registry (OCR) and the voting disk depending on the version of Oracle RAC.

Oracle RAC 11g Release 2 The OCR and the voting disk can reside on ASM or in directories in a cluster file system.

Note: You can use CVM raw volumes to create ASM disk groups.

Use one of the following ways to create the storage:

SF Oracle RAC script-based installer	See “Creating storage for OCR and voting disk using the SF Oracle RAC script-based installer” on page 276.
SF Oracle RAC Web-based installer	See “Creating storage for OCR and voting disk using the SF Oracle RAC Web-based installer” on page 281.

Manual

See [“Creating storage for OCR and voting disk manually”](#) on page 286.

You need to create CVM volumes or a CFS mount point for database file storage later in the installation process:

See [“Creating the Oracle RAC database”](#) on page 356.

Creating storage for OCR and voting disk using the SF Oracle RAC script-based installer

The SF Oracle RAC installer enables you to create OCR and voting disk storage on CVM raw volumes or on a clustered file system. After creating the storage, the installer adds the storage configuration to VCS for high availability.

If you are creating the OCR and voting disk storage on CVM raw volumes, the installer performs the following tasks:

- Creates CVM volumes for OCR and voting disk (two-way mirrored or unmirrored)
- Creates the OCR and voting disk volumes and sets the ownership
- Starts the volumes
- Adds the CVMVolDg resource to the VCS configuration in the cvm group so that the volumes are brought online automatically when the node starts
- Brings the CVMVolDg resource online

If you are creating the OCR and voting disk storage on CFS, the installer performs the following tasks:

- Creates CVM volumes for OCR and voting disk (two-way mirrored or unmirrored)
- Creates the OCR and voting disk volumes and sets the ownership
- Starts the volumes
- Creates the mount point and mounts it on all the nodes
- Sets the ownership for the CFS mount point
- Adds the CFSSMount and CVMVolDg resources to the VCS configuration in the cvm group so that the resources are brought online automatically when the nodes start
- Brings the CFSSMount and CVMVolDg resources online

To create storage for OCR and voting disk using the SF Oracle RAC installer

- 1 From the SF Oracle RAC installer menu, enter **2** to select the option **Create Storage for OCR and Voting disk**.

The following menu displays:

```

1) Create Oracle User and Group
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 2

```

- 2 Enter **y** to create the storage.

```

Do you want the installer to assist you in creating disk groups,
volumes and file systems for Oracle? (Mirroring enabled
by default) [y,n,q] (n) y

```

If you want to create the storage manually, enter **n**. The installer displays instructions for creating the storage manually. You may skip the remaining steps.

- 3 Select an appropriate option for the disk group.

```

1) Create a disk group
2) Use an existing disk group
b) Back to previous menu
Choose option: [1-2,b,q]

```

If you want to mirror the CVM volumes, click **Yes** at the following prompt:

```

Do you want to enable mirroring? [y,n,q] (y) y

```

- If you choose to create a disk group, the installer displays the list of existing disks that do not belong to any disk group. Specify the disk (by entering the serial numbers displayed next to the disk name) that you want to use to create the disk group. If you chose to mirror the CVM volumes, select at least two disks.

Enter the name of the disk group.

```

Enter the disk group name: [b] (ocrvotedg)

```

- If you choose to use an existing disk group, the installer displays the list of existing disk groups. Select a disk group. If you chose to mirror the CVM

volumes, select an existing disk group that contains at least two disks for mirroring.

4 Review and confirm the configuration information displayed:

```
CVM Master node: galaxy
Selected disks (including mirroring):
  1. Disk_2
  2. Disk_3
Disk group name: ocrvotedg
Is this information correct? [y,n,q] (y)
```

The installer initializes the disk groups.

5 Choose the type of storage.

- To create the storage on CVM raw volumes:
See [“Creating the OCR and voting disk storage on CVM raw volumes”](#) on page 278.
- To create the storage on CFS:
See [“Creating the OCR and voting disk storage on CFS”](#) on page 279.

Creating the OCR and voting disk storage on CVM raw volumes

Perform the steps in the following procedure to create the storage for OCR and voting disk on CVM raw volumes.

Note: For Oracle RAC 11g Release 2, you may create CVM raw volumes to create ASM disk groups that may be used to store the OCR and voting disk information.

To create the OCR and voting disk storage on CVM raw volumes

1 Enter 1 to select the option **CVM Raw Volume**.

```
1 CVM Raw Volume
2 Clustered File System
b Back to previous menu
Select the storage scheme to be used: [1-2,b,q] 1
```

2 Enter the name and size of the volume on which you want to store OCR information.

```
Enter the volume name for OCR: [b] (ocrvol)
Enter the volume size for OCR (in MB): [b] (320)
```

- 3 Enter the name and size of the volume on which you want to store voting disk information.

```
Enter the volume name for Vote: [b] (votevol)
Enter the volume size for Vote (in MB): [b] (320)
```

- 4 Enter the Oracle UNIX user name.

```
Enter Oracle UNIX user name: [b] oracle
```

- 5 Enter the Oracle UNIX group name.

```
Enter Oracle UNIX group name: [b] (oinstall)
```

- 6 Press **Return** to continue.
- 7 Review and confirm the configuration information. The installer creates the volumes and brings the corresponding resources online.
Press **Return** to continue.
- 8 Verify that the corresponding resource is online on all nodes in the cluster.

Note: It takes a few minutes for the CVMVolDg resource to come online.

```
# hares -state ocrvotevol_resname
```

Creating the OCR and voting disk storage on CFS

Perform the steps in the following procedure to create the storage for OCR and voting disk on CFS.

To create the OCR and voting disk storage on CFS

- 1 Enter **2** to select the option **Clustered File System**.

```
1 CVM Raw Volume
2 Clustered File System
b Back to previous menu
Select the storage scheme to be used: [1-2,b,q] 2
```

- 2 Specify whether you want to create separate file systems for OCR and voting disk:

```
Do you want to create separate filesystems for ocr and vote?
[y,n,q] (y)
```

3 Enter the name and size of the volume on which you want to store OCR and voting disk information.

- If you have chosen to create a shared file system for OCR and voting disk, enter the following information:

```
Enter the volume name for OCR and Voting disk:
[b] (ocrvotevol)
Enter the volume size for OCR and Voting disk (in MB):
[b] (640)
```

- If you have chosen to create separate file systems for OCR and voting disk, enter the following information:

```
Enter the volume name for OCR: [b] (ocrvol)
Enter the volume size for OCR (in MB): [b] (320)
Enter the volume name for Vote: [b] (votevol)
Enter the volume size for Vote (in MB): [b] (320)
```

4 Enter the Oracle UNIX user name.

```
Enter Oracle UNIX user name: [b] oracle
```

5 Enter the Oracle UNIX group name.

```
Enter Oracle UNIX group name: [b] (oinstall)
```

6 Press **Return** to continue.

7 Review and confirm the configuration information. The installer creates and starts the volumes on all nodes in the cluster.

8 Enter the CFS mount point for OCR and voting disk information.

- If you have chosen to create a shared file system for OCR and voting disk, enter the following information:

```
Enter the mount point location for CFS
(common for all the nodes) [b] (/ocrvote)
```

- If you have chosen to create separate file systems for OCR and voting disk, enter the following information:

```
Enter the mount point location for OCR storage
(common for all the nodes): [b] (/ocr)
Enter the mount point location for Vote storage
(common for all the nodes): [b] (/vote)
```


The installer creates the CFS mount points and sets the ownership. Press **Return** to continue.

- 9 Verify that the corresponding CVMVolDg and CFSMount resources are online on all nodes in the cluster:

Note: It takes a few minutes for the CVMVolDg resource to come online.

```
# hares -state ocrvotemnt_resname

# hares -state ocrvotevol_resname
```

Creating storage for OCR and voting disk using the SF Oracle RAC Web-based installer

The SF Oracle RAC installer enables you to create OCR and voting disk storage on CVM raw volumes or on a clustered file system. After creating the storage, the installer adds the storage configuration to VCS for high availability.

If you are creating the OCR and voting disk storage on CVM raw volumes, the installer performs the following tasks:

- Creates CVM volumes for OCR and voting disk (two-way mirrored or unmirrored) and sets the ownership
- Starts the volumes
- Adds the CVMVolDg resource to the VCS configuration in the cvm group so that the volumes are brought online automatically when the nodes start
- Brings the CVMVolDg resource online

If you are creating the OCR and voting disk storage on CFS, the installer performs the following tasks:

- Creates CVM volumes for OCR and voting disk (two-way mirrored or unmirrored)
- Starts the volumes
- Creates the mount point and mounts it on all the nodes
- Sets the ownership for the CFS mount point
- Adds the CFSMount and CVMVolDg resources to the VCS configuration in the cvm group so that the resources are brought online automatically when the node starts
- Brings the CFSMount and CVMVolDg resources online

To create storage for OCR and voting disk

- 1 From the SF Oracle RAC installer menu, select the option **Create Storage for OCR and Voting disk**. Click **Next**.

Confirm that SF Oracle RAC is running.

- 2 Confirm whether you want to create the storage using the installer or manually:

Do you want the installer to assist you in creating disk groups, volumes and file systems for Oracle?

Click **Yes** to create the storage.

If you want to create the storage manually, click **No**. The installer displays instructions for creating the storage manually. Follow the displayed instructions to create the storage.

- 3 Select an appropriate option for the disk group:

- **Create a disk group**
- **Use an existing disk group**

If you want to mirror the CVM volumes, click **Yes** at the following prompt:

Do you want to enable mirroring?

- 4 Depending on the disk group option selected, do one of the following:

Create a disk group

Enter the name of the disk group in the **New Disk Group Name** text box.

Example: **ocrvotedg**.

If you choose to create a disk group, the installer displays the list of existing disks that do not belong to any disk group. Select the disk that you want to use to create the disk group. If you chose to mirror the CVM volumes, select at least two disks.

Use an existing disk group

If you choose to use an existing disk group, the installer displays the list of existing disk groups. Select a disk group. If you chose to mirror the CVM volumes, select an existing disk group that contains at least two disks for mirroring. Click **Next**.

- 5 Click **Yes** to confirm the configuration information displayed:

The installer initializes the disk groups.

- 6 Choose the type of storage.

- To create the storage on CVM raw volumes:
See [“Creating the OCR and voting disk storage on CVM raw volumes”](#) on page 283.
- To create the storage on CFS:
See [“Creating the OCR and voting disk storage on CFS”](#) on page 284.

Creating the OCR and voting disk storage on CVM raw volumes

Perform the steps in the following procedure to create the storage for OCR and voting disk on CVM raw volumes.

Note: For Oracle RAC 11g Release 2, you may create CVM raw volumes to create ASM disk groups that may be used to store the OCR and voting disk information.

To create the OCR and voting disk storage on CVM raw volumes

- 1 Select the option **CVM raw volumes**. Click **Next**.
- 2 Provide the following information:

Enter the volume name for OCR	Enter the name of the volume on which you want to store OCR information. Example: ocrvol
Enter the volume name for Vote	Enter the name of the volume on which you want to store voting disk information. Example: votevol
Enter the volume size for OCR (in MB)	Enter the size of the volume on which you want to store OCR information. A minimum of 320 MB volume size is required. Example: 320
Enter the volume size for Vote (in MB)	Enter the size of the volume on which you want to store voting disk information. A minimum of 320 MB volume size is required. Example: 320

Click **Next**.

3 Provide the Oracle user information:

Enter the Oracle UNIX Enter the name of the Oracle user.

user name

Example:

grid

Enter the Oracle UNIX Enter the name of the Oracle group.

group name

Example: **oinstall**

Click **Next**.

- 4 Click **Yes** to confirm the configuration information. The installer creates the volumes and brings the corresponding resources online.
- 5 Verify that the resource is online on all nodes in the cluster.

Note: It takes a few minutes for the CVMVolDg resource to come online.

```
# hares -state ocrvotevol_resname
```

Creating the OCR and voting disk storage on CFS

Perform the steps in the following procedure to create the storage for OCR and voting disk on CFS.

To create the OCR and voting disk storage on CFS

- 1 Select the option **Clustered File System**. Click **Next**.
- 2 Specify whether you want to create separate file systems for OCR and voting disk:

```
Do you want to create separate filesystems for ocr and vote?
[y,n,q] (y)
```

- 3 Enter the name and size of the volume on which you want to store OCR and voting disk information.
 - If you have chosen to create a shared file system for OCR and voting disk, enter the following information:

Enter the volume name Enter the name of the volume on which you want to store OCR and voting disk information.

Example: **ocrvotevol**

Enter the volume size for OCR and Voting disk (in MB) Enter the size of the volume on which you want to store OCR and voting disk information. A minimum of 640 MB volume size is required.

Example: **640**

- If you have chosen to create separate file systems for OCR and voting disk, enter the following information:

Enter the volume name for OCR Enter the name of the volume on which you want to store OCR information.

Example: **ocrvol**

Enter the volume name for Vote Enter the name of the volume on which you want to store voting disk information.

Example: **votevol**

Enter the volume size for OCR and Voting disk (in MB) Enter the size of the volume on which you want to store OCR information. A minimum of 320 MB volume size is required.

Example: **320**

Enter the volume size for Vote (in MB) Enter the size of the volume on which you want to store voting disk information. A minimum of 320 MB volume size is required.

Example: **320**

Click **Next**.

- 4 Provide the Oracle user information:

Enter the Oracle UNIX user name Enter the name of the Oracle user.

Example:

grid

Enter the Oracle UNIX group name Enter the name of the Oracle group.

Example: **oinstall**

Click **Next**.

- 5 Click **Yes** to confirm the configuration information. The installer creates and starts the volumes on all nodes in the cluster.
- 6 Enter the CFS mount point for OCR and voting disk information.

- If you have chosen to create a shared file system for OCR and voting disk, enter the following information:

Enter the mount point location for CFS Example: `/ocrvote`
(common for all the nodes):

- If you have chosen to create separate file systems for OCR and voting disk, enter the following information:

Enter the mount point location for OCR Example: `/ocr`
(common for all the nodes):

Enter the mount point location for Vote Example: `/vote`
(common for all the nodes):

Click **Next**.

The installer creates the CFS mount points and sets the ownership.

- 7 Verify that the corresponding CVMVolDg and CFSSMount resources are online on all nodes in the cluster:

Note: It takes a few minutes for the CVMVolDg resource to come online.

For example:

```
# hares -state ocrvotemnt_resname
# hares -state ocrvotevol_resname
```

Creating storage for OCR and voting disk manually

Use one of the following storage options to create the OCR and voting disk storage:

CVM raw volumes	See “To create OCR and voting disk volumes on raw volumes” on page 287.
CFS	See “To create the storage for OCR and voting disks on CFS” on page 287.

Note: Whether you create volumes or file system directories, you can add them to the VCS configuration to make them highly available.

To create OCR and voting disk volumes on raw volumes

1 Log in as the root user.

2 Create a shared disk group:

```
# vxdg -s init ocrvote_dgname disk_name2 disk_name3
```

3 Create mirrored volumes in the shared group for OCR and voting disk:

```
# vxassist -g ocrvote_dgname make ocr_volname 300M nmirrors=2
```

```
# vxassist -g ocrvote_dgname make vote_volname 300M nmirrors=2
```

4 Set the ownership for the volumes:

```
# vxedit -g ocrvote_dgname set group=grp_name \
user=user_name mode=660 ocr_volname
# vxedit -g ocrvote_dgname set group=grp_name \
user=user_name mode=660 vote_volname
```

5 Add the storage resources to the VCS configuration to make them highly available.

See [“Adding the storage resources to the VCS configuration”](#) on page 288.

To create the storage for OCR and voting disks on CFS

1 Create a shared VxVM disk group:

```
# vxdg -s init ocrvote_dgname disk_name2 disk_name3
```

2 From the CVM master, create a mirrored volume (for example `ocrvotevol`) for OCR and voting disk:

```
# vxassist -g ocrvote_dgname make ocrvote_volname 640M nmirrors=2
```

3 From the CVM master, create a file system with the volume (`ocrvotevol`).

```
# mkfs -t vxfs /dev/vx/rdisk/ocrvote_dgname/ocrvote_volname
```

4 On each system, create a directory (for example, `/ocrvote`) on which to mount the file system containing OCR and voting disk.

```
# mkdir /ocrvote
```

- 5 On each system, mount the file system containing OCR and voting disk:

```
# mount -t vxfs -o cluster /dev/vx/dsk/ocrvote_dgname/\
ocrvote_volname ocrvote_mnt
```

- 6 From any system, change permissions on the file system containing OCR and voting disk.

For example:

```
# chown -R grid:oinstall ocrvote_mnt
```

- 7 Add the storage resources to the VCS configuration to make them highly available.

See [“Adding the storage resources to the VCS configuration”](#) on page 288.

Adding the storage resources to the VCS configuration

The type of storage resource you add to the VCS configuration depends on whether you chose to create the OCR and voting disk storage on raw volumes or CFS. If you chose to create the storage on raw volumes, you need to add a CVMVolDg resource to the VCS configuration. If you chose to create the storage on CFS, you need to add the CVMVolDg and CFMount resources to the VCS configuration.

Depending on the type of storage, follow the steps in one of the following procedures to add the storage resources to the VCS configuration using the command line interface (CLI):

For storage resources on CFS See [“To add the storage resources created on CFS to the VCS configuration”](#) on page 289.

For storage resources on CVM See [“To add the storage resources created on raw volumes to the VCS configuration”](#) on page 291.

Note: Set the attribute "Critical" to "0" for all the resources in the cvm service group. This ensures that critical CVM and CFS resources are always online.

To add the storage resources created on CFS to the VCS configuration

- 1 Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

- 2 Configure the CVM volumes under VCS:

```
# hares -add ocrvotevol_resname CVMVolDg cvm_grpname

# hares -modify ocrvotevol_resname Critical 0

# hares -modify ocrvotevol_resname CVMDiskGroup ocrvote_dgname

# hares -modify ocrvotevol_resname CVMVolume -add ocrvote_volname

# hares -modify ocrvotevol_resname CVMActivation sw
```

- 3 Set up the file system under VCS:

```
# hares -add ocrvotemnt_resname CFSSMount cvm_grpname

# hares -modify ocrvotemnt_resname Critical 0

# hares -modify ocrvotemnt_resname MountPoint "ocrvote_mnt"

# hares -modify ocrvotemnt_resname BlockDevice \
"/dev/vx/dsk/ocrvote_dgname/ocrvote_volname"
```

- 4 Link the parent and child resources:

```
# hares -link ocrvotevol_resname cvm_clus

# hares -link ocrvotemnt_resname ocrvotevol_resname

# hares -link ocrvotemnt_resname vxfsckd
```

- 5 Enable the resources:

```
# hares -modify ocrvotevol_resname Enabled 1

# hares -modify ocrvotemnt_resname Enabled 1

# haconf -dump -makero
```

- 6 Verify the configuration of the CVMVolDg and CFSSMount resources in the main.cf file.

For example:

```
CFSMount ocrvote_mnt_ocrvotedg (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg_ocrvotedg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

ocrvote_mnt_ocrvotedg requires ocrvote_voldg_ocrvotedg
ocrvote_mnt_ocrvotedg requires vxfsckd
ocrvote_voldg_ocrvotedg requires cvm_clus
```

- 7 Bring the CFSMount and CVMVolDg resources online on all systems in the cluster:

```
# hares -online ocrvotevol_resname -sys node_name
# hares -online ocrvotemnt_resname -sys node_name
```

Verify that the resources are online on all systems in the cluster:

```
# hares -state ocrvotevol_resname
# hares -state ocrvotemnt_resname
```

To add the storage resources created on raw volumes to the VCS configuration

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the CVM volumes under VCS:

```
# hares -add ocrvotevol_resname CVMVolDg cvm_grpname
```

```
# hares -modify ocrvotevol_resname Critical 0
```

```
# hares -modify ocrvotevol_resname CVMDiskGroup ocrvote_dgname
```

```
# hares -modify ocrvotevol_resname CVMVolume -add ocr_volname
```

```
# hares -modify ocrvotevol_resname CVMVolume -add vote_volname
```

```
# hares -modify ocrvotevol_resname CVMActivation sw
```

- 3 Link the parent and child resources:

```
# hares -link ocrvotevol_resname cvm_clus
```

- 4 Enable the resources:

```
# hares -modify ocrvotevol_resname Enabled 1
```

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

5 Verify the configuration of the CVMVolDg resource in the main.cf file.

For example:

```
CVMVolDg ocrvote_voldg_ocrvotedg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvol, votevol }
    CVMActivation = sw
)

ocrvote_voldg_ocrvotedg requires cvm_clus
```

6 Bring the ocrvote_voldg_ocrvotedg resource online on all systems in the cluster:

```
# hares -online ocrvotevol_resname -sys node_name
```

Verify that the resource is online on all systems in the cluster:

```
# hares -state ocrvotevol_resname
```

Configuring private IP addresses for Oracle RAC

Private IP addresses are required by Oracle RAC to provide communication between the cluster nodes. Depending on your private network configuration, you may need one or more IP addresses. You can configure the private IP addresses for high availability using the PrivNIC or MultiPrivNIC agents.

Note: IPv6 addresses are not supported in this release.

[Table 18-1](#) lists the available options for configuring the private network for Oracle RAC. Use one of the following options to configure the private network.

Table 18-1 Options for configuring the private network for Oracle RAC

Option	Description
PrivNIC configuration	Perform this configuration if you plan to use a common IP address for Oracle Clusterware/Grid Infrastructure heartbeat and Oracle RAC database cache fusion and you plan to configure the IP address for high availability using the PrivNIC agent. For instructions: See “Configuring the private IP address and PrivNIC resource” on page 293.

Table 18-1 Options for configuring the private network for Oracle RAC
(continued)

Option	Description
MultiPrivNIC configuration	Perform this configuration if you plan to: ■ Use multiple IP addresses for Oracle database cache fusion for load balancing ■ And, configure the IP addresses for high availability using the MultiPrivNIC agent For instructions: See “Configuring the private IP address information and MultiPrivNIC resource” on page 300.

Configuring the private IP address and PrivNIC resource

You need to configure the following information:

- An IP address on each node
- The PrivNIC agent for failing over IP addresses in the event of link failures. The Oracle Clusterware/Grid Infrastructure interconnects need to be protected against NIC failures and link failures. The installer discovers the existing private NICs on which LLT is configured. For maximum failover options, all available LLT links are used for PrivNIC configuration.

Note: The PrivNIC agent is not supported with Oracle RAC 11.2.0.2. For more information, see <http://www.symantec.com/business/support/index?page=content&id=TECH145261>.

Use one of the following ways to configure the PrivNIC and private IP address information:

SF Oracle RAC script-based installer	See “Configuring the private IP address and PrivNIC using the SF Oracle RAC script-based installer” on page 294.
SF Oracle RAC Web-based installer	See “Configuring the private IP address and PrivNIC using the SF Oracle RAC Web-based installer” on page 296.
Manual	See “Configuring the private IP address and PrivNIC resource manually” on page 299.

Configuring the private IP address and PrivNIC using the SF Oracle RAC script-based installer

The SF Oracle RAC installer performs the following tasks:

- Backs up the /etc/hosts file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the PrivNIC resource in the CVM group.

Perform the steps in the following procedure to configure the PrivNIC and private IP address using the installer.

To configure the PrivNIC and private IP address information

- 1 From the SF Oracle RAC menu, enter **3** to select the option **Oracle Network Configuration**.

```

1) Create Oracle User and Group
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 3
  
```

- 2 Enter **1** to select the option **Configure private IP addresses (PrivNIC Configuration)**.

```

1) Configure private IP addresses (PrivNIC Configuration)
2) Configure private IP addresses (MultiPrivNIC Configuration)
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 1
  
```

The installer discovers available LLT links and PrivNIC resources.

If PrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

Note: The installer only removes the corresponding PrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the /etc/hosts file.

- 3 Enter the name for the PrivNIC resource.

```
Enter the PrivNIC resource name: [b] (ora_priv)
```

- 4 Enter **y** to modify the priority of the PrivNIC interfaces.

Note: The priority you set determines the interface that the PrivNIC agent chooses during failover.

```
Do you want to update the priority of the PrivNIC
interfaces? [y,n,q] (n) y
```

- 5 Set the interface priority in decreasing order. The PrivNIC agent will assign the highest priority to the first interface specified in the list.

```
Enter the interface name in the decreasing priority order,
separated by a space: [b] (eth2 eth3) eth3 eth2
```

- 6
- Enter **y** to add the IP addresses to the `/etc/hosts` file.
 - Enter **n** if you choose to add the IP addresses to the file manually. Go to step 8.
- 7 Perform this step only if you enabled the installer to add the IP address to the `/etc/hosts` file in the previous step.

Provide the private IP address and the private node name for the IP address that must be added to the file.

Note: All IP addresses must be in the same subnet, failing which Oracle Clusterware/Grid Infrastructure will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the `/etc/hosts` file on one of nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

```
Enter the private IP for galaxy: [b] 192.168.12.1
Enter Hostname alias for the above IP address: [b] galaxy-priv
Enter the private IP for nebula: [b] 192.168.12.2
Enter Hostname alias for the above IP address: [b] nebula-priv
```

Go to step 9.

- 8 Perform this step only if you have chosen to add the IP address to the /etc/hosts file manually.

Enter the private IP address information.

```
Enter the private IP for galaxy: [b] 192.168.12.1
Enter the private IP for nebula: [b] 192.168.12.2
```

- 9 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer now displays the configuration information.

- 10 Enter **y** to review and confirm the configuration information. The installer adds the PrivNIC resources to the VCS configuration and updates the /etc/hosts file (if you chose an installer-based update).
- 11 If you chose to add the IP address information to the /etc/hosts file manually, proceed to update the file as described in the following procedure.
See [“Adding private IP addresses to the /etc/hosts file manually”](#) on page 309.
- 12 Verify the PrivNIC configuration updates made by the program in the main.cf file.
See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 310.

Configuring the private IP address and PrivNIC using the SF Oracle RAC Web-based installer

The SF Oracle RAC installer performs the following tasks:

- Backs up the /etc/hosts file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the PrivNIC resource in the CVM group.

Perform the steps in the following procedure to configure the PrivNIC and private IP address using the installer.

To configure the PrivNIC and private IP address information

- 1 From the SF Oracle RAC menu, select the option **Oracle Network Configuration**. Click **Next**.
- 2 Select the option **Configure private IP addresses (PrivNIC Configuration)** from the Select a Task drop-down list. Click **Next**.

The installer discovers available LLT links and PrivNIC resources.

If PrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

Note: The installer only removes the corresponding PrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the `/etc/hosts` file.

- 3 Enter the name for the PrivNIC resource.

Enter the PrivNIC resource name (ora_priv):

Click **Next**.

- 4 Click **Yes** in the confirmation dialog box to modify the priority of the PrivNIC interfaces.

Note: The priority you set determines the interface that the PrivNIC agent chooses during failover.

Set the interface priority in decreasing order. The PrivNIC agent will assign the highest priority to the first interface specified in the list.

Enter the interface name in the decreasing priority order, separated by a space (eth2 eth3): **eth3 eth2**

Click **Next**.

- 5
 - Click **Yes** to add the IP addresses to the `/etc/hosts` file.
 - Click **No** if you choose to add the IP addresses to the file manually. Go to step 7.

- 6 Perform this step only if you enabled the installer to add the IP address to the /etc/hosts file in the previous step.

Provide the private IP address and the private node name for the IP address that must be added to the file.

Note: All IP addresses must be in the same subnet, failing which Oracle Clusterware/Grid Infrastructure will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the /etc/hosts file on one of nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

```
Enter the private IP for galaxy: 192.168.12.1
Enter Hostname alias for the above IP address: galaxy-priv
Enter the private IP for nebula: 192.168.12.2
Enter Hostname alias for the above IP address: nebula-priv
```

Click **Next**.

Go to step 8.

- 7 Perform this step only if you have chosen to add the IP address to the /etc/hosts file manually.

Enter the private IP address information.

```
Enter the private IP for galaxy: 192.168.12.1
Enter the private IP for nebula: 192.168.12.2
```

- 8 Enter the netmask information for the private network:

```
Enter the Netmask for private network (255.255.255.0):
```

Click **Next**.

The SF Oracle RAC installer now displays the configuration information.

- 9 Click **Yes** to review and confirm the configuration information. The installer adds the PrivNIC resources to the VCS configuration and updates the /etc/hosts file (if you chose an installer-based update).

- 10 If you chose to add the IP address information to the `/etc/hosts` file manually, proceed to update the file as described in the following procedure.

See [“Adding private IP addresses to the `/etc/hosts` file manually”](#) on page 309.

- 11 Verify the PrivNIC configuration updates made by the program in the `main.cf` file.

See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 310.

Configuring the private IP address and PrivNIC resource manually

Perform the steps in the following procedure to configure the private IP address and PrivNIC resource manually. Configure the PrivNIC resource in the VCS group where you have configured the OCR and voting disk resources.

The sample procedure creates the PrivNIC resource in the `cvm` group. The PrivNIC agent plumbs the IP address to the specified network interface.

To configure the private IP address and PrivNIC resource manually

- 1 Log in as the root user on one of the nodes in the cluster.
- 2 Change the cluster configuration to read-write mode:
- 3 Create the PrivNIC resource and add the resource to the same group in which you plan to configure the `cssd` resource:

```
# haconf -makerw
```

```
# hares -add privnic_resname PrivNIC cvm_grpname
```

4 Modify the PrivNIC resource:

```
# hares -modify privnic_resname Critical 0
# hares -local privnic_resname Device
# hares -local privnic_resname Address
# hares -modify privnic_resname \
Device -add nic1_node1 0 -sys node_name1
# hares -modify privnic_resname \
Device -add nic2_node1 1 -sys node_name1
# hares -modify privnic_resname \
Address privnic_ip_node1 -sys node_name1
# hares -modify privnic_resname \
Device -add nic1_node2 0 -sys node_name2
# hares -modify privnic_resname \
Device -add nic2_node2 1 -sys node_name2
# hares -modify privnic_resname \
Address privnic_ip_node2 -sys node_name2
# hares -modify privnic_resname \
NetMask netmask_ip
# hares -modify privnic_resname Enabled 1
```

5 Change the cluster configuration to read-only mode:

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

Configuring the private IP address information and MultiPrivNIC resource

You need to configure the following information:

- An IP address on each node for Oracle Clusterware/Grid Infrastructure heartbeats
- One or more IP addresses on each node for the Oracle database
- The MultiPrivNIC agent for failing over IP addresses in the event of link failures. The Oracle Clusterware/Grid Infrastructure interconnects need to be protected against NIC failures and link failures. The MultiPrivNIC agent protects the links against failures, if multiple links are available. The installer discovers the existing private NICs on which LLT is configured. For maximum failover options, all available LLT links are used for MultiPrivNIC configuration.

Note: The MultiPrivNIC agent is not supported with Oracle RAC 11.2.0.2. For more information, see <http://www.symantec.com/business/support/index?page=content&id=TECH145261>

Use one of the following ways to configure the MultiPrivNIC and private IP address information:

SF Oracle RAC script-based installer	See “Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC script-based installer” on page 301.
SF Oracle RAC Web-based installer	See “Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC Web-based installer” on page 304.
Manual	See “Configuring MultiPrivNIC and private IP addresses manually ” on page 307.

Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC script-based installer

Perform the steps in the following procedure to configure the MultiPrivNIC and private IP address using the installer.

The installer performs the following tasks:

- Backs up the /etc/hosts file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the MultiPrivNIC resource in the CVM group.

To configure the MultiPrivNIC and private IP address information

1 From the SF Oracle RAC menu, enter 3 to select the option **Oracle Network Configuration.**

```
1) Create Oracle User and Group
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 3
```

2 Enter 2 to select the option **Configure private IP addresses (MultiPrivNIC Configuration).**

```
1) Configure private IP addresses (PrivNIC Configuration)
2) Configure private IP addresses (MultiPrivNIC Configuration)
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

The installer discovers available LLT links and MultiPrivNIC resources. If MultiPrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

Note: The installer only removes the corresponding MultiPrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the `/etc/hosts` file.

3 Enter the name for the MultiPrivNIC resource.

```
Enter the MultiPrivNIC resource name: [b] (ora_mpriv) multi_priv
```

- 4**
- Enter **y** to add the IP addresses to the `/etc/hosts` file.
 - Enter **n** if you choose to add the IP addresses to the file manually. Go to step 6.

- 5 Perform this step only if you enabled the installer to add the IP address to the /etc/hosts file in the previous step.

Provide the private IP address and the private node name for the IP addresses that must be added to the file. When you do not want to enter information at the prompts, enter **x**.

Note: The IP addresses used for a particular NIC on all nodes of a cluster must be in the same subnet. This subnet must be different from the subnets for the IP addresses on other NICs. Otherwise, Oracle Clusterware/Grid Infrastructure and UDP IPC will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the /etc/hosts file on one of the nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

```
Enter IP addresses for galaxy for eth1 separated by space: [b,q,?]
192.168.12.1
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] galaxy-priv
Enter IP addresses for galaxy for eth2
separated by space: [b,q,?] 192.168.2.1
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] galaxy-priv1
Enter IP addresses for nebula for eth1
separated by space: [b,q,?] 192.168.12.2
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] nebula-priv
Enter IP addresses for nebula for eth2
separated by space: [b,q,?] 192.168.2.2
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] nebula-priv1
```

Go to step [7](#).

- 6 Perform this step only if you have chosen to add the IP address to the /etc/hosts file manually.

Enter the private IP address information.

```
Enter IP addresses for galaxy for eth1
separated by space: [b,q,?] 192.168.12.1
Enter IP addresses for galaxy for eth2
separated by space: [b,q,?] 192.168.2.1
Enter IP addresses for nebula for eth1
separated by space: [b,q,?] 192.168.12.2
Enter IP addresses for nebula for eth2
separated by space: [b,q,?] 192.168.2.2
```

- 7 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer displays the configured parameters.

- 8 Enter **y** to review and confirm the configuration information. The installer adds the MultiPrivNIC resources and updates the /etc/hosts file (if you chose installer-based update).
- 9 If you chose to add the IP address information to the /etc/hosts file manually, proceed to update the file as described in the following procedure.
See [“Adding private IP addresses to the /etc/hosts file manually”](#) on page 309.
- 10 Verify the MultiPrivNIC configuration updates made by the program in the main.cf file.
See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 310.

Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC Web-based installer

Perform the steps in the following procedure to configure the MultiPrivNIC and private IP address using the installer.

The installer performs the following tasks:

- Backs up the /etc/hosts file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the MultiPrivNIC resource in the CVM group.

To configure the MultiPrivNIC and private IP address information

- 1 From the SF Oracle RAC menu, select the option **Oracle Network Configuration**.
- 2 Select the option **Configure private IP addresses (MultiPrivNIC Configuration)** from the Select a Task drop-down list.

The installer discovers available LLT links and MultiPrivNIC resources. If MultiPrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

Note: The installer only removes the corresponding MultiPrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the `/etc/hosts` file.

- 3 Enter the name for the MultiPrivNIC resource.

Enter the MultiPrivNIC resource name (multi_priv): **multi_priv**

Click **Next**.

- 4
 - Click **Yes** to add the IP addresses to the `/etc/hosts` file.
 - Click **No** if you choose to add the IP addresses to the file manually. Go to step 6.

- 5 Perform this step only if you enabled the installer to add the IP address to the /etc/hosts file in the previous step.

Provide the private IP address and the private node name for the IP addresses that must be added to the file.

Note: The IP addresses used for a particular NIC on all nodes of a cluster must be in the same subnet. This subnet must be different from the subnets for the IP addresses on other NICs. Otherwise, Oracle Clusterware/Grid Infrastructure and UDP IPC will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the /etc/hosts file on one of the nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

Enter IP addresses for galaxy for eth1 separated by space:

192.168.12.1

Enter Hostname aliases for the above IP addresses

separated by space: **galaxy-priv**

Click **Next**.

Enter IP addresses for galaxy for eth2

separated by space: **192.168.2.1**

Enter Hostname aliases for the above IP addresses

separated by space: **galaxy-priv1**

Click **Next**.

Enter IP addresses for nebula for eth1

separated by space: **192.168.12.2**

Enter Hostname aliases for the above IP addresses

separated by space: **nebula-priv**

Click **Next**.

Enter IP addresses for nebula for eth2

separated by space: **192.168.2.2**

Enter Hostname aliases for the above IP addresses

separated by space: **nebula-priv1**

Go to step [7](#).

- 6 Perform this step only if you have chosen to add the IP address to the /etc/hosts file manually.

Enter the private IP address information.

Enter IP addresses for galaxy for eth1
separated by space: **192.168.12.1**

Click **Next**.

Enter IP addresses for galaxy for eth2
separated by space: **192.168.2.1**

Click **Next**.

Enter IP addresses for nebula for eth1
separated by space: **192.168.12.2**

Click **Next**.

Enter IP addresses for nebula for eth2
separated by space: **192.168.2.2**

Click **Next**.

- 7 Enter the netmask information for the private network:

Enter the Netmask for private network (255.255.255.0):

The SF Oracle RAC installer displays the configured parameters.

- 8 Click **Yes** to review and confirm the configuration information. The installer adds the MultiPrivNIC resources and updates the /etc/hosts file (if you chose installer-based update).
- 9 If you chose to add the IP address information to the /etc/hosts file manually, proceed to update the file as described in the following procedure.

See [“Adding private IP addresses to the /etc/hosts file manually”](#) on page 309.

- 10 Verify the MultiPrivNIC configuration updates made by the program in the main.cf file.

See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 310.

Configuring MultiPrivNIC and private IP addresses manually

Perform the steps in the following procedure to configure MultiPrivNIC and private IP addresses for Oracle Clusterware/Grid Infrastructure and UDP IPC manually. Configure the MultiPrivNIC resource in the VCS group where you have configured the OCR and voting disk resources.

Make sure that the number of links are the same for every node. For example, in the following configuration, if the eth1 and eth2 links fail, the MultiPrivNIC agent fails over the IP addresses to the eth3 link on galaxy; On nebula, the absence of a third link results in the agent failing over the link to eth1, which is already down. This results in loss of communication between the nodes and causes Oracle Clusterware/Grid Infrastructure to reboot the cluster.

```
MultiPrivNIC multi_priv (
    Critical = 0
    Device@galaxy = {eth1 = 0, eth2 = 1,
                    eth3 = 2}
    Device@nebula = {eth1 = 0, eth2 = 1}
    Address@galaxy = {"192.168.12.1" = 0,
                    "192.168.2.1" = 1}
    Address@nebula = {"192.168.12.2" = 0,
                    "192.168.2.2" = 1}
    NetMask = "255.255.255.0"
)
```

Note: Avoid configurations where the number of links differ between the nodes.

The sample procedure creates the MultiPrivNIC resource in the cvm group. The MultiPrivNIC agent plumbs the IP addresses to the appropriate network interfaces.

To configure the private IP address and MultiPrivNIC manually

- 1 Log in as the root user on one of the nodes in the cluster.
- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Create the MultiPrivNIC resource and add the resource to the same group in which you plan to configure the cssd resource:

```
# hares -add multipriv_resname MultiPrivNIC cvm_grpname
```

4 Modify the MultiPrivNIC resource:

```
# hares -modify multipriv_resname Critical 0
# hares -local multipriv_resname Device
# hares -local multipriv_resname Address
# hares -modify multipriv_resname \
Device -add nic1_node1 0 -sys node_name1
# hares -modify multipriv_resname \
Device -add nic2_node1 1 -sys node_name1
# hares -modify multipriv_resname \
Address -add multipriv_ip1_node1 0 -sys node_name1
# hares -modify multipriv_resname \
Address -add multipriv_ip2_node1 1 -sys node_name1
# hares -modify multipriv_resname \
Device -add nic1_node2 0 -sys node_name2
# hares -modify multipriv_resname \
Device -add nic2_node2 1 -sys node_name2
# hares -modify multipriv_resname \
Address -add multipriv_ip1_node2 0 -sys node_name2
# hares -modify multipriv_resname \
Address -add multipriv_ip2_node2 1 -sys node_name2
# hares -modify multipriv_resname \
NetMask netmask_ip
# hares -modify multipriv_resname Enabled 1
```

5 Change the cluster configuration to read-only mode:

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

Adding private IP addresses to the /etc/hosts file manually

Perform the steps in the following procedure only if you plan to install Oracle RAC 10g Release 2 and you chose to add the private IP address information manually to the `/etc/hosts` file at the time of configuring them using the SF Oracle RAC installer.

To add private IP addresses to the /etc/hosts file manually

- 1 Log in to each system as the root user.
- 2 For a configuration using the PrivNIC agent, add the following entries to the /etc/hosts file.

For example:

```
192.168.12.1 galaxy-priv
```

```
192.168.12.2 nebula-priv
```

For a configuration using the MultiPrivNIC agent, add the following entries to the /etc/hosts file:

For example:

```
192.168.12.1 galaxy-priv
```

```
192.168.2.1 galaxy-priv1
```

```
192.168.12.2 nebula-priv
```

```
192.168.2.2 nebula-priv1
```

Verifying the VCS configuration for PrivNIC and MultiPrivNIC

After you complete the steps for configuring PrivNIC/MultiPrivNIC and the private IP addresses, verify the configuration in the VCS main.cf configuration file.

To verify the VCS configuration for PrivNIC and MultiPrivNIC

- 1 View the main.cf file located in the directory /etc/VRTSvcs/conf/config:

```
# more /etc/VRTSvcs/conf/config/main.cf
```

- 2 For a configuration using the PrivNIC agent:
 - Verify that the PrivNIC resource displays in the file.
- For example:

```
PrivNIC ora_priv (
    Critical = 0
    Device @galaxy = {eth1= 0, eth2= 1}
    Device @nebula = {eth1= 0, eth2= 1}
    Address @galaxy = "192.168.12.1"
    Address @nebula = "192.168.12.2"
```

```
NetMask = "255.255.255.0"
)
```

- Verify that the PrivNIC resource is online on all nodes in the cluster:

```
# hares -state priv_resname
Resource Attribute System Value
ora_priv State galaxy ONLINE
ora_priv State nebula ONLINE
```

3 For a configuration using the MultiPrivNIC agent:

- Verify that the MultiPrivNIC resource displays in the file.
For example:

```
MultiPrivNIC multi_priv (
    Critical = 0
    Device @galaxy = {eth1= 0, eth2 = 1}
    Device @nebula = {eth1= 0, eth2 = 1}
    Address @galaxy = {"192.168.12.1" =0, "192.168.2.1" =1}
    Address @nebula = {"192.168.12.2" =0, "192.168.2.2" =1}
    NetMask = "255.255.255.0"
)
```

- Verify that the MultiPrivNIC resource is online on all systems in the cluster:

```
# hares -state multipriv_resname
Resource      Attribute      System      Value
multi_priv    State          galaxy      ONLINE
multi_priv    State          nebula      ONLINE
```

4 Make sure that the specified private IP addresses and devices are displayed when you run the following command:

```
# ifconfig -a
```

5 From each system, verify that the private IP addresses are operational using the ping command.

Verifying that multicast is functional on all private network interfaces

Perform this step if you plan to install Oracle RAC 11g Release 2.

Multicast network communication on the private interconnect network must be enabled and functioning on all nodes otherwise the installation or upgrade of Oracle Grid Infrastructure may fail.

For more information, see the Oracle Metalink document: 1212703.1

Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually

You can create the Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the local file system or on a local Veritas file system, or on a Veritas cluster file system. When the installer prompts for the home directories at the time of installing Oracle Clusterware/Grid Infrastructure and Oracle database, it creates the directories locally on each node, if they do not exist.

Note: Symantec recommends that Oracle Clusterware and Oracle database binaries be installed local to each node in the cluster. For Oracle Grid Infrastructure binaries, Oracle requires that they be installed only on a local file system. Refer to the Oracle documentation for size requirements.

Table 18-2 lists the Oracle RAC directories you need to create:

Table 18-2 List of directories

Directory	Description
Oracle Grid Infrastructure Home Directory (GRID_HOME) (For Oracle RAC 11g Release 2)	The path to the home directory that stores the Oracle Grid Infrastructure binaries. The Oracle Universal Installer (OUI) installs Oracle Grid Infrastructure and Oracle ASM into this directory, also referred to as GRID_HOME. The directory must be owned by the installation owner of Oracle Grid Infrastructure (oracle or grid), with the permission set to 755. The path to the grid home directory must be the same on all nodes. Follow Oracle Optimal Flexible Architecture (OFA) guidelines while choosing the path.

Table 18-2 List of directories (continued)

Directory	Description
Oracle base directory (ORACLE_BASE)	The base directory that contains all the Oracle installations. For Oracle RAC 11g Release 2, create separate Oracle base directories for the grid user and the Oracle user. It is recommended that installations of multiple databases maintain an Optimal Flexible Architecture (OFA) configuration. The path to the Oracle base directory must be the same on all nodes. The permission on the Oracle base directory must be at least 755.
Oracle home directory (ORACLE_HOME)	The directory in which the Oracle database software is installed. The path to the Oracle home directory must be the same on all nodes. The permission on the Oracle home directory must be at least 755.

Use one of the following options to create the directories:

Local file system	See “To create the directories on the local file system” on page 313.
Cluster File System	See “To create the file system and directories on cluster file system for Oracle Clusterware and Oracle database” on page 315.

To create the directories on the local file system

- 1 Log in as the root user on each node.
- 2 Create a local file system and mount it using one of the following methods:
 - Using native operating system commands
For instructions, see the operating system documentation.
 - Using Veritas File System (VxFS) commands

As the root user, create a VxVM local diskgroup on each node.	<pre># vxpdg init vxvm_dg \ dg_name</pre>
Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure binaries and Oracle binaries.	<pre># vxassist -g vxvm_dg make clus_volname size # vxassist -g vxvm_dg make ora_volname size</pre>
Create the file systems with the volumes.	<pre># mkfs -t vxfs /dev/vx/rdisk/vxvm_dg/clus_volname # mkfs -t vxfs /dev/vx/rdisk/vxvm_dg/ora_volname</pre>

Mount the file system.

```
# mount -t vxfs /dev/vx/dsk/vxvm_dg/clus_volname \
clus_home
# mount -t vxfs /dev/vx/dsk/vxvm_dg/ora_volname \
oracle_home
```

3 Create the directories for Oracle RAC.

```
# mkdir -p grid_base
# mkdir -p clus_home
# mkdir -p oracle_base
# mkdir -p oracle_home
```

4 Set appropriate ownership and permissions for the directories.

```
# chown -R grid:oinstall grid_base
# chmod -R 775 grid_base
# chown -R grid:oinstall clus_home
# chmod -R 775 clus_home
# chown -R oracle:oinstall oracle_base
# chmod -R 775 oracle_base
# chown -R oracle:oinstall oracle_home
# chmod -R 775 oracle_home
```

5 Add the resources to the VCS configuration.

See [“To add the storage resources created on VxFS to the VCS configuration”](#) on page 314.

6 Repeat all the steps on each node of the cluster.

To add the storage resources created on VxFS to the VCS configuration

1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

2 Configure the VxVM volumes under VCS:

```
# hares -add dg_resname DiskGroup cvm
# hares -modify dg_resname DiskGroup vxvm_dg -sys node_name
# hares -modify dg_resname Enabled 1
```

3 Set up the file system under VCS:

```
# hares -add clusbin_mnt_resname Mount cvm

# hares -modify clusbin_mnt_resname MountPoint \
"clus_home"

# hares -modify clusbin_mnt_resname BlockDevice \
"/dev/vx/dsk/vxvm_dg/clus_volname" -sys node_name
# hares -modify clusbin_mnt_resname FSType vxfs
# hares -modify clusbin_mnt_resname FsckOpt "-n"
# hares -modify clusbin_mnt_resname Enabled 1
# hares -add orabin_mnt_resname Mount cvm

# hares -modify orabin_mnt_resname MountPoint \
"oracle_home"

# hares -modify orabin_mnt_resname BlockDevice \
"/dev/vx/dsk/vxvm_dg/ora_volname" -sys node_name
# hares -modify orabin_mnt_resname FSType vxfs
# hares -modify orabin_mnt_resname FsckOpt "-n"
# hares -modify orabin_mnt_resname Enabled 1
```

4 Link the parent and child resources:

```
# hares -link clusbin_mnt_resname vxvm_dg
# hares -link orabin_mnt_resname vxvm_dg
```

5 Repeat all the steps on each node of the cluster.

To create the file system and directories on cluster file system for Oracle Clusterware and Oracle database

Perform the following steps on the CVM master node in the cluster.

1 As the root user, create a VxVM shared disk group:

```
# vxdg -s init cvm_dg dg_name
```

2 Create separate volumes for Oracle Clusterware and Oracle database:

```
# vxassist -g cvm_dg make clus_volname size
# vxassist -g cvm_dg make ora_volname size
```

- 3 Create the Oracle base directory, clusterware home directory, and the Oracle home directory.

```
# mkdir -p oracle_base
# mkdir -p oracle_home
# mkdir -p clus_home
# mkdir -p grid_base
```

- 4 Create file systems with the volumes:

```
# mkfs -t vxfs /dev/vx/rdisk/cvm_dg/ora_volname
```

- 5 Mount the file systems. Perform this step on each node.

```
# mount -t vxfs -o cluster /dev/vx/dsk/cvm_dg/ora_volname \
oracle_home
```

- 6 Change the ownership and permissions on all nodes of the cluster.

```
# chown -R grid:oinstall grid_base
# chmod -R 775 grid_base
# chown -R grid:oinstall clus_home
# chmod -R 775 clus_home
# chown -R oracle:oinstall oracle_base
# chmod -R 775 oracle_base
# chown -R oracle:oinstall oracle_home
# chmod -R 775 oracle_home
```

- 7 Add the CVMVolDg and CFSSMount resources to the VCS configuration.

See [“To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI”](#) on page 317.

To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the CVM volumes under VCS:

```
# hares -add dg_resname CVMVolDg cvm

# hares -modify dg_resname Critical 0

# hares -modify dg_resname CVMDiskGroup cvm_dg

# hares -modify dg_resname CVMVolume -add clus_volname

# hares -modify dg_resname CVMVolume -add ora_volname

# hares -modify dg_resname CVMActivation sw
```

- 3 Set up the file system under VCS:

```
# hares -add clusbin_mnt_resname CFSMount cvm
# hares -modify clusbin_mnt_resname Critical 0
# hares -modify clusbin_mnt_resname MountPoint \
"clus_home"
# hares -modify clusbin_mnt_resname BlockDevice \
"/dev/vx/dsk/cvm_dg/clus_volname"
# hares -add orabin_mnt_resname CFSMount cvm
# hares -modify orabin_mnt_resname Critical 0
# hares -modify orabin_mnt_resname MountPoint \
"oracle_home"
# hares -modify orabin_mnt_resname BlockDevice \
"/dev/vx/dsk/cvm_dg/ora_volname"
```

- 4 Link the parent and child resources:

```
# hares -link dg_resname cvm_clus

# hares -link clusbin_mnt_resname dg_resname

# hares -link clusbin_mnt_resname vxfsckd

# hares -link orabin_mnt_resname dg_resname

# hares -link orabin_mnt_resname vxfsckd
```

5 Enable the resources:

```
# hares -modify dg_resname Enabled 1

# hares -modify clusbin_mnt_resname Enabled 1

# hares -modify orabin_mnt_resname Enabled 1

# haconf -dump -makero
```

6 Verify the resource configuration in the main.cf file.

The following is a sample resource configuration for Oracle RAC 10g Release 2:

```
CFSMount crsbin_mnt (
    Critical = 0
    MountPoint = "/u01/app/oracle/product/10.2.0/crsHOME"
    BlockDevice = "/dev/vx/dsk/bindg/crsbinvol"
)

CFSMount orabin_mnt (
    Critical = 0
    MountPoint = "/u01/app/oracle/product/10.2.0/dbHOME_1"
    BlockDevice = "/dev/vx/dsk/bindg/orabinvol"
)

CVMVolDg crsorabin_voldg (
    Critical = 0
    CVMGroup = bindg
    CVMVolume = { crsbinvol, orabinvol }
    CVMActivation = sw
)

crsbin_mnt requires crsorabin_voldg
crsbin_mnt requires vxfsckd
orabin_mnt requires crsorabin_voldg
orabin_mnt requires vxfsckd
crsorabin_voldg requires cvm_clus
```

7 Verify that the resources are online on all systems in the cluster.

```
# hares -state dg_resname

# hares -state clusbin_mnt_resname

# hares -state orabin_mnt_resname
```

Note: At this point, the crsorabin_voldg resource is reported offline, and the underlying volumes are online. Therefore, you need to manually bring the resource online on each node.

To bring the resource online manually:

```
# hares -online dg_resname -sys node_name
```

Setting up user equivalence

You must establish grid user (Oracle RAC 11g Release 2) and Oracle user equivalence on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

Set up passwordless SSH communication between the cluster nodes for the Oracle user and the grid user.

For more information, see the Oracle documentation.

Verifying the systems for Oracle RAC installation

Run the SF Oracle RAC installer as described in this section to verify that the cluster configuration settings are conducive for a successful Oracle RAC installation.

[Table 18-3](#) lists the checks performed by the installer on each node. For information on resolving issues that you may encounter during the process, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter: Troubleshooting SF Oracle RAC, Section: Troubleshooting installation and configuration check failures.

Table 18-3 List of checks

Check	Description
Time synchronization	The time settings are synchronized across nodes. If the time difference exceeds five seconds, the check fails.
System checks	All the nodes in the cluster have the same architecture. All the nodes in the cluster have the same number and type of processors.
Operating system and patch level checks	All the nodes in the cluster have the same operating system version and patch levels.
Kernel parameter	Verifies that the kernel parameter settings and tunables are in accordance with Oracle RAC recommendations.
User "nobody" check	Verifies that the user "nobody" exists on each node in the cluster. The check fails if the user does not exist.

Table 18-3 List of checks (*continued*)

Check	Description
LLT checks	The LLT checks are performed on only two links, the first and the last link specified in the <code>/etc/llttab</code> file. Verifies that the media speed setting for the links are set to full duplex. If the links are not set to full duplex, the check fails. Verifies that the media speed and auto-negotiation settings are the same across all nodes in the cluster. The check fails if the settings vary across the nodes and if auto-negotiation is enabled for the interface on any node in the cluster. Verifies that the payload (MTU) for the jumbo frames is between 1500 and 9000 bytes. Verifies that this setting is the same for all LLT links in the cluster. The check fails if the value is different for different LLT links in the cluster. Verifies that the links are not cross-connected. If the links are cross-connected, the check fails. Make sure that the links do not have a direct path of communication through switch interconnects. Verifies that the cluster has at least two high-priority links that do not belong to the public network. Verifies that the cluster configuration in the <code>/etc/llthosts</code> file is the same on all nodes in the cluster. Verifies that the cluster ID in the <code>/etc/llttab</code> file is the same on all the nodes. Note: The following checks are skipped if LLT is configured over UDP: LLT links' full duplex setting LLT link jumbo frame setting (MTU) LLT links' cross connection LLT links' speed and auto negotiation settings

Table 18-3 List of checks (*continued*)

Check	Description
I/O fencing checks	Verifies that I/O fencing is enabled on each node in the cluster. Note: The verification applies only for disk-based fencing. Verifies that all nodes in the cluster have registered keys with the coordinator disks. If registration keys are missing for any node in the cluster, the check fails for that node. The possible causes for check failure are: ■ Fencing is not configured on the cluster. ■ Fencing is configured in disabled mode. ■ Fencing is configured using CPS.
ODM checks	Verifies that port d is up and running on the nodes.The check fails if the port is not up. The possible causes for check failure are: ■ ODM is not running in the cluster. ■ ODM is running in standalone mode.
VCSMM checks	Verifies that port o is up and running on the nodes.The check fails if port o is not up.
GAB checks	Verifies that the ports a, b, d, f, h, o, u, v, w, y are up and running on the nodes. The check fails if any of the SF Oracle RAC components are not running.

To verify the systems for Oracle RAC installation

1 Start the SF Oracle RAC installer:

```
# cd /opt/VRTS/install

# ./installsfrc -configure galaxy nebula
```

2 Enter 2 to select the option **SF Oracle RAC Installation and Configuration Checks**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 2
```

The installer displays the location of the temporary log file and starts the verification process.

3 Enter n to return to the menu.

```
The installer was unable to detect an Oracle database
installation due to the absence of the ORACLE_HOME
environment variable. Is it installed? [y,n,q]n
```

Installing Oracle RAC

This chapter includes the following topics:

- [About installing Oracle RAC](#)
- [Installing the Oracle Clusterware/Grid Infrastructure software](#)
- [Configuring LLT links in the GPnP profile](#)
- [Installing the Oracle RAC database software](#)
- [Verifying the Oracle Clusterware/Grid Infrastructure and database installation](#)

About installing Oracle RAC

You can install Oracle RAC on shared storage or locally on each node.

Note: SF Oracle RAC supports the clusterware installation of Oracle versions 10g Release 2 and 11g Release 2. Oracle 11g Release 1 Clusterware is not supported. All database versions starting from Oracle 10g Release 2 and later are supported.

Use one of the following ways to install Oracle RAC:

SF Oracle RAC installer The SF Oracle RAC installer is available as a script-based or Web-based installer.

The SF Oracle RAC installer starts the installation process and prompts for information that is required by the Oracle Universal Installer. The Oracle Universal Installer launches with these installation values pre-filled and installs Oracle RAC.

You need to invoke the SF Oracle RAC script-based or Web-based installer to start the installation.

Oracle Universal Installer	The Oracle Universal Installer installs Oracle RAC. The installation values must be manually entered at the time of installation. You need to invoke the Oracle Universal Installer to install Oracle RAC.
Response files	The SF Oracle RAC script-based installer supports silent installation of Oracle RAC using its own response files. You need to modify the SF Oracle RAC response files to include the path information of the Oracle RAC software binaries and the Oracle RAC response files. Note: The SF Oracle RAC Web-based installer does not support silent installation of Oracle RAC. For more information and instructions: See “About response files” on page 373. For instructions, see the chapter “Installing Oracle RAC using a response file” in this document.

Note: The instructions in this chapter use variables and sample values wherever required. Replace these variables and sample values with values that conform to your installation requirements.

Before you start the installation:

- Keep the Oracle worksheets handy as you perform the installation tasks.
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 628.
See [“Required installation information for Oracle database”](#) on page 629.
- Review your Oracle installation manuals and the appropriate Oracle support Web sites for additional information required during the installation.

Installing the Oracle Clusterware/Grid Infrastructure software

This section provides instructions for installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer. The SF Oracle RAC installer prompts for information required to invoke the Oracle Universal Installer and launches it. The responses provided to the SF Oracle RAC installer are pre-filled in the Oracle Universal Installer wizard. When you step through the installation, review or change these installation values in the Oracle Universal Installer.

Note: Before you begin the installation, verify that the nodes in the cluster are connected with network links using similar network devices. For example, if you use eth0 as a public link on one node in the cluster, all other nodes in the cluster must also use eth0 as the public link. Similarly, if you use eth1 as a private link on one node in the cluster, all other nodes in the cluster must also use eth1 as the private link.

Oracle Grid Infrastructure software is installed on each node in the GRID_HOME location.

Note: If you want to install Oracle Clusterware/Grid Infrastructure on VxFS or CFS, make sure that you created the appropriate storage before proceeding with the installation.

See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 312.

Install the software using one of the following methods:

SF Oracle RAC script-based installer	Using script-based installer: See “Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC script-based installer” on page 326. Using response files: See the chapter <i>Installation of SF Oracle RAC and Oracle RAC using a response file</i> in this document.
SF Oracle RAC Web-based installer	See “Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC Web-based installer” on page 329.
Oracle Universal Installer	See “Installing Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer” on page 332.

Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC script-based installer

The SF Oracle RAC installer performs the following tasks:

- Invokes the Oracle Universal Installer to install Oracle Clusterware/Grid Infrastructure
- Verifies the Oracle Clusterware/Grid Infrastructure installation

To install Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC script-based installer

- 1 Make sure that you completed the required pre-installation steps.
- 2 Return to the following SF Oracle RAC installer menu and select the option **Install Oracle Clusterware/Grid Infrastructure and Database**.

```
1)  Configure SF Oracle RAC sub-components
2)  SF Oracle RAC Installation and Configuration Checks
3)  Prepare to Install Oracle
4)  Install Oracle Clusterware/Grid Infrastructure and Database
5)  Post Oracle Installation Tasks
6)  Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 4
```
- 3 Select the option **Install Oracle Clusterware/Grid Infrastructure**.

```
1)  Install Oracle Clusterware/Grid Infrastructure
2)  Install Oracle Database
3)  Exit SF Oracle RAC Configuration
b)  Back to previous menu
Choose option: [1-3,b,q] (1) 1
```
- 4 Verify that the nodes in the cluster are connected with network links using similar network devices. Review the related information on screen and press **Enter** to confirm.

```
Do you want to continue? [y,n,q] (y)
```
- 5 Set the DISPLAY environment variable that is required for the Oracle Universal Installer:

```
Enter DISPLAY environment variable: [b] 10.20.12.150:0.0
```

where 10.20.12.150 is the IP address of X server where you want to export the display for the installer.
- 6 Enter the Oracle UNIX user name. The Oracle UNIX user name was previously set up during the pre-installation process.

```
Enter Oracle UNIX user name: [b] (grid)
```

- 7 Enter Oracle UNIX group name. The Oracle UNIX group name was previously set up during the pre-installation process.

```
Enter Oracle UNIX group name: [b] (oinstall)
```

- 8 Enter the full path of the Oracle base directory.

Note: The ORACLE_BASE directory must be a local directory.

- 9 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.
If the Oracle Clusterware/Grid Infrastructure home directory you specified does not exist, the installer creates the directory locally on each node and sets appropriate permissions for the Oracle user.

- 10 Enter the full path of the Oracle Clusterware/Grid Infrastructure installation image. Press **Return** to proceed.

The installer detects the version of the Oracle Clusterware/Grid Infrastructure software.

- 11 Enter **y** to continue with the installation.
- 12 Review and confirm the configuration information. The installer invokes the Oracle Universal Installer:
- 13 Enter the required information when prompted by the Oracle Universal Installer.
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 628.
- 14 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware/Grid Infrastructure installation.

- 15 For Oracle RAC 11g Release 2: On SLES 11 nodes, verify that the RPMs `libcap1-1.10-6.10.x86_64.rpm` and `libcap2-2.11-2.15.x86_64.rpm` are installed:

```
# rpm -qa | grep libcap
libcap1-1.10-6.10
libcap2-2.11-2.15
```

If they are not installed, install them using the instructions in the Oracle Metalink document: 952051.1

- 16 At the end of the Oracle Clusterware/Grid Infrastructure installation, run the following configuration scripts as the root user from each node of the cluster, in the listed order.

- `oraInstRoot.sh` (located in the `oraInventory` directory)
Make sure the script exists on each node before proceeding.
 - `root.sh` (located in the `CRS_HOME` or `GRID_HOME` directory, depending on your Oracle RAC version)
- Do not run the scripts simultaneously on the cluster nodes.
- 17 Return to the Oracle Universal Installer window and click **OK** to continue.
The Oracle Universal Installer informs you that the Oracle Clusterware/Grid Infrastructure installation was successful.
 - 18 Return to the SF Oracle RAC installer and press Return to proceed. The installer verifies whether Oracle Clusterware/Grid Infrastructure is installed properly.
This completes the Oracle Clusterware/Grid Infrastructure installation.

Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC Web-based installer

The SF Oracle RAC installer performs the following tasks:

- Invokes the Oracle Universal Installer to install Oracle Clusterware/Grid Infrastructure
- Verifies the Oracle Clusterware/Grid Infrastructure installation

To install Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC Web-based installer

- 1 Make sure that you completed the required pre-installation steps.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 3 Select the following menu options from the Select a task and product page:

Task	Configure a Product
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 4 Indicate the systems on which to install the software. Enter one or more system names, separated by spaces. Click **Next**.
After the validation completes successfully, click **Next**.
- 5 Select **Install Oracle Clusterware/Grid Infrastructure and Database** from the Select a Task page. Click **Next**.

- Select **Install Oracle Clusterware/Grid Infrastructure** from the Select a Task page. Click **Next**. Review the information on network interfaces and confirm to proceed.
- Provide the following information:

Enter **DISPLAY** environment variable

Set the DISPLAY environment variable that is required for the Oracle Universal Installer.

For example, **10.20.12.150:0.0**

where 10.20.12.150 is the IP address of X server where you want to export the display for the installer.

Enter Oracle UNIX user name

For example, **grid**

Enter Oracle UNIX group name

For example, **oinstall**

Enter absolute path of Oracle Base directory

Enter the full path of the Oracle base directory.

Note: The ORACLE_BASE directory must be a local directory.

Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory

Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

If the Oracle Clusterware/Grid Infrastructure home directory you specified does not exist, the installer creates the specified directory locally on each node and sets appropriate permissions for the Oracle user.

Enter absolute path of Oracle Clusterware/Grid install image

Enter the full path of the Oracle Clusterware/Grid Infrastructure installation image.

The installer detects the version of the Oracle Clusterware/Grid Infrastructure software.

Enter the arguments to be passed to the Oracle installer

Enter one or more of the following arguments (multiple argument values must be separated by a space):


```
-ignorePrereq
-ignoreSysPrereqs
-ignoreInternalDriverError
```


Note: Failure to pass one or more of the arguments during Oracle RAC 10g Release 2 installation causes issues with the installation.

Confirm the detected Oracle version to proceed.

- 8 Review and confirm the configuration information. The installer invokes the Oracle Universal Installer.
- 9 Enter the required information when prompted by the Oracle Universal Installer.
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 628.
- 10 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware/Grid Infrastructure installation.

- 11 For Oracle RAC 11g Release 2: On SLES 11 nodes, verify that the RPMs `libcap1-1.10-6.10.x86_64.rpm` and `libcap2-2.11-2.15.x86_64.rpm` are installed:

```
# rpm -qa | grep libcap
libcap1-1.10-6.10
libcap2-2.11-2.15
```

If they are not installed, install them using the instructions in the Oracle Metalink document: 952051.1

- 12 At the end of the Oracle Clusterware/Grid Infrastructure installation, run the following configuration scripts as the root user from each node of the cluster, in the listed order.
 - `oraInstRoot.sh` (located in the `oraInventory` directory)
Make sure the script exists on each node before proceeding.
 - `root.sh` (located in the `CRS_HOME` or `GRID_HOME` directory, depending on your Oracle RAC version)

Do not run the scripts simultaneously on the cluster nodes.

- 13 Return to the Oracle Universal Installer window and click **OK** to continue.
The Oracle Universal Installer informs you that the Oracle Clusterware/Grid Infrastructure installation was successful.
- 14 Return to the SF Oracle RAC installer and click **OK** to proceed. The installer verifies whether Oracle Clusterware/Grid Infrastructure is installed properly.
This completes the Oracle Clusterware/Grid Infrastructure installation.

Installing Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer

This section provides instructions for installing the Oracle Clusterware/Grid Infrastructure software using the Oracle Universal Installer. The software is installed on each node in the Oracle Clusterware/Grid Infrastructure home directory.

To install Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer

- 1 Log in as the Oracle grid user (Oracle RAC 11g Release 2) or as the Oracle user (Oracle RAC 10g Release 2). On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

where 10.20.12.150 is the IP address of X server where you want to export the display for the installer.

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

where 10.20.12.150 is the IP address of X server where you want to export the display for the installer.

- 2 Start the Oracle Universal Installer on the first node.

```
$ ./runInstaller -ignoreSysPrereqs
```

- 3 Enter the required information when prompted by the Oracle Universal Installer.

See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 628.

- 4 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware/Grid Infrastructure installation.

Note: For Oracle RAC 11g Release 2: If you want to save the Oracle Grid Infrastructure installation configuration into a response file for future installations, click the **Save Response File** option on the Summary page of the Oracle Universal Installer.

- 5 Run the orainstRoot.sh script as prompted by the Oracle Universal Installer.

- 6 For Oracle RAC 11g Release 2: On SLES 11 nodes, verify that the RPMs `libcap1-1.10-6.10.x86_64.rpm` and `libcap2-2.11-2.15.x86_64.rpm` are installed:

```
# rpm -qa | grep libcap
libcap1-1.10-6.10
libcap2-2.11-2.15
```

If they are not installed, install them using the instructions in the Oracle Metalink document: 952051.1

- 7 Run the `root.sh` script on each node as prompted by the Oracle Universal Installer:

The Oracle Clusterware daemons are started on the node.

Configuring LLT links in the GPnP profile

Perform this step only for Oracle RAC 11g Release 2 installations.

Update the GPnP profile to include the remaining LLT links that were not added to the profile during the Oracle Grid Infrastructure installation.

To configure the LLT links in the GPnP profile

1 View the currently configured interfaces:

```
# $GRID_HOME/bin/oifcfg getif
eth0 10.2.156.0          global      public
eth1 192.168.12.0       global      cluster_interconnect
```

The interfaces that are currently stored in the GPnP profile, their subnets, and their role (public or cluster_interconnect) are displayed.

2 Add the remaining LLT links to the GPnP profile:

```
# $GRID_HOME/bin/oifcfg setif -global \
eth2/192.168.12.0:cluster_interconnect
```

If you are using multiple IP addresses on different subnet for cluster interconnect (for load balancing), add the remaining interface subnets to the GPnP profile.

```
# $GRID_HOME/bin/oifcfg setif -global \
eth2/192.168.2.0:cluster_interconnect
# $GRID_HOME/bin/oifcfg setif -global \
eth1/192.168.2.0:cluster_interconnect
```

3 Verify that the correct interface subnet is in use:

```
# $GRID_HOME/bin/oifcfg getif
eth0 10.2.156.0          global      public
eth1 192.168.12.0       global      cluster_interconnect
eth2 192.168.12.0       global      cluster_interconnect
eth1 192.168.2.0        global      cluster_interconnect
eth2 192.168.2.0        global      cluster_interconnect
```

Make sure all the LLT links are configured and listed in the GPnP profile.

Installing the Oracle RAC database software

Before you start the installation of Oracle database, make sure that Oracle Clusterware/Grid Infrastructure is up and running. Symantec recommends you to install the Oracle database locally on each node.

Note: If you want to install Oracle database on VxFS or CFS, make sure that you created the appropriate storage before proceeding with the installation.

See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 312.

Install the software using one of the following methods:

SF Oracle RAC script-based installer

Using script-based installer:

See [“Installing the Oracle RAC database using the SF Oracle RAC script-based installer”](#) on page 335.

Using response files:

See the chapter *Installation of SF Oracle RAC and Oracle RAC using a response file* in this document.

SF Oracle RAC Web-based installer

See [“Installing the Oracle RAC database using the SF Oracle RAC Web-based installer”](#) on page 337.

Oracle Universal Installer

See [“Installing the Oracle RAC database using the Oracle Universal Installer”](#) on page 340.

Installing the Oracle RAC database using the SF Oracle RAC script-based installer

The SF Oracle RAC installer performs the following tasks:

- Verifies the status of Oracle Clusterware/Grid Infrastructure on all nodes
- Invokes the Oracle Universal Installer to install the Oracle database
- Verifies the Oracle RAC database installation

To install the Oracle RAC database using the SF Oracle RAC script-based installer**1 Return to the SF Oracle RAC installer and type 4 to select the option **Install Oracle Clusterware/Grid Infrastructure and Database**.**

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 4
```

2 Select the option **Install Oracle Database.**

```
1) Install Oracle Clusterware/Grid Infrastructure
2) Install Oracle Database
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

3 Set the DISPLAY environment variable that is required for the Oracle Universal Installer.

```
Enter the DISPLAY environment variable: [b] 10.20.12.150:0.0
```

where **10.20.12.150** is the IP address of X server where you want to export the display for the installer.

4 Enter Oracle UNIX user name. The Oracle UNIX user name was previously set up during the pre-installation process.

```
Enter Oracle UNIX user name: [b] (oracle)
```

5 Enter Oracle UNIX group name. The Oracle UNIX group name was previously set up during the pre-installation process.

```
Enter Oracle UNIX group name: [b] (oinstall)
```

6 Enter the full path of the Oracle base directory.

```
Enter absolute path of Oracle Base directory: [b]
```


- 7 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory: [b]

- 8 Enter the full path of the Oracle database home directory.

Enter absolute path of Oracle Database Home directory: [b]

If the Oracle RAC database home directory you specified does not exist, the installer creates the directory locally on each node and sets appropriate permissions for the Oracle user.

- 9 Enter the full path of the database installation image.

Enter absolute path of Oracle Database install image: [b]

The installer determines the version of the Oracle software from the binaries.

- 10 Enter **y** to proceed with the installation.

- 11 Review and confirm the configuration information.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running and invokes the Oracle Universal Installer:

- 12 Enter the required information when prompted by the Oracle Universal Installer.

See [“Required installation information for Oracle database”](#) on page 629.

- 13 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle database installation.

- 14 Run the root.sh script as the root user on the cluster nodes:

Return to the Oracle Universal Installer window and click **OK** to continue.

- 15 Return to the SF Oracle RAC installer and press **Return** to continue. The installer verifies the Oracle database installation.

Note: After the installation completes successfully, the installer prompts you to relink the database binaries. Symantec recommends you to relink the SF Oracle RAC libraries only after completing all the required patch additions, if any. See the Oracle documentation for patch updates that may be required.

Installing the Oracle RAC database using the SF Oracle RAC Web-based installer

The SF Oracle RAC installer performs the following tasks:

- Verifies the status of Oracle Clusterware/Grid Infrastructure on all nodes
- Invokes the Oracle Universal Installer to install the Oracle database
- Verifies the Oracle RAC database installation

To install the Oracle RAC database using the SF Oracle RAC Web-based installer

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 2 Select the following menu options from the Select a task and product page:

Task	Configure a Product
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.
After the validation completes successfully, click **Next**.
- 4 Select **Install Oracle Clusterware/Grid Infrastructure and Database** from the Select a Task page. Click **Next**.
- 5 Select **Install Oracle Database** from the Select a Task page. Click **Next**.
Confirm to proceed.

6 Provide the following information:

Enter DISPLAY environment variable	Set the DISPLAY environment variable that is required for the Oracle Universal Installer. For example, 10.20.12.150:0.0 where 10.20.12.150 is the IP address of X server where you want to export the display for the installer.
Enter Oracle UNIX user name	For example, oracle
Enter Oracle UNIX group name	For example, oinstall
Enter absolute path of Oracle Base directory	Enter the full path of the Oracle base directory. Note: The ORACLE_BASE directory must be a local directory.
Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory	Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.
Enter absolute path of Oracle Database Home directory	Enter the full path of the Oracle database home directory. If the Oracle RAC database home directory you specified does not exist, the installer creates the directory locally on each node and sets appropriate permissions for the Oracle user.
Enter absolute path of Oracle Database install image	Enter the full path of the database installation image. The installer determines the version of the Oracle software from the binaries.

Confirm the detected Oracle version to proceed.

7 Review and confirm the configuration information.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running and invokes the Oracle Universal Installer:

- 8 Enter the required information when prompted by the Oracle Universal Installer.
See [“Required installation information for Oracle database”](#) on page 629.
- 9 Review the configuration summary presented by the Oracle Universal Installer.
The Oracle Universal Installer begins the Oracle database installation.

- 10 Run the `root.sh` script as the root user on the cluster nodes:
Return to the Oracle Universal Installer window and click **OK** to continue.
- 11 Return to the SF Oracle RAC installer and click **OK** to continue. The installer verifies the Oracle database installation.

Note: After the installation completes successfully, the installer prompts you to relink the database binaries. Symantec recommends you to relink the SF Oracle RAC libraries only after completing all the required patch additions, if any. See the Oracle documentation for patch updates that may be required.

Installing the Oracle RAC database using the Oracle Universal Installer

The following procedure describes how to install the Oracle RAC database using the Oracle Universal Installer. Symantec recommends that you install the Oracle RAC database locally on each node.

To install Oracle RAC database using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the `DISPLAY` variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

where **10.20.12.150** is the IP address of X server where you want to export the display for the installer.

- 2 Start the Oracle Universal Installer.

```
$ ./runInstaller -ignoreSysPrereqs
```

- 3 Enter the required information when prompted by the Oracle Universal Installer.
See [“Required installation information for Oracle database”](#) on page 629.

- 4 Review the configuration summary presented by the Oracle Universal Installer.

Note: For Oracle RAC 11g Release 2: If you want to save the Oracle RAC database installation configuration into a response file for future installations, click the **Save Response File** option on the Summary page of the Oracle Universal Installer.

The Oracle Universal Installer begins the Oracle database installation.

- 5 Run the root.sh script as prompted by the Oracle Universal Installer.

```
# cd $ORACLE_HOME
# ./root.sh
```

Verifying the Oracle Clusterware/Grid Infrastructure and database installation

The following procedure verifies the Oracle Clusterware/Grid Infrastructure and Oracle RAC database installation by verifying that the Oracle processes are running on all nodes.

To verify the installation, run the following command from any node in the cluster. Verify in the command output that the Oracle Clusterware/Grid Infrastructure processes are online on the nodes.

```
# $GRID_HOME/bin/crsctl stat res -t
```

NAME	TARGET	STATE	SERVER	STATE_DETAILS
Local Resources				
ora.LISTENER.lsnr				
	ONLINE	ONLINE	galaxy	
	ONLINE	ONLINE	nebula	
ora.asm				
	OFFLINE	OFFLINE	galaxy	
	OFFLINE	OFFLINE	nebula	
.				
.				

```
.
-----
Cluster Resources
-----
ora.LISTENER_SCAN1.lsnr
      1          ONLINE  ONLINE      nebula
ora.LISTENER_SCAN2.lsnr
      1          ONLINE  ONLINE      galaxy
ora.LISTENER_SCAN3.lsnr
      1          ONLINE  ONLINE      galaxy
.
.
.
```

To verify the Oracle RAC database installation, check the oraInventory logs.

Performing Oracle RAC post-installation tasks

This chapter includes the following topics:

- [Adding Oracle RAC patches or patchsets](#)
- [Configuring the CSSD resource](#)
- [Preventing automatic startup of Oracle Clusterware/Grid Infrastructure](#)
- [Relinking the SF Oracle RAC libraries with Oracle RAC](#)
- [Updating the CSS misscount setting](#)
- [Creating the Oracle RAC database](#)
- [Configuring VCS service groups for Oracle RAC](#)
- [Preventing automatic database startup](#)
- [Removing permissions for communication](#)
- [Configuring the SFDB repository database after installation](#)

Adding Oracle RAC patches or patchsets

If you installed Oracle 11.2.0.1, apply the Oracle patch 8649805 to integrate Oracle Clusterware/Grid Infrastructure with VCS.

To install the required patches or patchsets, review the notes that accompany the patch or patchset.

Before installing any Oracle RAC patch or patchset software:

- Review the latest information on supported Oracle RAC patches and patchsets:

www.symantec.com/docs/DOC4039

- You must have installed the base version of the Oracle RAC software.

Configuring the CSSD resource

You must configure the CSSD resource to ensure that the CSSD dependencies on the resources that manage OCR and voting disk and the private IP address are satisfied before Oracle Clusterware/Grid Infrastructure starts.

Note: It is mandatory to use CSSD agent in SF Oracle RAC installations. Using the CSSD agent ensures adequate handling of inter-dependencies, thus preventing the premature startup of Oracle Clusterware/Grid Infrastructure, which causes cluster failures.

Before you configure the CSSD resource, make sure that the following requirements are satisfied:


1. Oracle Clusterware/Grid Infrastructure is up and running.
2. OCR and voting disk is configured on CVM raw volumes or CFS and managed by VCS.
3. The private IP address for Oracle Clusterware/Grid Infrastructure is configured under the PrivNIC or MultiPrivNIC resource in the same VCS group as that of OCR and voting disk.

Note: The SF Oracle RAC installer configures the OCR, voting disk, and PrivNIC/MultiPrivNIC resources in the cvm group. If you configured one of these resources manually, make sure that these resources are placed in the cvm group. If the resources are not in the same group, configure the CSSD resource manually.

See [“Configuring the CSSD resource manually”](#) on page 348.

Figure 20-1 illustrates the configuration performed by the SF Oracle RAC installer. In the figure, the CSSD resource is configured under the CVM group.

Figure 20-1 CSSD configuration by SF Oracle RAC installer



Use one of the following ways to configure the CSSD resource:

SF Oracle RAC script-based installer

See [“Configuring the CSSD resource using the SF Oracle RAC script-based installer”](#) on page 345.

SF Oracle RAC Web-based installer

See [“Configuring the CSSD resource using the SF Oracle RAC Web-based installer”](#) on page 347.

Manual

See [“Configuring the CSSD resource manually”](#) on page 348.

Configuring the CSSD resource using the SF Oracle RAC script-based installer

Configure the CSSD resource using the SF Oracle RAC installer if the OCR and voting disk storage is configured on CFS.

Note: If the OCR and voting disk storage is configured on ASM disk groups, configure the CSSD resource manually.

The installer performs the following configuration tasks:

- Adds the CSSD resource to the VCS configuration in the cvm group.

Note: If the CSSD resource already exists, the installer enables reconfiguration of the resource by deleting the existing resource.

- Sets the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resource that manages the private IP address for Oracle Clusterware/Grid Infrastructure.
- Sets the dependency of the CSSD resource on the CFSMount or CVMVolDg resources that manage OCR and voting disk.
- Enables the CSSD resource and saves the new configuration.

To configure the CSSD resource using the SF Oracle RAC script-based installer

- 1 Start the SF Oracle RAC installer, if it is not already running. Select the option **Post Oracle Installation Tasks**.

```
1)  Configure SF Oracle RAC sub-components
2)  SF Oracle RAC Installation and Configuration Checks
3)  Prepare to Install Oracle
4)  Install Oracle Clusterware/Grid Infrastructure and Database
5)  Post Oracle Installation Tasks
6)  Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 5
```

- 2 Select the option **Configure CSSD agent**.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running on all the nodes.

- 3 Press **Return** to continue with the configuration. The installer reads the resource and group mappings for the CSSD agent from the VCS configuration file and displays the information.
- 4 Enter **y** to continue with the configuration. Review the messages as the installer configures the CSSD agent and sets the appropriate dependencies.
- 5 Press **Return** to return to the installer menu.
- 6 If the Oracle Clusterware/Grid Infrastructure and the Oracle database binaries are on CFS, set the dependencies between the cssd resource and the CFSMount resources for the binaries manually:

```
# hares -link cssd crsbin_mnt
# hares -link cssd orabin_mnt
```

Configuring the CSSD resource using the SF Oracle RAC Web-based installer

Configure the CSSD resource using the SF Oracle RAC installer if the OCR and voting disk storage is configured on CFS.

Note: If the OCR and voting disk storage is configured on ASM disk groups, configure the CSSD resource manually.

The installer performs the following configuration tasks:

- Adds the CSSD resource to the VCS configuration in the cvm group.

Note: If the CSSD resource already exists, the installer enables reconfiguration of the resource by deleting the existing resource.

- Sets the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resource that manages the private IP address for Oracle Clusterware/Grid Infrastructure.
- Sets the dependency of the CSSD resource on the CFSSMount or CVMVolDg resources that manage OCR and voting disk.
- Enables the CSSD resource and saves the new configuration.

To configure the CSSD resource using the SF Oracle RAC Web-based installer

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 2 Select the following menu options from the Select a task and product page:

Task	Configure a Product
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.
- 4 After the validation completes successfully, click **Next**.
- 5 Select **Post Oracle Installation Tasks** from the Select a Task page. Click **Next**.
- 6 Select **Configure CSSD agent** from the Select a Task page. Click **Next**.

- 7 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory. Click **Next**.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running on all the nodes.

- 8 Review and confirm the configuration information. The installer reads the resource and group mappings for the CSSD agent from the VCS configuration file and displays the information.
- 9 Click **Yes** to continue with the configuration.
- 10 If the Oracle Clusterware/Grid Infrastructure and the Oracle database binaries are on CFS, set the dependencies between the cssd resource and the CFSMount resources for the binaries manually:

```
# hares -link cssd crsbin_mnt
# hares -link cssd orabin_mnt
```

Configuring the CSSD resource manually

Add the `cssd` resource to the VCS configuration and set CSSD dependencies on the resources that manage OCR and voting disk and the private IP addresses for Oracle Clusterware/Grid Infrastructure.

Note: It is recommended that the OCR, voting disk, and PrivNIC/MultiPrivNIC resources be configured in the same VCS group as that of the `cssd` resource. If the resources are not in the same group, set the appropriate dependencies between the service groups.

To configure the CSSD resource

- 1 Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

- 2 Add the CSSD resource to the `cvm` group:

```
# hares -add cssd_resname Application cvm_grpname
```

3 Modify the CSSD resource attributes:

```
# hares -modify cssd_resname StartProgram \  
/opt/VRTSvcs/rac/bin/cssd-online  
# hares -modify cssd_resname StopProgram \  
/opt/VRTSvcs/rac/bin/cssd-offline  
# hares -modify cssd_resname MonitorProgram \  
/opt/VRTSvcs/rac/bin/cssd-monitor  
# hares -modify cssd_resname CleanProgram \  
/opt/VRTSvcs/rac/bin/cssd-clean  
# hares -modify cssd_resname Critical 0  
# hares -override cssd_resname OnlineWaitLimit  
# hares -modify cssd_resname OnlineWaitLimit 5  
# hares -override cssd_resname OfflineWaitLimit  
# hares -modify cssd_resname OfflineWaitLimit 3
```

4 Enable the CSSD resource:

```
# hares -modify cssd_resname Enabled 1
```

5 Set the dependency of the CSSD resource on the CFSSMount or CVMVolDg resources that manage OCR and voting disk.

If you configured OCR and voting disk on CVM raw volumes:

```
# hares -link cssd ocrvotevol_resname
```

If you configured OCR and voting disk on CFS:

```
# hares -link cssd ocrvotemnt_resname
```

6 Set the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resources that manage the private IP address for Oracle Clusterware/Grid Infrastructure.

If you configured the PrivNIC resource:

```
# hares -link cssd priv_resname
```

If you configured the MultiPrivNIC resource:

```
# hares -link cssd multipriv_resname
```

- 7 If the Oracle Clusterware/Grid Infrastructure and the Oracle database binaries are on CFS, set the dependencies between the CSSD resource and the CFSMount resources for the binaries manually:

```
# hares -link cssd clusbin_mnt_resname  
# hares -link cssd orabin_mnt_resname
```

- 8 Change the permission on the VCS configuration file to read-only mode:

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

Preventing automatic startup of Oracle Clusterware/Grid Infrastructure

The use of the CSSD agent is mandatory to ensure adequate handling of service group inter-dependencies and thereby prevent the premature startup of Oracle Clusterware/Grid Infrastructure. Therefore, disable automatic startup of Oracle Clusterware/Grid Infrastructure when the system starts.

To prevent automatic startup of Oracle Clusterware/Grid Infrastructure

- 1 Log in as the root user on each node in the cluster.
- 2 Disable automatic startup of Oracle Clusterware/Grid Infrastructure.

```
# $GRID_HOME/bin/crsctl disable crs
```

Relinking the SF Oracle RAC libraries with Oracle RAC

If you added or upgraded the Oracle patches, you must relink the SF Oracle RAC libraries—VCSMM and ODM—with Oracle. Relinking the libraries enables coordinated exchange of cluster membership information and protection of data.

Note: Symantec recommends that you relink the SF Oracle RAC libraries only after completing all the required patch additions, if any.

This release does not support Oracle RAC 11g Release 1 Clusterware. References to Oracle RAC 11g Release 1 in the procedures apply to the Oracle database alone.

Use one of the following ways to relink the libraries:

SF Oracle RAC script-based installer	See “Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC script-based installer” on page 351.
SF Oracle RAC Web-based installer	See “Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC Web-based installer” on page 352.
Manual	See “Relinking SF Oracle RAC libraries with Oracle RAC manually” on page 353.

Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC script-based installer

Perform the steps in the following procedure to relink the libraries using the SF Oracle RAC script-based installer.

To relink the SF Oracle RAC libraries with Oracle RAC

1 Return to the SF Oracle RAC installer menu, and select the option **Post Oracle Installation**.

- 1) Configure SF Oracle RAC sub-components
- 2) SF Oracle RAC Installation and Configuration Checks
- 3) Prepare to Install Oracle
- 4) Install Oracle Clusterware/Grid Infrastructure and Database
- 5) Post Oracle Installation Tasks
- 6) Exit SF Oracle RAC Configuration

Choose option: [1-6,q] (1) **5**

2 Select the option **Relink Oracle Database Binary**.

- 1) Configure CSSD agent
- 2) Relink Oracle Database Binary
- 3) Exit SF Oracle RAC Configuration
- b) Back to previous menu

Choose option: [1-3,b,q] (1) **2**

- 3 Provide the Oracle environment information—user name, group name, CRS_HOME (or GRID_HOME), ORACLE_HOME. Based on this information, the installer detects the version of Oracle installed.

Note: You need to provide the Oracle environment information only if you quit the installer after installing Oracle Clusterware/Grid Infrastructure and the Oracle database.

```
Enter Oracle UNIX user name: [b] oracle
Enter Oracle UNIX group name: [b] (oinstall)
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
Enter absolute path of Oracle Database Home directory: [b]
.
.
Do you want to continue? [y,n,q] (y)
```

- 4 Review and confirm the Oracle database information.
The installer starts relinking the libraries.

Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC Web-based installer

Perform the steps in the following procedure to relink the libraries using the SF Oracle RAC Web-based installer.

To relink the SF Oracle RAC libraries with Oracle RAC

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 86.
- 2 Select the following menu options from the Select a task and product page:

Task	Configure a Product
Product	Veritas Storage Foundation for Oracle RAC

Click **Next**.

- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.
- 4 After the validation completes successfully, click **Next**.

- 5 Select **Post Oracle Installation Tasks** from the Select a Task page. Click **Next**.
- 6 Select **Relink Oracle Database Binary** from the Select a Task page. Click **Next**.
- 7 Provide the following information:

Enter Oracle UNIX user name For example, **oracle**

Enter Oracle UNIX group name For example, **oinstall**

Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

Enter absolute path of Oracle Database Home directory Enter the full path of the Oracle database home directory.

Click **Next**. Confirm the detected Oracle version to proceed.

Note: You need to provide the Oracle environment information only if you quit the installer after installing Oracle Clusterware/Grid Infrastructure and the Oracle database.

- 8 Review and confirm the Oracle database information.
The installer starts relinking the libraries.

Relinking SF Oracle RAC libraries with Oracle RAC manually

The relinking process involves the following tasks:

- [Linking Veritas Membership library](#)
- [Linking the ODM library](#)

The Oracle RAC and Veritas library locations for linking the Oracle libraries with SF Oracle RAC are listed in the following tables. You may use these tables as reference when you relink the libraries as described in this section.

[Table 20-1](#) lists the Oracle RAC and Veritas library locations for Oracle RAC 11g Release 2.

Table 20-1 Oracle RAC and Veritas library locations for VCSMM and ODM - Oracle RAC 11g Release 2

Libraries	Oracle RAC library	Veritas library
VCSMM	<code>\$GRID_HOME/lib/libskgxn2.so</code>	<code>/opt/ORCLcluster/lib/libskgxn2.so</code>
ODM	<code>\$ORACLE_HOME/lib/libodm11.so</code>	<code>/usr/lib64/libodm.so</code>

Linking Veritas Membership library

You need to link the SF Oracle RAC libraries with Oracle RAC to enable coordinated exchange of cluster membership information and protection of data integrity. Oracle provides an API called `skgxn` (system kernel generic interface node membership) to obtain information on membership. SF Oracle RAC implements the API as a library linked to Oracle Grid Infrastructure after the Oracle Grid Infrastructure installation. Oracle uses the linked `skgxn` library (`libskgxn`) to make `ioctl` calls to VCSMM, which in turn obtains membership information for clusters and instances.

To link Veritas Membership library

- 1 Log in as the Oracle grid user.
- 2 Change to the `$GRID_HOME/lib` directory:
- 3 If the `$GRID_HOME/lib/libskgxn2.so` library is not linked to `/opt/ORCLcluster/lib/libskgxn2.so`, perform the following steps to link the library:

```
$ cd $GRID_HOME/lib

$ mv libskgxn2.so libskgxn2.so.oracle.`date +%m_%d_%Y-%H_%M_%S`

$ ln -s /opt/ORCLcluster/lib/libskgxn2.so libskgxn2.so
```

Linking the ODM library

Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries

on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

To link the Veritas ODM library

- 1 Log in as the Oracle user.
- 2 Change to the `$ORACLE_HOME/lib` directory:

```
$ cd $ORACLE_HOME/lib
```

- 3 Back up Oracle's ODM library:

```
$ mv libodm11.so libodm11.so.oracle-`date +%m_%d_%Y-%H_%M_%S`
```

- 4 Link the Veritas ODM library with Oracle's `libodm` library:

```
$ cp /usr/lib64/libodm.so libodm11.so
```

Updating the CSS misscount setting

Oracle RAC requires the CSS misscount value to be set to 600 seconds with vendor clusterware.

Check the current CSS misscount setting:

```
# $GRID_HOME/bin/crsctl get css misscount
```

If the current CSS misscount setting is not set to 600 seconds, update the setting as described in the following procedure.

To update the CSS misscount setting

- 1 Shut down Oracle Clusterware/Grid Infrastructure on all but one node.
- 2 Log in as the root user on the active node and set the CSS misscount interval to 600 seconds:

```
# $GRID_HOME/bin/crsctl set css misscount 600
```

- 3 Restart Oracle Clusterware/Grid Infrastructure.
- 4 Verify that the setting is updated:

```
# $GRID_HOME/bin/crsctl get css misscount  
600
```

Creating the Oracle RAC database

Create the Oracle RAC database on CVM raw volumes or CFS. Use your own tools or scripts, or review the guidelines on using the Oracle DBCA (Database Creation Assistant) tool to create the database.

For instructions on creating an Oracle RAC database:

See [“About creating a test database”](#) on page 673.

For more information, see the Oracle RAC documentation.

Note: If you plan to configure global clusters, then set up the Oracle RAC database only on the primary site. On the secondary site, the database will be replicated.

Configuring VCS service groups for Oracle RAC

Before you configure VCS service groups for Oracle databases, review the following information:

- Supported types of database management
See [“Supported types of database management”](#) on page 356.
- Sample service group configurations
See [“Sample service group configurations”](#) on page 357.

To configure the VCS service groups for Oracle RAC:

See [“Configuring VCS service groups manually for Oracle databases”](#) on page 361.

Supported types of database management

[Table 20-2](#) lists the database management options in Oracle RAC that are supported by SF Oracle RAC.

Table 20-2 Supported types of database management

Management type	Description
Policy-managed databases	In policy-managed databases, administrators specify the server pool on which the database instances run. Oracle Grid Infrastructure determines the server on which the database instances run. For more information, see the Oracle documentation.

Table 20-2 Supported types of database management (*continued*)

Management type	Description
Administrator-managed databases	In administrator-managed databases environments, the administrator specifies the servers on which the databases instances run. For more information, see the Oracle documentation.

Sample service group configurations

You can set up the Oracle database to be managed by one of the following clusterwares:

- Veritas Cluster Server

Note: Symantec recommends that the Oracle database be configured under VCS.

When the database is configured under VCS:

You can choose to configure the service group in a way that insulates all the databases from failure in any of the databases in the group.

VCS manages the start and stop sequence of the applications and the database.

See [“Sample service group configurations with the VCS Oracle agent”](#) on page 357.

- Oracle Clusterware/Grid Infrastructure

See [“Sample service group configurations without the VCS Oracle agent”](#) on page 360.

Sample service group configurations with the VCS Oracle agent

This section illustrates sample service group configurations with the VCS Oracle agent for multiple databases.

[Figure 20-2](#) illustrates a service group configuration with the VCS Oracle agent.

Figure 20-2 Service group configuration with the VCS Oracle agent





Figure 20-3 illustrates an alternate service group configuration with the VCS Oracle agent.

Figure 20-3 Service group configuration with the VCS Oracle agent (alternate configuration)



To configure the Oracle database under VCS, create Oracle service groups after installing Oracle RAC and creating a database.

See “Configuring VCS service groups manually for Oracle databases” on page 361.

Sample service group configurations without the VCS Oracle agent

Figure 20-4 illustrates a sample configuration in the absence of the VCS Oracle agent for single database configurations.

Figure 20-4 Service group configuration without VCS Oracle agent (single database)





Figure 20-5 illustrates a sample service group configuration in the absence of the VCS Oracle agent for multiple database configurations.

Figure 20-5 Service group configuration without the VCS Oracle agent (multiple databases)



In a service group configuration without the VCS Oracle agent, Oracle Clusterware/Grid Infrastructure controls the database. An online local firm dependency exists between the Oracle Clusterware/Grid Infrastructure group and the CVM group. When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware/Grid Infrastructure group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware/Grid Infrastructure, and starts Oracle Clusterware/Grid Infrastructure. Oracle Clusterware/Grid Infrastructure starts the database and the application is brought online.

Note: In a service group configuration without the VCS Oracle agent, when the system starts, all volumes and mount points **MUST** be online for the dependent service groups to be online.

Configuring VCS service groups manually for Oracle databases

This section describes the steps to configure the VCS service group manually for Oracle databases.

See [Figure 20-2](#) on page 358.

The following procedure assumes that you have created the database.

To configure the Oracle service group manually for Oracle databases

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the service group to the VCS configuration:

```
# hagrps -add oradb_grpname
```

- 3 Modify the attributes of the service group:

```
# hagrps -modify oradb_grpname Parallel 1
```

```
# hagrps -modify oradb_grpname SystemList node_name1 0 node_name2 1
```

```
# hagrps -modify oradb_grpname AutoStartList node_name1 node_name2
```

- 4 Add the CVMVolDg resource for the service group:

```
# hares -add oradb_dg_resname CVMVolDg oradb_grpname
```

- 5 Modify the attributes of the CVMVolDg resource for the service group:

```
# hares -modify oradb_dg_resname CVMGroup oradb_dgname
```

```
# hares -modify oradb_dg_resname CVMActivation sw
```

```
# hares -modify oradb_dg_resname CVMVolume oradb_volname
```

- 6 Add the CFSSMount resource for the service group:

```
# hares -add oradbmnt_resname CFSSMount oradb_grpname
```

- 7 Modify the attributes of the CFSSMount resource for the service group:

```
# hares -modify oradbmnt_resname MountPoint "oradb_mnt"
```

```
# hares -modify oradbmnt_resname BlockDevice \  
"/dev/vx/dsk/oradb_dgname/oradb_volname"
```

- 8 Add the Oracle RAC database instance to the service group:

```
# hares -add db_resname Oracle oradb_grpname
```

9 Modify the attributes of the Oracle resource for the service group:

```
# hares -modify db_resname Owner oracle
# hares -modify db_resname Home "db_home"
# hares -modify db_resname StartUpOpt SRVCTLSTART
# hares -modify db_resname ShutDownOpt SRVCTLSTOP
```

10 For administrator-managed databases, perform the following steps:

- Localize the Sid attribute for the Oracle resource:

```
# hares -local db_resname Sid
```

- Set the Sid attributes for the Oracle resource on each system:

```
# hares -modify db_resname Sid oradb_sid_node1 -sys node_name1
# hares -modify db_resname Sid oradb_sid_node2 -sys node_name2
```

11 For policy-managed databases, perform the following steps:

- Modify the attributes of the Oracle resource for the service group:

```
# hares -modify db_resname DBName db_name
# hares -modify db_resname ManagedBy POLICY
```

- Set the Sid attribute to the Sid prefix for the Oracle resource on all systems:

```
# hares -modify db_resname Sid oradb_sid_prefix
```

Note: The Sid prefix is displayed on the confirmation page during database creation. The prefix can also be determined by running the following command :

```
# grid_home/bin/crsctl status resource ora.db_name.db -f | grep
GEN_USR_ORA_INST_NAME@ | tail -1 | sed 's/.*=//' | sed
's/_[0-9]$//'
```

- Set the IntentionalOffline attribute for the resource to 1 and make sure that the health check monitoring is disabled:

```
# hares -override db_resname IntentionalOffline
# hares -modify db_resname IntentionalOffline 1
# hares -modify db_resname MonitorOption 0
```

- 12 Set the dependencies between the CFSSMount resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link oradbmnt_resname oradbdbg_resname
```

- 13 Set the dependencies between the Oracle resource and the CFSSMount resource for the Oracle service group:

```
# hares -link db_resname oradbmnt_resname
```

- 14 Create an online local firm dependency between the oradb1_grp service group and the cvm service group:

```
# hagrps -link oradb_grpname cvm_grpname online local firm
```

- 15 Enable the Oracle service group:

```
# hagrps -enableresources oradb_grpname
```

- 16 Change the cluster configuration to the read-only mode:

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

- 17 Bring the Oracle service group online on all the nodes:

```
# hagrps -online oradb_grpname -any
```

Note: For policy-managed databases: When VCS starts or when the administrator attempts to bring the Oracle resource online, if the server is not part of the server pool associated with the database, the resource will remain offline. If Oracle Grid Infrastructure decides to move the server from the server pool, the database will be brought offline by the Oracle Grid Infrastructure and the oracle resource moves to offline state.

For more information and instructions on configuring the service groups using the CLI:

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

Managing database restart after failure

When a database instance faults, it is observed that both Oracle Clusterware and VCS (through the Oracle agent) respond to the resource fault. When Oracle Clusterware detects an instance failure, it attempts to restart the instance based

on the value set for the `RESTART_ATTEMPTS` attribute in the database resource profile of Oracle Clusterware. At the same time, VCS detects the instance failure and faults the resource and takes action depending on the VCS configuration.

To avoid both clusterwares from independently responding to the fault, you need to modify the appropriate resource parameters as described in the following sections.

To manage database restart after failure

- 1 Log into one of the nodes in the cluster as the root user.
- 2 Perform one of the following steps:
 - Disable Oracle Clusterware/Grid Infrastructure from restarting the instance by modifying the `RESTART_ATTEMPTS` attribute as follows:

```
# crsctl modify resource db_resname -attr "RESTART_ATTEMPTS=0"
```

- Prevent VCS from faulting the resource until Oracle Clusterware/Grid Infrastructure exhausts its restart attempts.

```
# haconf -makerw
# hares -modify db_resname ToleranceLimit = tl_value
# haconf -dump -makero
```

where *tl_value* is a value greater than or equal to the value of the `RESTART_ATTEMPTS` attribute.

Location of VCS log files

You may want to review the log files at `/var/VRTSvcs/log/engine_A.log` for errors or status messages. When large amounts of data are written, multiple log files may be written, such as `engine_B.log`, `engine_C.log`, and so on. The `engine_A.log` contains the most recent data.

Preventing automatic database startup

Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent. If you configure the VCS service groups for Oracle, you need to prevent the Oracle database from starting automatically. The Oracle Clusterware/Grid Infrastructure and Oracle agent may attempt to start the database instance at the same time if the database mount is available. To prevent the Oracle database from starting automatically, you must change the management policy for the database from automatic to manual using the Oracle `SRVCTL` command. The command changes the `AUTO_START` attribute of the Oracle database and instance resources.

To prevent automatic database startup

- 1 Register the database, if not already registered.

```
$ srvctl add database -d db_name -o db_home \  
-p location_parameterfile -y manual
```

- 2 Once the database is registered, change the management policy for the database to manual.

Note: For Oracle RAC 11g Release 2 policy-managed databases, retain the default management policy, which is `automatic`.

```
$ srvctl stop database -d db_name  
$ srvctl modify database -d db_name -y manual
```

- 3 Start the database.

```
$ srvctl start database -d db_name
```

Removing permissions for communication

Make sure you completed the installation of SF Oracle RAC and the verification of disk support for I/O fencing. If you used `rsh`, remove the temporary `rsh` access permissions that you set for the nodes and restore the connections to the public network.

If the nodes use `ssh` for secure communications, and you temporarily removed the connections to the public network, restore the connections.

Configuring the SFDB repository database after installation

If you want to use the Storage Foundation for Databases (SFDB) tools, you must set up the SFDB repository after installing and configuring SF Oracle RAC and Oracle. For SFDB repository set up procedures:

See *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*

Upgrading Oracle RAC

This chapter includes the following topics:

- [Supported upgrade paths](#)
- [Preparing to upgrade Oracle RAC](#)
- [Upgrading Oracle RAC binaries](#)
- [Migrating the Oracle RAC database](#)

Supported upgrade paths

The information in this section and the following sections in this chapter are relevant only if you upgraded from Storage Foundation products to SF Oracle RAC 6.0.

[Table 21-1](#) lists the upgrade paths for Oracle RAC.

Table 21-1 Supported upgrade paths for Oracle RAC

From current version	Upgrade to
Oracle RAC 9i Release 2	Oracle RAC 10g Release 2 Oracle RAC 11g Release 2 Grid Infrastructure and Oracle RAC 11g Release 1 database
Oracle RAC 10g Release 1	Oracle RAC 10g Release 2 Oracle RAC 11g Release 2 Grid Infrastructure and Oracle RAC 11g Release 1 database Oracle RAC 11g Release 2
Oracle RAC 10g Release 2	Oracle RAC 11g Release 2 Grid Infrastructure and Oracle RAC 11g Release 1 database Oracle RAC 11g Release 2

Table 21-1 Supported upgrade paths for Oracle RAC (*continued*)

From current version	Upgrade to
Oracle RAC 11g Release 1	Oracle RAC 11g Release 2

Note: When you upgrade to a different version of Oracle RAC, make sure that the full path of the Oracle Clusterware/Grid Infrastructure home directory and the Oracle database home directory is different from the path where the existing version of Oracle RAC resides.

The upgrade procedure assumes that the beginning configuration includes the following components, and that these components are running on the cluster nodes:

- SF Oracle RAC 6.0
- A supported version of the operating system

Preparing to upgrade Oracle RAC

Complete the following preparatory tasks before you upgrade Oracle RAC:

1. Depending on the version you are upgrading from, perform the steps in one of the following sections:
See [“Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g”](#) on page 368.
2. Verify that the cluster configuration is compatible for upgrading Oracle RAC.
See [“Verifying the systems for Oracle RAC installation”](#) on page 320.

Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g

Perform the following tasks before upgrading Oracle RAC.

To prepare for upgrade from Oracle RAC 10g or Oracle RAC 11g

- 1 Take a hot or cold backup of the existing database.
- 2 Back up the existing Oracle home and central inventory.

- 3 If the Oracle RAC database is under VCS control, freeze the Oracle service groups to prevent VCS from reporting the resource as faulted when Oracle RAC stops and starts the database during the upgrade:

```
# haconf -makerw  
  
# hagrps -freeze oradb_grpname -persistent
```

- 4 Freeze the cvm service group to prevent VCS from reporting the resource as faulted when Oracle Clusterware is stopped and started during the upgrade:

```
# hagrps -freeze cvm_grpname -persistent  
  
# haconf -dump -makero
```

Upgrading Oracle RAC binaries

Review your Oracle installation manuals and the appropriate Oracle support Web sites before upgrading Oracle RAC.

To upgrade Oracle RAC binaries

- 1 Upgrade Oracle Clusterware.

Note: If you upgrade to Oracle RAC 11g Release 2, upgrade Oracle Clusterware to a new directory called the Oracle Grid Infrastructure home directory (GRID_HOME).

Starting with Oracle RAC 11g Release 2, ASM must reside in the Oracle Grid Infrastructure home directory. If you plan to upgrade ASM to Oracle RAC 11g Release 2, make sure that you upgrade it to run in the Oracle Grid Infrastructure home directory.

For instructions, see the Oracle RAC documentation.

- 2 Make sure that Oracle Clusterware is running.
- 3 Install the Oracle RAC database binaries.

For instructions, see the Oracle RAC documentation.

- 4 Complete the following post-installation tasks:
 - Add Oracle RAC patches or patchsets.
See [“Adding Oracle RAC patches or patchsets”](#) on page 343.
 - Relink the SF Oracle RAC libraries with Oracle RAC.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC”](#) on page 350.

- For upgrades from Oracle RAC 9i: Add the CSSD resource to the VCS configuration.

Migrating the Oracle RAC database

For instructions on migrating the existing Oracle RAC database, see the Oracle metalink documentation.

After migrating the database, complete the post-upgrade tasks:

See [“Performing post-upgrade tasks”](#) on page 370.

Performing post-upgrade tasks

Perform the steps in the following procedure to complete the upgrade.

To perform post-upgrade tasks

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 For upgrades from Oracle RAC 9i: Modify the resources in the Oracle service groups:

```
# hares -modify db_resname StartUpOpt SRVCTLSTART
# hares -modify db_resname ShutDownOpt SRVCTLSTOP
# hares -modify db_resname pfile "" -sys node_name1
# hares -modify db_resname pfile "" -sys node_name2
```

- 3 Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.
- 4 Modify the Oracle RAC configuration to prevent automatic startup of Oracle Clusterware.

See [“Preventing automatic startup of Oracle Clusterware/Grid Infrastructure”](#) on page 350.

- 5 Modify the Oracle RAC database configuration to prevent automatic database startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.

See [“Preventing automatic database startup”](#) on page 365.

- 6 Unfreeze the VCS service groups that were frozen earlier.

As root user, enter:

```
# hagrps -unfreeze oradb_grpname -persistent  
  
# hagrps -unfreeze cvm_grpname -persistent  
  
# haconf -dump -makero
```

Automated installation using response files

- [Chapter 22. About response files](#)
- [Chapter 23. Installing and configuring SF Oracle RAC using a response file](#)
- [Chapter 24. Configuring I/O fencing using a response file](#)
- [Chapter 25. Installing Oracle RAC using a response file](#)
- [Chapter 26. Installing SF Oracle RAC and Oracle RAC using a response file](#)
- [Chapter 27. Response file variable definitions](#)

About response files

This chapter includes the following topics:

- [About response files](#)
- [Response file syntax](#)
- [Guidelines for creating the SF Oracle RAC response file](#)
- [Installation scenarios for response files](#)

About response files

Use response files to standardize and automate installations on multiple clusters.

You can perform the following installation activities using a response file:

- Installing and configuring SF Oracle RAC
- Preparing the nodes for installation of Oracle RAC
- Installing Oracle RAC
- Relinking the SF Oracle RAC libraries with Oracle RAC
- Configuring the CSSD agent
- Upgrading SF Oracle RAC
- Uninstalling SF Oracle RAC

You can perform end-to-end installations or modular deployments.

For more information on modular deployments:

See [“Modular deployments using response files”](#) on page 375.

[Table 22-1](#) lists the various options available for creating or obtaining a response file.

Table 22-1 Options for obtaining a response file

Option	Description
Create a response file	Create a response file based on the response file template provided with SF Oracle RAC. The file is located at <code>/opt/VRTSvcs/rac/install</code>.
Reuse or customize the response files generated by an installation	The Veritas installation programs generate a response file during the installation, configuration, upgrade, or uninstallation of SF Oracle RAC and for installer-based pre-installation tasks for Oracle RAC. The response file generated by the installer is located in the following directory: <pre>/opt/VRTS/install/logs/installsfrac-installernumber\ /installsfrac-installernumber.response file</pre> You can reuse or customize the response files to perform a complete end-to-end installation or to perform a specialized activity, such as setting up the PrivNIC or MultiPrivNIC configuration on your clusters. All you need to do is update the response file variable definitions to enable or disable the options, depending on the task you want to perform. Note: Response files are not created if the tasks terminated abruptly or if you entered q to quit the installation. To generate the response file when you plan to discontinue a task, use the Exit SF Oracle RAC configuration option.
Use the <code>-makeresponsefile</code> option with the SF Oracle RAC installer	Create a response file by specifying the <code>-makeresponsefile</code> option with the SF Oracle RAC installer. Mount the product disc and navigate to the folder that contains the installation program. Start the installation program. <pre># ./installsfrac -makeresponsefile</pre> Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option. The response file is created in the directory <code>/opt/VRTS/install/logs/</code>. Note: You can use the <code>-makeresponsefile</code> option to create response files only for installing, configuring, or uninstalling SF Oracle RAC. For more information: See “About the <code>-makeresponsefile</code> option” on page 376.

At the end of the SF Oracle RAC installation, the following files are created:

- A log file that contains executed system commands and output.
- A summary file that contains the output of the installation scripts.
- Response files to be used with the `-responsefile` option of the installer.

Note: The SF Oracle RAC response files also contain VCS variables used for the installation and configuration of VCS.

For the VCS variable definitions, see the *Veritas Cluster Server Installation Guide*.

Modular deployments using response files

Modular deployments offer flexibility in planning your installation and configuration activities. You can choose to perform any installation or configuration task independent of other related activities. For example, you can disable Oracle user and group creation across all nodes in a cluster while installing Oracle RAC.

Modular deployments are supported for the following installation and configuration activities:

- Installing SF Oracle RAC
- Configuring SF Oracle RAC
- Creating Oracle user and group
- Creating storage for OCR and voting disk
- Configuring private IP addresses for Oracle RAC
- Installing Oracle Clusterware
- Installing Oracle database
- Oracle post-installation tasks:
 - Configuring CSSD agent
 - Relinking Oracle RAC libraries

The variables required for each of the above tasks are described in the chapter *Response file variable definitions* of this document.

All modular deployment activities require the following variable definitions:

```
$CFG{prod}  
$CFG{systems}
```

About the `-makeresponsefile` option

The SF Oracle RAC installer includes the option (`-makeresponsefile`) to generate sample response files that can be customized to perform installation, configuration, or uninstallation of the product.

Note: No actual installation, configuration, or uninstallation occurs when you use this option.

Use the option in the following situations:

- To understand the information that is required when you install, configure, or uninstall SF Oracle RAC
- To create a response file
The option creates a response file that can be used as a template for installing, configuring, or uninstalling SF Oracle RAC. You can customize the response file, as required.

To generate the response file for installing and configuring the product, specify the `-makeresponsefile` option with the installer or product installation script at the command line.

To generate the response file for uninstalling the product, specify the `-makeresponsefile` option with the installer or the product uninstall script at the command line.

Response file syntax

The Perl statement syntax that is included in the response file varies, depending on whether “Scalar” or “List” values are required by the variables.

For example,

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```


Guidelines for creating the SF Oracle RAC response file

This section provides guidelines for creating the SF Oracle RAC response file.

1. Create a response file using one of the available options.

For various options on creating or obtaining an SF Oracle RAC response file:

See [“About response files”](#) on page 373.

2. Set the following master values to 1 to enable SF Oracle RAC installation and configuration.

Note: The master settings must be set to 1 to enable the installer to read dependent variable definitions. For example, if the value `$CFG{opt}{install}` is not set to 1, the other dependent installation values in the response file will be disregarded. This is true for any master setting.

The following is the list of master values that must be set for installing and configuring SF Oracle RAC.

Installing SF Oracle RAC `$CFG{opt}{install}=1;`

Configuring SF Oracle RAC `$CFG{opt}{configure}=1;`
 `$CFG{config_sfrac_subcomponents}=1;`

3. Now, set the appropriate value in the dependent variable definitions for installing and configuring SF Oracle RAC.

The set of minimum definitions for a successful installation and configuration is as follows:

```
$CFG{accepteula}=1;
$CFG{opt}{install}=1;
$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SFRAC60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
```

```
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_11tlink1}{galaxy}="eth1";
$CFG{vcs_11tlink1}{nebula}="eth1";
$CFG{vcs_11tlink2}{galaxy}="eth2";
$CFG{vcs_11tlink2}{nebula}="eth2";
$CFG{vcs_userenpw}=[ qw(gpqIpkPmqLqqOyqKpn) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];
```

You can add more variable definitions, as required.

4. Set the following master values to 1 to enable the Oracle pre-installation tasks.

Note: This step assumes that you have already set the master value `$CFG{opt}{configure}` to 1 in the previous step.

The following is the list of additional master values that must be set for completing Oracle pre-installation tasks:

Private IP address configuration	If you plan to use the PrivNIC agent, use the following master settings: \$CFG{config_privnic}=1; \$CFG{config_multiprivnic}=0; If you plan to use the MultiPrivNIC agent, use the following master settings: \$CFG{config_multiprivnic}=1; \$CFG{config_privnic}=0;
OCR and voting disk storage	\$CFG{create_ocr_vote_storage}=1;
Oracle Clusterware/Grid Infrastructure installation	\$CFG{install_oracle_clusterware}=1;
Oracle database installation	\$CFG{install_oracle_database}=1;
Post Oracle tasks	\$CFG{config_cssd_agent}=1; \$CFG{relink_oracle_database}=1;

Set the master value for the `$CFG{create_oracle_user_group}` variable to 0 as follows:

```
$CFG{create_oracle_user_group}=0;
```

Note: If you are performing an end-to-end installation, the creation of Oracle user and group must be completed before starting the installation. This is because user equivalence must be manually established on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

5. Set the appropriate value for the dependent variable definitions for the Oracle pre-installation tasks.

Note: The following table discusses the variable definitions with sample values. Replace the sample values with those that are appropriate for your installation requirements.

Private IP address
configuration

If you plan to use the PrivNIC agent, provide the following definitions:

```
$CFG{privnic_resname} = "ora_priv";
$CFG{privnic_interface_priority}="eth2 eth3";
$CFG{galaxy}{privnicip} = "192.168.12.1";
$CFG{galaxy}{hostname_for_ip} = "galaxy-priv";
$CFG{nebula}{privnicip} = "192.168.12.2";
$CFG{nebula}{hostname_for_ip} = "nebula-priv";
$CFG{nic_netmask} = "255.255.255.0";
$CFG{nic_add_ip_to_files} = 1;
```

If you plan to use the MultiPrivNIC agent, use the following master settings:

```
$CFG{multiprivnic_resname} = "ora_priv";
$CFG{nic_add_ip_to_files} = 1;
$CFG{galaxy}{eth1}{multiprivnicip}
= "192.168.12.1";
$CFG{galaxy}{eth1}{hostname_for_ip}
= "galaxy-priv";
$CFG{galaxy}{eth2}{multiprivnicip}
= "192.168.2.1";
$CFG{galaxy}{eth2}{hostname_for_ip}
= "galaxy-priv1";
$CFG{nebula}{eth1}{multiprivnicip}
= "192.168.12.2";
$CFG{nebula}{eth1}{hostname_for_ip}
= "nebula-priv";
$CFG{nebula}{eth2}{multiprivnicip}
= "192.168.2.2";
$CFG{nebula}{eth2}{hostname_for_ip}
= "nebula-priv1";
$CFG{nic_netmask} = "255.255.255.0";
```

OCR and voting disk storage If you choose to create the OCR and voting disk storage:

```
$CFG{ocrvotedgoption}=0;
$CFG{ocrvotescheme} = 1;
$CFG{enable_mirroring} = 1;
$CFG->{enable_sep_filesys} = 0
$CFG{ocrvotedisks} = [ qw(Disk_1 Disk_2) ];
$CFG{ocrvotedgname} = "ocrvotedg";
$CFG{ocrvotevolname} = "ocrvotevol";
$CFG{ocrvotevolsize} = 640;
$CFG{ocrvotemount} = "/ocrvote";
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
```

If you choose to use an existing storage for the OCR and voting disk storage:

```
$CFG{ocrvotedgoption}=1;
$CFG{ocrvotescheme} = 0;
$CFG{enable_mirroring} = 1;
$CFG{ocrvotedgname} = "ocrvotedg";
$CFG{ocrvotevolname} = "ocrvotevol";
$CFG{ocrvotevolsize} = 640;
$CFG{ocrvotemount} = "/ocrvote";
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
```

**Oracle
Clusterware/Grid
Infrastructure
installation**

```
$CFG{oracle_user}="oracle";
$CFG{oracle_group} = "oinstall";
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home} = "/u01/app/11.2.0/grid";
$CFG{crs_installpath} = "/cdrom/oracle/\
clusterware";
$CFG{oracle_version} = "11.2.0.2";
$CFG{crs_responsefile} = "/oracle/crs.rsp";
```

**Oracle database
installation**

```
$CFG{oracle_user}="oracle";
$CFG{oracle_group} = "oinstall";
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home} = "/u01/app/11.2.0/grid";
$CFG{db_home} = "/u02/app/oracle/product/\
11.2.0/dbhome_1";
$CFG{db_installpath} = "/cdrom/oracle/database";
$CFG{oracle_version} = "11.2.0.2";
$CFG{db_responsefile} = "/oracle/db.rsp";
```

Post Oracle tasks

The master value (set in the previous step) suffices for CSSD resource configuration.

For relinking libraries, the dependent variable settings are as follows:

```
$CFG{oracle_user}="oracle";  
$CFG{oracle_group} = "oinstall";  
$CFG{crs_home} = "/u01/app/11.2.0/grid";  
$CFG{db_home} = "/u02/app/oracle/product/\n11.2.0/dbhome_1";  
$CFG{oracle_version} = "11.2.0.2";
```

Installation scenarios for response files

The chapters in this section cover the following installation scenarios using response files:

- Installing and configuring SF Oracle RAC
See [“Installing and configuring SF Oracle RAC”](#) on page 383.
- Installing Oracle RAC
See [“About installing Oracle RAC using response files”](#) on page 389.
- Installing both SF Oracle RAC and Oracle RAC
See [“Installing SF Oracle RAC and Oracle RAC”](#) on page 398.

Installing and configuring SF Oracle RAC using a response file

This chapter includes the following topics:

- [Installing and configuring SF Oracle RAC](#)
- [Sample response file for installing and configuring SF Oracle RAC](#)

Installing and configuring SF Oracle RAC

You can create a single response file or separate response files for installing and configuring SF Oracle RAC.

The installer performs the following tasks:

- Installs SF Oracle RAC.
- Configures SF Oracle RAC.

The following sample procedure uses a single response file for installing and configuring SF Oracle RAC.

To install and configure SF Oracle RAC using response files

- 1 Make sure that the systems meet the installation requirements.
- 2 Complete the preparatory steps before starting the installation.

For instructions, see the chapter "Preparing to install SF Oracle RAC" in this document.

- 3 Create a response file using one of the available options.

For information on various options available for creating a response file:

See [“About response files”](#) on page 373.

Note: You must replace the host names in the response file with that of the new systems in the cluster.

For guidelines on creating a response file:

See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 377.

For a sample response file:

See [“Sample response file for installing and configuring SF Oracle RAC”](#) on page 385.

- 4 Mount the product disc and navigate to the product directory that contains the installation program.
- 5 Start the installation and configuration:

```
# ./installsfrc -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

Note: Fencing is configured in disabled mode. You need to configure fencing after the configuration.

- 6 Configure I/O fencing.

Note: Before you configure I/O fencing, make sure that you complete the required pre-configuration tasks. For instructions, see the chapter *Preparing to configure SF Oracle RAC* in this document.

For instructions on configuring I/O fencing using a response file, see the chapter *Configuring I/O fencing using a response file* in this document.

- 7 Complete the SF Oracle RAC post-installation tasks.

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

Sample response file for installing and configuring SF Oracle RAC

The following sample response file installs and configures SF Oracle RAC on two nodes, galaxy and nebula.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{prod}="SFRAC60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents} = 1;
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_lltlink1}{galaxy}="eth1";
$CFG{vcs_lltlink1}{nebula}="eth1";
$CFG{vcs_lltlink2}{galaxy}="eth2";
$CFG{vcs_lltlink2}{nebula}="eth2";
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];
$CFG{vcs_userenpw}=[ qw(gpqIpkPmqLqqOyqKpn) ];
$CFG{uploadlogs}=0;

1;
```

Configuring I/O fencing using a response file

This chapter includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Sample response file for configuring server-based I/O fencing](#)

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for SF Oracle RAC.

To configure I/O fencing using response files

- 1 Make sure that SF Oracle RAC is configured.
- 2 Based on whether you want to configure disk-based or server-based I/O fencing, make sure you have completed the preparatory tasks.

See [“About planning to configure I/O fencing”](#) on page 59.

- 3 Copy the response file to one of the cluster systems where you want to configure I/O fencing.

See [“Sample response file for configuring disk-based I/O fencing”](#) on page 387.

See [“Sample response file for configuring server-based I/O fencing”](#) on page 387.

- 4 Edit the values of the response file variables as necessary.
 See [“Response file variables to configure disk-based I/O fencing”](#) on page 414.
 See [“Response file variables to configure server-based I/O fencing”](#) on page 415.
- 5 Start the configuration from the system to which you copied the response file.
 For example:

```
# /opt/VRTS/install/installsfrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 414.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SFRAC60";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{fencing_dgname}="fendg";
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{fencing_newdg_disks}=
  [ qw(sdx sdy sdz) ];
$CFG{fencing_option}=2;
```

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips>{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_scsi3_disk_policy}="raw";
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_ports>{"10.200.117.145"}=14250;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFRAC60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{fencing_option}=1;
```

Installing Oracle RAC using a response file

This chapter includes the following topics:

- [About installing Oracle RAC using response files](#)
- [Before you install](#)
- [Installing Oracle RAC](#)
- [Sample response file for installing Oracle RAC](#)

About installing Oracle RAC using response files

You can perform a silent pre-configuration of your systems by running the SF Oracle RAC installer with the `-responsefile` option. This capability in tandem with the Oracle response files for Oracle Clusterware and Oracle database installation enables you to standardize and automate Oracle RAC deployments in your cluster.

Before you install

Make sure that you complete the tasks in the following procedure before starting the silent installation of Oracle RAC.

To prepare the systems for installing Oracle RAC using response files

- 1 Make sure that the systems meet the installation requirements.
For information on requirements, see the chapter *System requirements* in this document.
- 2 Complete the following preparatory tasks on the nodes manually or by using the SF Oracle RAC installer:

- Identify the public virtual IP addresses for use by Oracle
- Plan for the storage and the private network configuration for Oracle RAC
- Set the kernel parameters
- Verify the user "nobody" exists
- Create Oracle user and groups
- Set up Oracle user equivalence
- Edit the Oracle user profile

For instructions, see the chapter *Before installing Oracle RAC* in this document.

- 3 Create response files for Oracle Clusterware/Grid Infrastructure and Oracle database installation using the response file template provided by Oracle RAC.

For instructions on using the response file template, see the Oracle RAC documentation.

Note: Keep at hand the full path of the directory where these response files will be saved. The full path of the Oracle Clusterware/Grid Infrastructure response file must be set in the SF Oracle RAC response file variable `$CFG{crs_responsefile};`. The full path of the Oracle database response file must be set in the SF Oracle RAC response file variable `$CFG{db_responsefile};`.

- 4 Create an SF Oracle RAC response file.

For information on various options available for creating a response file:

See [“About response files”](#) on page 373.

Make sure that you provide the full path information of the Oracle Clusterware/Grid Infrastructure and Oracle database response files.

For guidelines on creating a response file:

See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 377.

For information on the list of required and optional variables:

See [“Response file variable definitions for Oracle RAC”](#) on page 416.

For a sample response file:

See [“Sample response file for installing Oracle RAC”](#) on page 392.

- 5 Make sure that the Oracle user has appropriate read and write permissions on the response files.
- 6 Make sure that passwordless communication between Oracle users is set up on the nodes of the cluster.

Installing Oracle RAC

To perform a silent installation of Oracle RAC, run the SF Oracle RAC installer.

The installer supports completion of the following tasks using the response file:

- Creates storage for OCR and voting disk
- Configures the private network for Oracle RAC
- Installs Oracle Clusterware/Grid Infrastructure by leveraging the corresponding Oracle response file
- Installs Oracle database by leveraging the Oracle database response file
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

The sample procedure assumes that Oracle user equivalence is established on all the nodes in the cluster.

To perform Oracle pre-installation and installation on the nodes

- 1 Navigate to the directory containing the SF Oracle RAC installer:

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

- 2 Start the installation:

```
# ./installsfrc -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the SF Oracle RAC response file.

- 3 Run the `root.sh` script as the root user on the cluster nodes.

Note: Do not run the script simultaneously on your cluster nodes.

- 4 Complete the following Oracle post-installation tasks:
 - Add any patches or patchsets required by Oracle RAC .
 - Create the Oracle RAC database.

- Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.
- Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.
- Configure the VCS service groups for Oracle RAC.
- Verify the cluster.
- Remove the temporary communication permissions.

For instructions:

See the chapter *Performing Oracle RAC post-installation tasks*.

Sample response file for installing Oracle RAC

The following sample response file installs Oracle RAC 11g Release 2 and performs the following Oracle RAC pre-installation and installation tasks on two nodes, galaxy and nebula, in the cluster:

- Creates a disk group for OCR and voting disk storage
- Creates OCR and voting disk storage on raw volumes
- Configures the PrivNIC agent for high availability on both nodes
- Installs Oracle Clusterware
- Installs Oracle database
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

```
#
# Configuration Values:
#
our %CFG;

$CFG{prod}="SFRAC60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{configure}=1;

$CFG{config_privnic}=1;
$CFG{privnic_resname}="ora_priv";
$CFG{privnic_interface_priority}="eth2 eth3";
```



```
$CFG{galaxy}{privnicip}=" 192.168.12.1";
$CFG{galaxy}{hostname_for_ip}="galaxy-priv";
$CFG{nebula}{privnicip}=" 192.168.12.2";
$CFG{nebula}{hostname_for_ip}="nebula-priv";
$CFG{nic_netmask}="255.255.255.0";
$CFG{nic_add_ip_to_files}=1;
$CFG{nic_reuseip}=1;
$CFG{nic_reusealias}=1;

$CFG{create_ocr_vote_storage}=1;
$CFG{ocrvotedgoption}=0;
$CFG{oracle_group}="dba";
$CFG{grid_user}="grid";
$CFG{oracle_user}="oracle";
$CFG{ocrvolname}="ocrvotevol";
$CFG{ocrvolsize}=640;
$CFG{enable_mirroring} = 1;
$CFG{ocrvotedgname}="ocrvotedg";
$CFG{ocrvotedisks}=[qw(Disk_1 Disk_2)];
$CFG{ocrvotescheme}=0;
$CFG{votevolname}="votevol";
$CFG{votevolsize}=500;

$CFG{install_oracle_clusterware}=1;
$CFG{install_oracle_database}=1;
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home}="/u01/app/oracle/product/11.2.0/crshome";
$CFG{db_home}="/u01/app/oracle/product/11.2.0/dbhome_1";
$CFG{crs_installpath}="/cdrom/oracle/clusterware";
$CFG{db_installpath}="/cdrom/oracle/database";
$CFG{crs_responsefile}="/oracle/crs.rsp";
$CFG{db_responsefile} = "/oracle/db.rsp";

$CFG{config_cssd_agent}=1;
$CFG{relink_oracle_database}=1;
$CFG{uploadlogs}=0;

1;
```

Installing SF Oracle RAC and Oracle RAC using a response file

This chapter includes the following topics:

- [About installing SF Oracle RAC and Oracle RAC using response files](#)
- [Before you install](#)
- [Installing SF Oracle RAC and Oracle RAC](#)
- [Sample response file for installing SF Oracle RAC and Oracle RAC](#)

About installing SF Oracle RAC and Oracle RAC using response files

SF Oracle RAC integrates with the silent installation capabilities of the Oracle Universal Installer (OUI) to perform an end-to-end silent installation of SF Oracle RAC and Oracle RAC. The installation process uses the SF Oracle RAC response file to perform the installation, configuration, and Oracle pre-installation tasks in tandem with the Oracle RAC response files to install Oracle Clusterware and Oracle database. The SF Oracle RAC response file must be supplied with the full path of the location where the Oracle Clusterware and Oracle database response files reside.

You must have the following response files to perform a complete end-to-end installation:

An SF Oracle RAC installation and configuration response file	Contains the information listed in the following section: See “Information required in the SF Oracle RAC response file” on page 395. For various options on creating or obtaining an SF Oracle RAC response file: See “About response files” on page 373.
An Oracle Clusterware installation response file	Contains Oracle Clusterware installation information The response file may be created as described in the Oracle RAC documentation.
An Oracle database installation response file	Contains Oracle database installation information The response file may be created as described in the Oracle RAC documentation.

To perform the installation, you need to invoke the SF Oracle RAC installer with the `-responsefile` option.

Note: Some pre-installation and post-installation tasks must be manually completed before and after the installation, respectively. Complete these tasks as described in the procedures in this chapter.

Information required in the SF Oracle RAC response file

The SF Oracle RAC response file must contain the following information for an end-to-end installation:

- SF Oracle RAC installation information
- SF Oracle RAC configuration information
- Oracle RAC pre-installation information, such as follows:
 - Grid user information
 - Oracle user and group information
 - OCR and voting disk storage information
 - Private IP address configuration information
 - Full path to the Oracle Clusterware and Oracle database installation binaries
 - Full path of the Oracle Clusterware and Oracle database home directories
 - Full path of the Oracle Clusterware response file

- Full path of the Oracle database response file
- Oracle RAC post-installation information, such as follows:
 - CSSD resource configuration
 - Relinking Oracle RAC libraries with SF Oracle RAC

Before you install

SF Oracle RAC integrates with the silent installation capabilities of the Oracle Universal Installer (OUI) to provide an end-to-end installation of SF Oracle RAC and Oracle RAC. Complete the steps in the following procedure before you install SF Oracle RAC and Oracle RAC using response files.

To prepare to install SF Oracle RAC and Oracle RAC using response files

- 1 Make sure that the systems meet the installation requirements.
For information on requirements, see the chapter *System requirements* in this document.
- 2 Complete the preparatory tasks for installing SF Oracle RAC.
See [“Preparing to install Oracle RAC using the SF Oracle RAC installer or manually”](#) on page 266.
- 3 Set the kernel parameters for Oracle RAC.
For instructions, see the Oracle RAC documentation.
- 4 Create the users and groups required by Oracle.
See [“Creating users and groups for Oracle RAC”](#) on page 271.
- 5 Set up Oracle user equivalence on all nodes.
See [“Setting up user equivalence”](#) on page 320.

- 6 Create response files for Oracle Clusterware and Oracle database installation using the response file template provided by Oracle RAC.

For instructions on creating the response files for Oracle Clusterware and Oracle database, see the Oracle RAC documentation.

Make sure that the Oracle user has read/write permissions on the Oracle Clusterware and Oracle database response files.

Note: Keep at hand the full path of the directory where these response files are located. The full path of the Oracle Clusterware response file must be set in the SF Oracle RAC response file variable `$CFG{crs_responsefile};`. The full path of the Oracle database response file must be set in the SF Oracle RAC response file variable `$CFG{db_responsefile};`.

- 7 Create an SF Oracle RAC response file. Make sure that you provide the full path information of the Oracle Clusterware and Oracle database response files.

For guidelines on creating an SF Oracle RAC response file:

See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 377.

For the list of variable definitions:

For SF Oracle RAC variable definitions

For installing SF Oracle RAC

See [“Response file variables for installing SF Oracle RAC”](#) on page 402.

For configuring SF Oracle RAC

See [“Response file variables to configure Veritas Storage Foundation for Oracle RAC”](#) on page 405.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 414.

For Oracle RAC variable definitions

See [“Response file variable definitions for Oracle RAC”](#) on page 416.

For a sample response file:

See [“Sample response file for installing SF Oracle RAC and Oracle RAC”](#) on page 399.

Installing SF Oracle RAC and Oracle RAC

The installer performs the following tasks using the response file:

- Installs and configures SF Oracle RAC.

Note: Fencing is configured in disabled mode. Configure fencing after the installation of SF Oracle RAC and Oracle RAC is complete.

- Installs Oracle Clusterware/Grid Infrastructure and Oracle database.
- Configures the CSSD agent.
- Relinks the Oracle RAC libraries with SF Oracle RAC.

If any of the installation tasks fail, the installer quits the installation. To resume the installation:

- Review the log file to identify the cause of failure.
- Resolve the issue.
- Review the messages on the console to identify the tasks that have completed successfully. Disable these tasks in the response file.
- Run the installer again using the modified response file.

To install and configure SF Oracle RAC and Oracle RAC using a response file

- 1 Mount the product disc and navigate to the product directory that contains the installation program.
- 2 Start the installation and configuration:

```
# ./installsfprac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

Note: Fencing is configured in disabled mode. You need to configure fencing after the configuration.

- 3 Run the `root.sh` script as the root user on the cluster nodes.

Note: Do not run the script simultaneously on your cluster nodes.

- 4 After the installation completes, initialize disks as VxVM disks and configure fencing.

See [“Initializing disks as VxVM disks”](#) on page 140.

For instructions on configuring I/O fencing using a response file, see the chapter *Configuring I/O fencing using a response file* in this document.

- 5 Complete the SF Oracle RAC post-installation tasks.

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

- 6 Complete the following Oracle post-installation tasks:

- Add any patches or patchsets required by Oracle RAC.
- Create the Oracle RAC database.
- Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.
- Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.
- Configure the VCS service groups for Oracle RAC.
- Verify the cluster.
- Remove the temporary communication permissions.

For instructions:

See the chapter *Performing Oracle RAC post-installation tasks* in this document.

Sample response file for installing SF Oracle RAC and Oracle RAC

The sample response file installs SF Oracle RAC 6.0 and Oracle RAC 11g Release 2 and performs the following installation and configuration tasks on two nodes, galaxy and nebula, in the cluster:

- Installs all SF Oracle RAC RPMs
- Configures SF Oracle RAC with two private interconnects on each node
- Creates a disk group for OCR and voting disk storage
- Creates OCR and voting disk storage on CFS (located on same file system)

- Configures the PrivNIC agent for high availability on both nodes
- Installs Oracle Clusterware
- Installs Oracle database
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{prod}="SFRAC60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vxkeyless}=1;

$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents} = 1;
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_lltlink1}{galaxy}="eth1";
$CFG{vcs_lltlink1}{nebula}="eth1";
$CFG{vcs_lltlink2}{galaxy}="eth2";
$CFG{vcs_lltlink2}{nebula}="eth2";

$CFG{create_oracle_user_group}=0;

$CFG{config_privnic}=1;
$CFG{privnic_resname}="ora_priv";
$CFG{privnic_interface_priority}="eth2 eth3";
$CFG{galaxy}{hostname_for_ip}="galaxy-priv";
$CFG{galaxy}{privnicip}="192.168.12.1";
$CFG{nebula}{hostname_for_ip}="nebula-priv";
$CFG{nebula}{privnicip}="192.168.12.2";
$CFG{nic_netmask}="255.255.255.0";
$CFG{nic_add_ip_to_files}=1;

$CFG{create_ocr_vote_storage}=1;
```



```
$CFG{ocrvotedgoption}=0;
$CFG{oracle_group}="dba";
$CFG{grid_user}="grid";
$CFG{oracle_user}="oracle";
$CFG{enable_mirroring} = 1;
$CFG->{enable_sep_filesys} = 0
$CFG{ocrvotedgname}="ocrvotedg";
$CFG{ocrvotedisks}=[ qw(Disk_1 Disk_2) ];
$CFG{ocrvotemount}="/ocrvote";
$CFG{ocrvotescheme}=1;
$CFG{ocrvotevolname}="ocrvotevol";
$CFG{ocrvotevolsize}=640;

$CFG{install_oracle_clusterware}=1;
$CFG{install_oracle_database}=1;
$CFG{oracle_base}="/u01/app/oracle";
$CFG{crs_home}="/u01/app/oracle/product/11.2.0/crshome";
$CFG{db_home}="/u01/app/oracle/product/11.2.0/dbhome_1";
$CFG{crs_installpath}="/cdrom/oracle/clusterware";
$CFG{db_installpath}="/cdrom/oracle/database";
$CFG{crs_responsefile} = "/oracle/crs.rsp";
$CFG{db_responsefile} = "/oracle/db.rsp";
$CFG{oracle_display}="10.20.12.150:0.0";

$CFG{config_cssd_agent}=1;
$CFG{relink_oracle_database}=1;

$CFG{uploadlogs}=0;

1;
```

Response file variable definitions

This chapter includes the following topics:

- [Response file variables for installing SF Oracle RAC](#)
- [Response file variables to configure Veritas Storage Foundation for Oracle RAC](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Response file variables to configure server-based I/O fencing](#)
- [Response file variable definitions for Oracle RAC](#)

Response file variables for installing SF Oracle RAC

[Table 27-1](#) lists the response file variables that you can define to install SF Oracle RAC.

Table 27-1 Response file variables for installing SF Oracle RAC

Variable	List or Scalar	Description
CFG{opt}{install}	Scalar	Required Installs SF Oracle RAC RPMs.
CFG{opt}{systems}	List	Required List of systems on which the product is to be installed.

Table 27-1 Response file variables for installing SF Oracle RAC (*continued*)

Variable	List or Scalar	Description
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	Scalar	Required Instructs the installer to install SF Oracle RAC RPMs based on the variable that has the value set to 1: ■ installallpkgs: Installs all RPMs ■ installrecpkgs: Installs recommended RPMs ■ installminpkgs: Installs minimum RPMs Note: Set only one of these variable values to 1. In addition to setting the value of one of these variables, you must set the variable <code>\$CFG{opt}{install}</code> to 1.
CFG{opt}{rsh}	Scalar	Optional Defines that remote shell must be used instead of secure shell as the communication method between systems.
CFG{opt}{gco}	Scalar	Optional Defines that the installer must enable the global cluster option. You must set this variable value to 1 if you want to configure global clusters.
CFG{opt}{keyfile}	Scalar	Optional Defines the location of an ssh keyfile that is used to communicate with all remote systems.
CFG{opt}{pkgpath}	Scalar	Optional Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems.

Table 27-1 Response file variables for installing SF Oracle RAC (*continued*)

Variable	List or Scalar	Description
CFG{opt}{tmppath}	Scalar	Optional Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp.
CFG{opt}{logpath}	Scalar	Optional Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location.
CFG{opt}{donotinstall} {RPM}	List	Optional Instructs the installation to not install the optional RPMs in the list.
\$CFG{opt}{vxkeyless}	Scalar	Optional Defines a Boolean value 0 or 1. The value 1 indicates keyless installation. The value 0 indicates that a license key is required for the product. You must set the variable <i>\$CFG{keys}{hostname}</i> with appropriate values.
CFG{keys} {hostname}	Scalar	Optional List of keys to be registered on the system if the variable <i>\$CFG{opt}{vxkeyless}</i> is set to 0.

Table 27-1 Response file variables for installing SF Oracle RAC (*continued*)

Variable	List or Scalar	Description
\$CFG{uploadlogs}	Scalar	Optional Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site.

Response file variables to configure Veritas Storage Foundation for Oracle RAC

[Table 27-2](#) lists the response file variables that you can define to configure SF Oracle RAC.

Table 27-2 Response file variables specific to configuring Veritas Storage Foundation for Oracle RAC

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (Required) Set the value to 1 to configure SF Oracle RAC.
\$CFG{config_sfrac_subcomponents}	Scalar	Set the variable to 1 to configure the SF Oracle RAC components. (Required) Note: You must set the <i>\$CFG{opt}{configure}</i> variable to 1.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)

Table 27-2 Response file variables specific to configuring Veritas Storage Foundation for Oracle RAC (*continued*)

Variable	List or Scalar	Description
CFG{systems}	List	List of systems on which the product is to be configured. (Required)
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS60 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group

(csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtpsrsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table 27-3 lists the response file variables that specify the required information to configure a basic SF Oracle RAC cluster.

Table 27-3 Response file variables specific to configuring a basic SF Oracle RAC cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)

Table 27-4 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 27-4 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} { "system" }	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. You must enclose the system name within double quotes. (Required)

Table 27-4 Response file variables specific to configuring private LLT over Ethernet (*continued*)

Variable	List or Scalar	Description
CFG{vcs_1ltlinklowpri#} { <i>"system"</i> }	Scalar	Defines a low priority heartbeat link. Typically, 1ltlinklowpri is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, 1ltlinklowpri1, 1ltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

[Table 27-5](#) lists the response file variables that specify the required information to configure LLT over UDP.

Table 27-5 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{1lttoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)
CFG{vcs_udplink<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 27-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 27-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 27-6 lists the response file variables that specify the required information to configure virtual IP for SF Oracle RAC cluster.

Table 27-6 Response file variables specific to configuring virtual IP for SF Oracle RAC cluster

Variable	List or Scalar	Description
CFG{vcs_csgnic} {system}	Scalar	Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_csgvip}	Scalar	Defines the virtual IP address for the cluster. (Optional)
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

Table 27-7 lists the response file variables that specify the required information to configure the SF Oracle RAC cluster in secure mode.

Table 27-7 Response file variables specific to configuring SF Oracle RAC cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonenode}	Scalar	Specifies that the securityonenode option is being used.
CFG{securityonenode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. 1—Configure the first node 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.

[Table 27-8](#) lists the response file variables that specify the required information to configure VCS users.

Table 27-8 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)

Table 27-8 Response file variables specific to configuring VCS users (*continued*)

Variable	List or Scalar	Description
CFG{vcs_userpriv}	List	List of privileges for VCS users Note: The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

[Table 27-9](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 27-9 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtpsevl}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

[Table 27-10](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 27-10 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

[Table 27-11](#) lists the response file variables that specify the required information to configure SF Oracle RAC global clusters.

Table 27-11 Response file variables specific to configuring SF Oracle RAC global clusters

Variable	List or Scalar	Description
CFG{vcs_gconic} {system}	Scalar	Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_gcovip}	Scalar	Defines the virtual IP address to that the Global Cluster Option uses. (Optional)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Response file variables to configure disk-based I/O fencing

Table 27-12 lists the response file variables that specify the required information to configure disk-based I/O fencing for SF Oracle RAC.

Table 27-12 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. 1—Coordination Point Server-based I/O fencing 2—Coordinator disk-based I/O fencing 3—Disabled mode 4—Fencing migration when the cluster is online (Required)
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the I/O fencing mechanism. This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the fencing_scsi3_disk_policy variable and either the fencing_dgname variable or the fencing_newdg_disks variable. (Optional)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.

Table 27-12 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

[Table 27-13](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 27-13 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.

Table 27-13 Coordination point server (CP server) based fencing response file definitions (*continued*)

Response file field	Definition
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as "\$CFG{fencing_reusedg}=0" or "\$CFG{fencing_reusedg}=1" before proceeding with a silent installation.
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG {fencing_scsi3_disk_policy}	The disk policy that the customized fencing uses. The value for this field is either "raw" or "dmp"

Response file variable definitions for Oracle RAC

The variable definitions for Oracle RAC are grouped in tabular format for the following Oracle tasks:

Creating Oracle user and group

See [Table 27-14](#) on page 417.

Creating storage for OCR and voting disk	See Table 27-15 on page 418.
Configuring the private IP address and PrivNIC resource under VCS	See Table 27-16 on page 422.
Configuring the private IP address and MultiPrivNIC resource under VCS	See Table 27-17 on page 426.
Installing Oracle Clusterware	See Table 27-18 on page 429.
Installing Oracle database	See Table 27-19 on page 430.
Configuring CSSD resource	See Table 27-20 on page 431.
Relinking Oracle RAC libraries	See Table 27-21 on page 432.

Note: Some of the variable definitions may occur in multiple sections, for example `$CFG{oracle_user}`. These variables need not be repeated if all the tasks are performed as a single installation activity. However, if you perform these tasks independently, make sure that all the required variables, as indicated in the table for each task, are supplied in the response file.

[Table 27-14](#) lists the variables that are used to create the Oracle user and group.

Table 27-14 Variables for creating Oracle user and group

Variable	List or Scalar	Description
<code>\$CFG{create_oracle_user_group}</code>	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that Oracle user and group will be created. The value 0 indicates that Oracle user and group will not be created.
<code>\$CFG{grid_user}</code>	Scalar	Required Defines the name of the grid user.
<code>\$CFG{oracle_user}</code>	Scalar	Required Defines the name of the Oracle user.
<code>\$CFG{oracle_uid}</code>	Scalar	Required Defines the user ID of the Oracle user.

Table 27-14 Variables for creating Oracle user and group (*continued*)

Variable	List or Scalar	Description
\$CFG{oracle_group}	Scalar	Required Defines the primary group of the Oracle user.
\$CFG{oracle_gid}	Scalar	Required Defines the group ID of the Oracle user.
\$CFG{oracle_user_home}	Scalar	Required Defines the full path of the Oracle user's home directory.
\$CFG{oracle_secondary_group}	List	Optional Defines the list of secondary groups for the Oracle user.
\$CFG{oracle_secondary_gid}	List	Optional Defines the list of secondary group IDs for the Oracle user. The elements of this variable must be in the same order as that of the elements in the variable \$CFG{oracle_secondary_group}.

[Table 27-15](#) lists the variables that are used to create the storage for OCR and voting disk.

Table 27-15 Variables for creating storage for OCR and voting disk

Variable	List or Scalar	Description
\$CFG{create_ocr_vote_storage}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the storage for OCR and voting disk will be created. The value 0 indicates that the storage for OCR and voting disk will not be created.

Table 27-15 Variables for creating storage for OCR and voting disk (*continued*)

Variable	List or Scalar	Description
\$CFG{enable_mirroring}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the storage for OCR and voting disk is mirrored. Provide two disks as input for the variable \$CFG{ocrvotedisks}. The value 0 indicates that the storage for OCR and voting disk is not mirrored.
\$CFG{ocrvotedgoption}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that an existing disk group will be used to create the storage for OCR and voting disk. Note: If you choose to use an existing disk group, use the \$CFG{ocrvotedgname} variable to specify the name of an existing disk group that has a minimum of two disks (for mirroring). The value 0 indicates that a new disk group will be created for OCR and voting disk storage. Note: If you choose to create a disk group, you must set the following variables: \$CFG{ocrvotedisks}, \$CFG{ocrvotedgname}
\$CFG{ocrvotescheme}	Scalar	Required Defines the storage scheme to be used for OCR and voting disk. The value 1 indicates Clustered File System. The value 0 indicates CVM raw volumes.

Table 27-15 Variables for creating storage for OCR and voting disk (*continued*)

Variable	List or Scalar	Description
\$CFG->{enable_sep_filesys}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that OCR and voting disk are located on separate file systems. Provide values for the following variables: <pre>\$CFG->{ocrvolname}, \$CFG->{ocrvolsize}, \$CFG->{votevolname}, \$CFG->{votevolsize}, \$cfg->{ocrmount}, \$cfg->{votemount}</pre> The value 0 indicates that OCR and voting disk are located on the same file system. Provide a single volume name using the variable \$CFG{ocrvotevolname}, a single mount point using the variable \$CFG{ocrvotemount}, and the size using the variable \$CFG{ocrvotevolsize}.
\$CFG{ocrvotedisks}	List	Required Defines the list of shared disks to be used for OCR and voting disk.
\$CFG{ocrvotedgname}	Scalar	Required Defines the name of the disk group to be used for OCR and voting disk.
\$CFG{ocrvotevolname}	Scalar	Required Defines the volume name for OCR and voting disk. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).
\$CFG{ocrvotevolsize}	Scalar	Required Defines the size of the OCR and voting disk volume. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).

Table 27-15 Variables for creating storage for OCR and voting disk (*continued*)

Variable	List or Scalar	Description
\$CFG{ocrvotemount}	Scalar	Required if you have chosen to locate OCR and voting disk on the same file system. Defines the full path to the CFS mount point. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).
\$CFG{ocrmount}	Scalar	Required if you have chosen to locate OCR and voting disk on separate file systems. Defines the full path to the CFS mount point for OCR. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).
\$CFG{votemount}	Scalar	Required if you have chosen to locate OCR and voting disk on separate file systems. Defines the full path to the CFS mount point for voting disk. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).
\$CFG{ocrvolname}	Scalar	Required Defines the volume name for OCR. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).
\$CFG{ocrvolsize}	Scalar	Required Defines the size of the OCR volume. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).
\$CFG{votevolname}	Scalar	Required Defines the volume name for voting disk. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).

Table 27-15 Variables for creating storage for OCR and voting disk (*continued*)

Variable	List or Scalar	Description
\$CFG{votevolsize}	Scalar	Required Defines the size of the voting disk volume. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).
\$CFG{oracle_user}	Scalar	Required Defines the name of the Oracle user.
\$CFG{oracle_group}	Scalar	Required Defines the primary group of the Oracle user.

[Table 27-16](#) lists the variables that are used to configure the private IP address and PrivNIC resource under VCS.

Table 27-16 Variables for configuring the private IP address and PrivNIC resource under VCS

Variable	List or Scalar	Description
\$CFG{config_privnic}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the PrivNIC and private IP address information will be configured for Oracle Clusterware. The value 0 indicates that the PrivNIC and private IP address information will not be configured for Oracle Clusterware.
\$CFG{privnic_resname}	Scalar	Required Defines the PrivNIC resource name in the main.cf file.

Table 27-16 Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

Variable	List or Scalar	Description
\$CFG{privnic_interface_priority}	String	Required Defines the priority that determines which NIC will be used in the event of a failover. Set the priority in decreasing order. For example, the following priority setting indicates that eth2 will be given priority in the event of a failover: \$CFG{privnic_interface_priority}="eth2 eth3";
\$CFG{host1}{privnicip}	Scalar	Required Defines the IP address to be configured for the PrivNIC resource on the node. Repeat this variable for each node in the cluster. For example, if you have two nodes in the cluster, you must provide this variable for each node. For example: <pre>\$CFG{galaxy}{privnicip} ="192.168.12.1" \$CFG{nebula}{privnicip} ="192.168.12.2"</pre>
\$CFG{nic_reuseip}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the existing IP addresses in the /etc/hosts file will be used. The value 0 indicates that the IP addresses will not be reused.

Table 27-16 Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

Variable	List or Scalar	Description
\$CFG{host}{hostname_for_ip}	Scalar	Required Defines the private node name of the IP address (<i>hostname_for_ip</i>) for the PrivNIC resource and the node (<i>system</i>) for which the resource is configured. Repeat this variable for each node in the cluster. For example, if you have two nodes in the cluster, you must provide this variable for each node. For example: <pre>\$CFG{galaxy}{hostname_for_ip} ="galaxy-priv" \$CFG{nebula}{hostname_for_ip} ="galaxy-priv"</pre>
\$CFG{nic_netmask}	Scalar	Required Defines the netmask for the private network.
\$CFG{nic_add_ip_to_files}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the IP addresses are added to the /etc/hosts file. Note: Make sure that the IP addresses for the NIC resource are not already present in the files or set the <i>\$CFG{nic_reuseip}</i> and <i>\$CFG{nic_reusealias}</i> variables, otherwise the network configuration step fails. The value 0 indicates that the IP addresses may already be present in the file.

Table 27-16 Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

Variable	List or Scalar	Description
\$CFG{nic_reconfigure_existing_resource}	Scalar	Optional Defines a boolean value 0 or 1. The value 1 indicates that the existing PrivNIC resource in the main.cf file will be deleted and reconfigured. The value 0 indicates that the existing PrivNIC resource in the main.cf file will be reused.
\$CFG{nic_reusealias}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the installer will not check the <code>/etc/hosts</code> file to determine whether the host name alias for the private IP addresses exist or not. The installer assumes that the host names alias information is present in the file. Make sure that the alias information is present in the file. The value 0 indicates that the installer checks whether the host name alias information is present in the <code>/etc/hosts</code> file. Make sure that the alias information is present in the file otherwise the installation fails.

[Table 27-17](#) lists the variables that are used to configure the private IP addresses and the MultiPrivNIC resource under VCS.

Table 27-17 Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS

Variable	List or Scalar	Description
\$CFG{config_multiprivnic}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the MultiPrivNIC and private IP address information will be configured for Oracle Clusterware. The value 0 indicates that the MultiPrivNIC and private IP address information will not be configured for Oracle Clusterware.
\$CFG{multiprivnic_resname}	Scalar	Required Defines the MultiPrivNIC resource name in the main.cf file.
\$CFG{nic_add_ip_to_files}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the IP addresses are added to the /etc/hosts file. Note: Make sure that the IP addresses for the NIC resource are not already present in the files or set the <i>\$CFG{nic_reuseip}</i> and <i>\$CFG{nic_reusealias}</i> variables, otherwise the network configuration step fails. The value 0 indicates that the IP addresses may already be present in the file.

Table 27-17 Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS (*continued*)

Variable	List or Scalar	Description
\$CFG{ <i>host1</i> }{ <i>NIC1</i> }{multiprivnicip}	List	Required Defines the list of IP addresses for the MultiPrivNIC resource. Note: The private IP addresses must be configured for each node and each interface in the cluster. For example, if you have two nodes <i>galaxy</i> and <i>nebula</i> in the cluster: <pre>\$CFG{galaxy}{eth1} {multiprivnicip}="192.168.12.1"; \$CFG{galaxy}{eth2} {multiprivnicip}=="192.168.2.1"; \$CFG{nebula}{eth1} {multiprivnicip}="192.168.12.2"; \$CFG{nebula}{eth2} {multiprivnicip}="192.168.2.2";</pre>
\$CFG{ <i>host1</i> }{ <i>NIC1</i> }{hostname_for_ip}	List	Required Defines the list of private node names (<i>hostname_for_ip</i>) for the IP addresses configured in the MultiPrivNIC resource for the interface (<i>inf</i>) on the node (<i>system</i>) . Note: The private IP address must be configured for each node and each interface in the cluster. For example, if you have two nodes <i>galaxy</i> and <i>nebula</i> in the cluster: <pre>\$CFG{galaxy}{eth1} {hostname_for_ip}="galaxy-priv"; \$CFG{galaxy}{eth2} {hostname_for_ip}="galaxy-priv1"; \$CFG{nebula}{eth1} {hostname_for_ip}="nebula-priv"; \$CFG{nebula}{eth2} {hostname_for_ip}="nebula-priv1";</pre>

Table 27-17 Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS (*continued*)

Variable	List or Scalar	Description
\$CFG{nic_netmask}	Scalar	Required Defines the netmask for the private network.
\$CFG{nic_reconfigure_existing_resource}	Scalar	Optional Defines a boolean value 0 or 1. The value 1 indicates that the existing MultiPrivNIC resource in the main.cf file will be deleted and reconfigured. The value 0 indicates that the existing MultiPrivNIC resource in the main.cf file will be reused.
\$CFG{nic_reuseip}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the existing IP addresses in the /etc/hosts file will be used. The value 1 indicates that the existing IP addresses in the /etc/hosts or /etc/inet/ipnodes files will be used. The value 0 indicates that the IP addresses will not be reused.
\$CFG{nic_reusealias}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the installer will not check the /etc/hosts file to determine whether the host name alias for the private IP addresses exist or not. The installer assumes that the host names alias information is present in the file. Make sure that the alias information is present in the file. The value 0 indicates that the installer checks whether the host name alias information is present in the /etc/hosts file. Make sure that the alias information is present in the file otherwise the installation fails.

[Table 27-18](#) lists the variables that are used to install Oracle Clusterware.

Table 27-18 Variables for installing Oracle Clusterware

Variable	List or Scalar	Description
\$CFG{install_oracle_clusterware}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that Oracle Clusterware will be configured. The value 0 indicates that Oracle Clusterware will not be configured.
\$CFG{oracle_user}	Scalar	Required Defines the name of the Oracle user.
\$CFG{oracle_group}	Scalar	Required Defines the primary group of the Oracle user.
\$CFG{oracle_base}	Scalar	Required Defines the base directory for the Oracle RAC installation.
\$CFG{crs_home}	Scalar	Required Defines the Oracle Clusterware home directory. The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.
\$CFG{crs_installpath}	Scalar	Required Defines the full path of the Oracle Clusterware installation binaries.
\$CFG{oracle_version}	Scalar	Required Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.

Table 27-18 Variables for installing Oracle Clusterware (*continued*)

Variable	List or Scalar	Description
\$CFG{crs_responsefile}	Scalar	Required Defines the full path of the Oracle Clusterware response file.

[Table 27-19](#) lists the variables that are used to install Oracle database.

Table 27-19 Variables for installing Oracle database

Variable	List or Scalar	Description
\$CFG{install_oracle_database}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the Oracle RAC database will be configured. The value 0 indicates that the Oracle RAC database will not be configured.
\$CFG{oracle_user}	Scalar	Required Defines the name of the Oracle user.
\$CFG{oracle_group}	Scalar	Required Defines the primary group of the Oracle user.
\$CFG{oracle_base}	Scalar	Required Defines the base directory for the Oracle RAC installation.
\$CFG{crs_home}	Scalar	Required Defines the Oracle Clusterware home directory. The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.

Table 27-19 Variables for installing Oracle database (*continued*)

Variable	List or Scalar	Description
\$CFG{db_home}	Scalar	Required Defines the Oracle RAC database home directory. The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle RAC database response file.
\$CFG{db_installpath}	Scalar	Required Defines the full path of the Oracle RAC database installation binaries.
\$CFG{oracle_version}	Scalar	Required Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.
\$CFG{db_responsefile}	Scalar	Required Defines the full path of the Oracle database response file.

[Table 27-20](#) lists the variables that are used to configure CSSD resource.

Table 27-20 Variables for configuring CSSD resource

Variable	List or Scalar	Description
\$CFG{config_cssd_agent}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the CSSD agent will be configured after Oracle RAC installation. The value 0 indicates that the CSSD agent will not be configured after Oracle RAC installation.

Table 27-20 Variables for configuring CSSD resource (*continued*)

Variable	List or Scalar	Description
\$CFG{reconfigure_cssd_resource}	Scalar	Required Defines a boolean value 0 or 1. The value 1 indicates that the SF Oracle RAC installer deletes the existing CSSD resource from the main.cf file and reconfigures it. The value 0 indicates that the SF Oracle RAC installer does not delete and reconfigure the resource. The installer exits the task with the message that the resource exists.

[Table 27-21](#) lists the variables that are used to relink Oracle RAC libraries.

Table 27-21 Variables for relinking Oracle RAC libraries

Variable	List or Scalar	Description
\$CFG{relink_oracle_database}	Scalar	Required Defines a Boolean value 0 or 1. The value 1 indicates that the SF Oracle RAC libraries will be relinked with the Oracle RAC database after Oracle RAC installation. The value 0 indicates that the SF Oracle RAC libraries will not be relinked with the Oracle RAC database after Oracle RAC installation.
\$CFG{oracle_user}	Scalar	Required Defines the name of the Oracle user.
\$CFG{oracle_group}	Scalar	Required Defines the primary group of the Oracle user.

Table 27-21 Variables for relinking Oracle RAC libraries (*continued*)

Variable	List or Scalar	Description
\$CFG{crs_home}	Scalar	Required Defines the Oracle Clusterware home directory. The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.
\$CFG{db_home}	Scalar	Required Defines the Oracle RAC database home directory. The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle RAC database response file.
\$CFG{oracle_version}	Scalar	Required Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.

Adding or removing nodes from an SF Oracle RAC cluster

- [Chapter 28. Adding a node to SF Oracle RAC clusters](#)
- [Chapter 29. Removing a node from SF Oracle RAC clusters](#)

Adding a node to SF Oracle RAC clusters

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Adding a node to a cluster using the product installer](#)
- [Adding a node using the Web-based installer](#)
- [Adding the node to a cluster manually](#)
- [Configuring CVM and CFS on the new node](#)
- [Preparing the new node for installing Oracle RAC](#)
- [Adding the new node to Oracle RAC](#)
- [Adding nodes to a cluster that is using authentication for SFDB tools](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)
- [Sample configuration file for adding a node to the cluster](#)

About adding a node to a cluster

After you install SF Oracle RAC and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 16 nodes.

The following procedure provides a summary of the tasks required to add a node to an existing SF Oracle RAC cluster.

Table 28-1 Tasks for adding a node to a cluster

Step	Description
Complete the prerequisites and preparatory tasks before adding a node to the cluster.	See “Before adding a node to a cluster” on page 436.
Add a new node to the cluster: ■ Using the product installer ■ Using the Web installer ■ Manually	See “Adding a node to a cluster using the product installer” on page 439. See “Adding a node using the Web-based installer” on page 442. See “Adding the node to a cluster manually” on page 443.
Complete the configuration of the new node after adding it to the cluster.	
Prepare the new node for installing Oracle.	See “Preparing the new node for installing Oracle RAC” on page 466.
Add the node to Oracle.	See “Adding the new node to Oracle RAC” on page 490.
If you are using the Storage Foundation for Databases (SFDB) tools, you must update the repository database.	See “Adding nodes to a cluster that is using authentication for SFDB tools” on page 490. See “Updating the Storage Foundation for Databases (SFDB) repository after adding a node” on page 491.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing cluster, perform the required preparations.

- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

- 1 Review the hardware and software requirements.
For Storage Foundation for Oracle RAC (SF for Oracle RAC) or Storage Foundation for Sybase ASE CE:
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 Verify the existing cluster is running SF Oracle RAC 6.0 on all nodes.
- 4 If the cluster is upgraded from a previous version, you must check the cluster protocol version to make sure it has the same version as the node to be added. If there is a protocol mismatch, the node is unable to join the existing cluster.

Check the cluster protocol version using:


```
# vxdctl protocolversion  
Cluster running at protocol 110
```

- 5 If the cluster protocol is below 110, upgrade it on the master node using:

```
# vxdctl upgrade [version]
```

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 28-1](#).

Figure 28-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SF Oracle RAC private Ethernet controllers.
 - When you add nodes to a cluster, use independent switches or hubs for the private network connections.
 - If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.
- [Figure 28-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.
- 2 Connect the system to the shared storage as required. Make sure that you meet the following requirements:
 - The new node must be connected to the same shared storage devices as the existing nodes.
 - The node must have private network connections to two independent switches for the cluster.
 - The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Prepare the new node before you add it to an existing cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
# ./installsf rac -precheck
```

You can also use the Web-based installer for the precheck.

- 2 Install SF Oracle RAC on the new system using the `-install` option to install SF Oracle RAC. Do not configure SF Oracle RAC when prompted.
- 3 You can restart the new node after installation is complete. Configure the new node using the configuration from the existing cluster nodes.

See [“About installing SF Oracle RAC ”](#) on page 80.

Adding a node to a cluster using the product installer

You can add a node using the `-addnode` option with the product installer.

The product installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:


```
/etc/llttab
/etc/VRTSvcs/conf/sysname
```
- Copies the following files on the new node:


```
/etc/llthosts
/etc/gabtab
/etc/VRTSvcs/conf/config/main.cf
```
- Copies the following files on the new node:


```
/etc/llthosts
/etc/gabtab
/etc/VRTSvcs/conf/config/main.cf
/etc/vxfenmode
/etc/vxfendg
/etc/vx/.uuids/clusuid
/etc/sysconfig/llt
```

```
/etc/sysconfig/gab
/etc/sysconfig/vxfen
/etc/vcsmmtab
```

- Generates security credentials on the new node if the CPS server of existing cluster is secure
- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.
Configures disk-based fencing.
- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

Note:

- Starts product processes and configures CVM and CFS on the new node.

Note:

At the end of the process, the new node joins the existing cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

Caution: If you plan to use the product installer for completing the Oracle RAC pre-installation tasks on the new node, do not quit the installer after adding the node to the cluster. If you quit the installer, you must perform the Oracle RAC pre-installation tasks manually.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the product installer with the `-addnode` option.

```
# cd /opt/VRTS/install
# ./installsfrc -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing cluster.

The installer uses the node information to identify the existing cluster.

Enter one node of the SF Oracle RAC cluster to which
 you would like to add one or more new nodes: **galaxy**

- 4 Review and confirm the cluster information.
- 5 Enter the names of the systems that you want to add as new nodes to the cluster.

Enter the system names separated by spaces
 to add to the cluster: **saturn**

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

Enter the NIC for the first private heartbeat
 link on saturn: [b,q,?] **eth1**

Enter the NIC for the second private heartbeat
 link on saturn: [b,q,?] **eth2**

Note: At least two private heartbeat links must be configured for high availability of the cluster.

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The product installer verifies the network interface settings and displays the information.

- 8 Review and confirm the information.
- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

Enter the NIC for VCS to use on saturn: **eth3**

- 10 The installer prompts you with an option to mount the shared volumes on the new node. Select **y** to mount them.

When completed, the installer confirms the volumes are mounted indicates the location of the log file with details of the actions performed.

- 11 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.

The installer then starts all the required Veritas processes and joins the new node to cluster.

Note: Do not quit the installer if you want to perform the Oracle pre-installation tasks using the product installer.

- 12 Confirm that the new node has joined the existing cluster using `lltstat -n` and `gabconfig -a` commands.

Adding a node using the Web-based installer

You can use the Web-based installer to add a node to a cluster.

To add a node to a cluster using the Web-based installer

- 1 From the Task pull-down menu, select **Add a Cluster node**.

From the product pull-down menu, select the product.

Click the **Next** button.

- 2 Click **OK** to confirm the prerequisites to add a node.

- 3

In the System Names field enter a name of a node in the cluster where you plan to add the node and click **OK**.

The installer program checks inter-system communications and compatibility. If the node fails any of the checks, review the error and fix the issue.

If prompted, review the cluster's name, ID, and its systems. Click the **Yes** button to proceed.
- 4

In the System Names field, enter the names of the systems that you want to add to the cluster as nodes. Separate system names with spaces. Click the **Next** button.

The installer program checks inter-system communications and compatibility. If the system fails any of the checks, review the error and fix the issue.

Click the **Next** button. If prompted, click the **Yes** button to add the system and to proceed.
- 5

From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.
- 6

Once the addition is complete, review the log files. Optionally send installation information to Symantec. Click the **Finish** button to complete the node's addition to the cluster.

Adding the node to a cluster manually

Perform the following tasks in order after you install to add the node to the cluster manually.

Table 28-2 Tasks for adding a cluster manually

Task	Description
Start the Volume Manager.	See “Starting Volume Manager on the new node” on page 456.
Add the cluster files.	Configure LLT, GAB, and VCSMM. See “Configuring cluster files on the new node” on page 456.
Generate security credentials.	If the CP server of the existing cluster is secure, generate security credentials on the new node. See “Setting up the node to run in secure mode” on page 458.

Table 28-2 Tasks for adding a cluster manually (*continued*)

Task	Description
Set up fencing.	Configure fencing for the new node to match the fencing configuration on the existing cluster. If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node. See “Starting fencing on the new node” on page 461.
Start VCS.	See “Starting VCS after adding the new node” on page 462.
Configure Cluster Volume Manager (CVM) and Cluster File System (CFS).	If CVM and CFS are configured on the existing cluster, configure them on the new node. See “Configuring CVM and CFS on the new node” on page 465.
Configure the ClusterService group.	If the ClusterService group is configured on the existing cluster, add the node to the group. See “Configuring the ClusterService group for the new node” on page 455.

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfrc` program.

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfrc` program.

To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
# vxinstall
```

- 2 Enter **n** when prompted to set up a system wide disk group for the system.
The installation completes.

- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

Configuring cluster files on the new node

Perform the steps in the following procedure to configure LLT, GAB, and VCSMM on the new node.

To configure LLT, GAB, and VCSMM on the new node

- 1 For Red Hat Linux, modify the file `/etc/sysctl.conf` on the new system to set the shared memory and other parameter required by Oracle; refer to the Oracle documentation for details. The value of the shared memory parameter is put to effect when the system restarts.

Do not apply for SUSE Linux.

- 2 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
```

- 3 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.
- 4 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - 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.

- 5 Use `vi` or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where `N` represents the number of systems in the cluster including the new node. For a three-system cluster, `N` would equal 3.

- 6 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.

- 7 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/llt
/etc/sysconfig/gab
/etc/sysconfig/vcs
/etc/sysconfig/vcsmm
```

- 8 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 9 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \
-from_sys galaxy -to_sys saturn
```

- 10 Start the LLT, GAB, VCSMM, and ODM drivers on the new node:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vcsmm start
# /etc/init.d/vxgms start
# /etc/init.d/vxglm start
# /etc/init.d/vxodm restart
```

- 11 On the new node, verify that the GAB port memberships are a, d, and o:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

```
Port a gen      df204 membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012
```

Port h may also be present. In that case, run:

```
# hstop -local
```

Configuring cluster files on the new node

Perform the steps in the following procedure to configure cluster files on the new node.

To configure cluster files on the new node

- 1 If you are using Red Hat Linux with SF Oracle RAC, modify the file `/etc/sysctl.conf` on the new system to set the shared memory and other parameter required by your application; refer to the your Oracle RAC documentation for details. The value of the shared memory parameter is put to effect when the system restarts.

Do not apply this step for SUSE Linux.

- 2 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
```

- 3 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.
- 4 Create an `/etc/llttab` file on the new system.

For example:

```
set-node saturn
set-cluster 101
```

For AIX:

For HP-UX:

For Linux:

```
link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

- 5 For Solaris:

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.

- 6 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 7 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 8 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/llt  
/etc/sysconfig/gab  
/etc/sysconfig/vcs  
/etc/sysconfig/vcsmm
```

- 9 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 10 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```


Start the cluster components

1 Start LLT.

For AIX:

For HP-UX:

For Linux:

```
# /etc/init.d/llt start
```

For Solaris:

2 Start GAB.

For AIX:

For HP-UX:

For Linux:

```
# /etc/init.d/gab start
```

For Solaris:

3 Start disk-based fencing.

For AIX:

For HP-UX:

For Linux:

For Solaris:

4 Start VCSMM.

For AIX:

For HP-UX:

For Linux:

```
# /etc/init.d/vcsmm start
```

```
# /etc/init.d/vxgms start
```

```
# /etc/init.d/vxglm start
```

For Solaris:

5 Start LMX.

For AIX:

For HP-UX:

For Solaris:

6 Start ODM.

For AIX:

For HP-UX:

For Linux:

```
# /etc/init.d/vxodm restart
```

For Solaris:

7 Start the LLT, GAB, VCSMM, and ODM drivers on the new node.

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vcsmm start
# /etc/init.d/vxgms start
# /etc/init.d/vxglm start
# /etc/init.d/vxodm restart
```

8 On the new node, verify that the GAB port memberships are a, d, and o:

```
# gabconfig -a
```

GAB Port Memberships

```
=====
Port a gen      df204 membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012
```

Port h may also be present. In that case, run:

```
# hastop -local
```

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

```
/etc/sysconfig/vxfen  
/etc/vxfendg  
/etc/vxfenmode
```

For Solaris:

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:
- 3 On the new node, verify that the GAB port memberships are a, b, d, h and o:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

Port a gen	df204 membership 012
Port b gen	df20d membership 012
Port d gen	df20a membership 012
Port o gen	df207 membership 012

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

```
/etc/sysconfig/vxfen  
/etc/vxfendg  
/etc/vxfenmode
```

For Solaris:

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:
- 3 On the new node, verify that the GAB port memberships are a, b, d, h and o:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012
```

Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

To configure CVM and CFS on the new node

- 1 Make a backup copy of the main.cf file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagrps -modify cvm SystemList -add saturn 2
# hagrps -modify cvm AutoStartList -add saturn
# hares -modify cvm_clus CVMNodeId -add saturn 2
# haconf -dump -makero
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
saturn:/etc/VRTSvcs/conf/config/main.cf
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

- 7 The `/etc/vx/tunefstab` file sets non-default tunables for local-mounted and cluster-mounted file systems.

If you have configured a `/etc/vx/tunefstab` file to tune cluster-mounted file systems on any of the existing cluster nodes, you may want the new node to adopt some or all of the same tunables.

To adopt some or all tunables, review the contents of the file, and copy either the file, or the portions desired, into the `/etc/vx/tunefstab` file on the new cluster node.

Starting VCS after adding the new node

If you added the new node to the cluster manually, before you starting VCS you must create the file `cssd-pretend-offline` on the new node and make the `cssd` resource non-critical. Failing this, the `cssd` resource lapses into an UNKNOWN state until Oracle Clusterware is installed on the new node, thus preventing the `cvm` group from coming online.

If you used the product installer to add the new node to the cluster, you will not need to perform this step.

Note: The `cssd` resource will remain in FAULTED/OFFLINE state till Oracle Clusterware is installed on the new node.

Start VCS on the new node.

To start VCS on the new node

- 1 If you installed the new node manually:
 - On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Create the file `cssd-pretend-offline` on the new node:

```
# touch /var/VRTSvcs/lock/cssd-pretend-offline
```

- 2 Start VCS on the new node:

```
# hstart
```

VCS brings the CVM group online.

- 3 Verify that the CVM group is online:

```
# hagr -state
```

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

- 1 On an existing node, for example galaxy, write-enable the configuration:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing ClusterService group.

```
# hagr -modify ClusterService SystemList -add saturn 2
```

```
# hagr -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device eth0 -sys saturn
```

```
# hares -modify gconic Device eth0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfprac` program.

To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
# vxinstall
```

- 2 Enter **n** when prompted to set up a system wide disk group for the system.
The installation completes.

- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

Configuring cluster files on the new node

Perform the steps in the following procedure to configure LLT, GAB, and VCSMM on the new node.

To configure LLT, GAB, and VCSMM on the new node

- 1 For Red Hat Linux, modify the file `/etc/sysctl.conf` on the new system to set the shared memory and other parameter required by Oracle; refer to the Oracle documentation for details. The value of the shared memory parameter is put to effect when the system restarts.

Do not apply for SUSE Linux.

- 2 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
```

- 3 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.

- 4 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - 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.

- 5 Use `vi` or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where `N` represents the number of systems in the cluster including the new node. For a three-system cluster, `N` would equal 3.

- 6 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 7 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/llt
/etc/sysconfig/gab
/etc/sysconfig/vcs
/etc/sysconfig/vcsmm
```

- 8 Use `vi` or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 9 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \
-from_sys galaxy -to_sys saturn
```

10 Start the LLT, GAB, VCSMM, and ODM drivers on the new node:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vcsmm start
# /etc/init.d/vxgms start
# /etc/init.d/vxglm start
# /etc/init.d/vxodm restart
```

11 On the new node, verify that the GAB port memberships are a, d, and o:

```
# gabconfig -a

GAB Port Memberships
=====
Port a gen      df204 membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012
```

Port h may also be present. In that case, run:

```
# hastop -local
```

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

[Table 28-3](#) uses the following information for the following command examples.

Table 28-3 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.

Configuring the authentication broker on node saturn

To configure the authentication broker on node saturn

- 1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvc/vcsauth/bkup  
# cd /tmp; gunzip -c /opt/VRTSvc/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# /cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvc/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvc/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \
./broker_setup.sh/tmp/eat_setup

/opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \
/VRTSatlocal.conf -b 'Security\Authentication \
\Authentication Broker' -k UpdatedDebugLogFileName \
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

4 Copy the broker credentials from one node in the cluster to saturn by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

CMDSERVER  CPSADM  CPSEVER  HAD  VCS_SERVICES  WAC
```

5 Import the `VCS_SERVICES` domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \
/VCS_SERVICES -p password
```

6 Import the credentials for `HAD`, `CMDSERVER`, `CPSADM`, `CPSEVER`, and `WAC`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \
/HAD -p password
```

7 Start the `vcsauthserver` process on saturn.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvc/vcsauth/data/CLIENT  
  
# mkdir /var/VRTSvc/vcsauth/data/TRUST  
  
# export EAT_DATA_DIR='/var/VRTSvc/vcsauth/data/TRUST'  
  
# /opt/VRTSvc/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \  
localhost:14149 -s high
```

9 Create the `/etc/VRTSvc/conf/config/.secure` file:

```
# touch /etc/VRTSvc/conf/config/.secure
```

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

```
/etc/sysconfig/vxfen
/etc/vxfendg
/etc/vxfenmode
```

For Solaris:

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:
- 3 On the new node, verify that the GAB port memberships are a, b, d, h and o:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen    df204 membership 012
Port b gen    df20d membership 012
Port d gen    df20a membership 012
Port o gen    df207 membership 012
```

Starting VCS after adding the new node

If you added the new node to the cluster manually, before you starting VCS you must create the file `cssd-pretend-offline` on the new node and make the `cssd` resource non-critical. Failing this, the `cssd` resource lapses into an UNKNOWN state until Oracle Clusterware is installed on the new node, thus preventing the `cvm` group from coming online.

If you used the product installer to add the new node to the cluster, you will not need to perform this step.

Note: The `cssd` resource will remain in FAULTED/OFFLINE state till Oracle Clusterware is installed on the new node.

Start VCS on the new node.

To start VCS on the new node

- 1 If you installed the new node manually:
 - On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Create the file `cssd-pretend-offline` on the new node:

```
# touch /var/VRTSvcs/lock/cssd-pretend-offline
```

- 2 Start VCS on the new node:

```
# hstart
```

VCS brings the CVM group online.

- 3 Verify that the CVM group is online:

```
# hagrps -state
```

Configuring server-based fencing on the new node

This section describes the procedures to configure server-based fencing on a new node. Depending on whether server-based fencing is configured in secure or non-secure mode on the existing cluster, perform the tasks in one of the following procedures:

- Server-based fencing in non-secure mode:
[To configure server-based fencing in non-secure mode on the new node](#)
- Server-based fencing in secure mode:
[To configure server-based fencing with security on the new node](#)

To configure server-based fencing in non-secure mode on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c rac_cluster101 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com \  
-a list_nodes
```

The new node must be listed in the command output.

- 4 Add the VCS user cpsclient@saturn to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx
```

```
User cpsclient@saturn successfully added
```

To configure server-based fencing with security on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c rac_cluster101 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes
```

The new node must be listed in the output.

Adding the new node to the vxfen service group

Perform the steps in the following procedure to add the new node to the vxfen service group.

To add the new node to the vxfen group using the CLI

- 1 On one of the nodes in the existing SF Oracle RAC cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing vxfen group.

```
# hagrps -modify vxfen SystemList -add saturn 2
```

- 3 Save the configuration by running the following command from any node in the SF Oracle RAC cluster:

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

Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

To configure CVM and CFS on the new node

- 1 Make a backup copy of the main.cf file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagr -modify cvm SystemList -add saturn 2
# hagr -modify cvm AutoStartList -add saturn
# hares -modify cvm_clus CVMNodeId -add saturn 2
# haconf -dump -makero
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
saturn:/etc/VRTSvcs/conf/config/main.cf
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

- 7 The `/etc/vx/tunefstab` file sets non-default tunables for local-mounted and cluster-mounted file systems.

If you have configured a `/etc/vx/tunefstab` file to tune cluster-mounted file systems on any of the existing cluster nodes, you may want the new node to adopt some or all of the same tunables.

To adopt some or all tunables, review the contents of the file, and copy either the file, or the portions desired, into the `/etc/vx/tunefstab` file on the new cluster node.

Preparing the new node for installing Oracle RAC

You must complete certain pre-installation tasks before you add the node to Oracle RAC.

Note: The instructions in this chapter use variables and sample values wherever required. Replace these variables and sample values with values that conform to your installation requirements.

Before you start the preparatory tasks, you may want to update the sample worksheet with the correct installation values and keep them handy during the process.

Use one of the following ways to complete the preparatory tasks:

Using the SF Oracle RAC script-based installer	See “Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC script-based installer” on page 467.
Using the SF Oracle RAC Web-based installer	See “Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC Web-based installer” on page 474.
Manual	See “Preparing the new node manually for installing Oracle RAC” on page 482.

Note: Some of the pre-installation tasks can be completed using the SF Oracle RAC installer while some of the tasks must be completed manually as indicated in the procedures.

Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC script-based installer

The SF Oracle RAC installer performs the following tasks:

- Creates the Oracle user and groups on the new node
- Configures the private IP addresses and the PrivNIC or MultiPrivNIC resources (if they are configured in the existing cluster).
- If the CFSMount and CVMVoIDg resources for OCR and voting disk are configured under the `cvm` service group, the installer brings them online after adding the node to the cluster.
 If the resources are configured in any other service group, make sure that you modify the service group to include the new node and bring the service group online.

Note: If OCR and voting disk are not configured under VCS, manually mount the OCR and voting disk after you finish the steps in the following procedure.

- Starts the CVM group on the new node.

To prepare to install Oracle RAC on the new node using the SF Oracle RAC installer

- 1 After you configure SF Oracle RAC on the new node, the installer displays the following options for configuring Oracle RAC:

- 1) Create Oracle User and Group
- 2) Configure private IP addresses (PrivNIC configuration)
- 3) Configure private IP addresses (MultiPrivNIC configuration)
- 4) Finish

Note: Depending on whether you configured the private IP addresses as PrivNIC or MultiPrivNIC resources in the existing cluster, either option 2 or 3 are displayed.

Enter **1** to select the option **Create Oracle User and Group** from the SF Oracle RAC installer menu.

See [“Creating Oracle user and groups on the new node”](#) on page 470.

- 2 Configure the private IP addresses and the PrivNIC resource for Oracle Clusterware (only if the IP addresses on the existing cluster are configured as PrivNIC resources).

You must manually update the PrivNIC resource configuration in the following cases:

- If the PrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the PrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {eth1=0, eth2=1}
Device@nebula = {eth1=0, eth2=1}
```

Enter **2** to select the option **Configure private IP addresses (PrivNIC configuration)**.

See [“Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware”](#) on page 471.

- 3 Configure private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC (only if the IP addresses on the existing cluster are configured as MultiPrivNIC resources).

You must manually update the MultiPrivNIC resource configuration in the following cases:

- If the MultiPrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the MultiPrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {eth1=0, eth2=1, eth3=2}
Device@nebula = {eth1=0, eth2=1, eth3=2}
```

Enter **3** to select the option **Configure private IP addresses (MultiPrivNIC configuration)**.

See [“Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC”](#) on page 473.

- 4 Select **Finish** to start the `cvm` group on the new node.

Note: The `cssd` resource appears FAULTED until the new node is added to Oracle Clusterware.

- 5 If the `cssd` resource is configured as a critical resource, the `cvm` group will be brought offline on the new node. Modify the configuration to make the `cssd` resource non-critical and bring the `cvm` group online.

- On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Bring the `cvm` group online:

```
# hares -online cvm -sys saturn
```

6 Verify that all the GAB ports are up:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen ada401 membership 012
Port b gen ada40d membership 012
Port d gen ada409 membership 012
Port f gen ada41c membership 012
Port h gen ada40f membership 012
Port o gen ada406 membership 012
Port u gen ada416 membership 012
Port v gen ada416 membership 012
Port w gen ada418 membership 012
Port y gen ada419 membership 012
```

7 Complete the following additional preparatory tasks using the instructions in the chapter "Preparing to install Oracle RAC":

- Identify public virtual IP addresses for use by Oracle RAC.
- Set the kernel parameters.
- Verify that the user "nobody" exists.
- Set up Oracle user equivalence for remote shell and remote copy environments.
- Edit the Oracle user profile.
- If the OCR and voting disk resources are not configured under VCS, mount the OCR and voting disk manually.

8 Create Oracle Clusterware and Oracle RAC database home directories manually.

Creating Oracle user and groups on the new node

Perform the steps in the following procedure to create Oracle user and groups on the new node.

Note: Set the password of the Oracle user manually before you configure secure shell or remote shell connection on the node.

To create Oracle user and groups on the new node

- 1 Enter the Oracle user name that is used for Oracle RAC operations on the existing cluster.

Note: If the Oracle user and groups already exist on the new node, make sure that the UID and GID of the Oracle user and groups are the same as that on the current cluster.

```
Enter Oracle UNIX user name:  [b] Oracle
```

The installer obtains the existing group and identifier information based on the Oracle user name.

- 2 Review and confirm the information.

The installer adds the user and groups to the new node.

- 3 Press **Return** to continue with the other configuration tasks.

Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware

Perform this step only if the private IP addresses are configured as PrivNIC resources in the existing cluster.

Note: Make sure that the network interface names of the private interconnects on the new node are the same as those of the existing cluster. For maximum failover options, all available LLT links are used for PrivNIC configuration.

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

To configure the private IP addresses and PrivNIC resource for Oracle Clusterware

- 1 Enter the name for the PrivNIC resource. The installer displays the names of the existing PrivNIC resources. Specify the name of an existing resource.

```
Enter the PrivNIC resource name: [b] (ora_priv)
```

- 2 Enter **y** if you want the installer to add the IP address to the `/etc/hosts` file.
 You can also add the IP address to the file manually after the configuration.
 The installer displays the existing PrivNIC resource configuration on the cluster.

```
Resource name: ora_priv
System: galaxy
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.1
      Alias for above IP: galaxy-priv
System: nebula
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.2
      Alias for above IP: nebula-priv
Is this information correct? [y,n,q] (y)
```

- 3 Review and confirm the information.
- 4 Enter the private IP address and its private node name for the new node.

```
Enter the private IP for saturn: [b] 192.168.12.5
Enter Hostname alias for the above IP address: [b] saturn-priv
```

The installer displays the resource configuration for the new node.

```
Resource name: ora_priv
System: saturn
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.5
      Alias for above IP: saturn-priv
Is this information correct? [y,n,q] (y)
```

- 5 Review and confirm the information.
 The installer updates the existing PrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).
- 6 Press **Return** to continue with the other configuration tasks.

Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

Perform this step only if the private IP addresses are configured as MultiPrivNIC resources in the existing cluster.

Note: Make sure that you configure the same interfaces (as those on the existinnebulag cluster) for private interconnects on the new node. For maximum failover options, all available LLT links are used for MultiPrivNIC configuration.

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

To configure the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

- 1 Enter the name for the MultiPrivNIC resource. The installer displays the names of the existing MultiPrivNIC resources. Specify the name of an existing resource.

```
Enter the MultiPrivNIC resource name: [b] (multi_priv)
```

- 2 Enter **y** if you want the installer to add the IP address to the `/etc/hosts` file.

You can also add the IP address to the file manually after the configuration.

The installer displays the existing MultiPrivNIC resource configuration on the cluster.

```
Resource name: multi_priv
System: galaxy
      Private Interfaces: eth1 eth2
      Private IPs on eth1: 192.168.12.1
      Aliases for above IPs: galaxy-priv
      Private IPs on eth2: 192.168.2.1
      Aliases for above IPs: galaxy-priv1
System: nebula
      Private Interfaces: eth1 eth2
      Private IPs on eth1: 192.168.12.2
      Aliases for above IPs: nebula-priv
      Private IPs on eth2: 192.168.2.2
      Aliases for above IPs: nebula-priv1
Is this information correct? [y,n,q] (y)
```

- 3 Review and confirm the information.

- 4 Enter the private IP address and the corresponding private node name for the eth1 interface on the new node.

```
Enter IP addresses for saturn for eth1
separated by space: [b,q,?] 192.168.12.5
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] saturn-priv
```

- 5 Enter the private IP address and the corresponding private node name for the eth2 interface on the new node.

```
Enter IP addresses for saturn for eth2
separated by space: [b,q,?] 192.168.2.6
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] saturn-priv1
```

The installer displays the resource configuration for the new node.

```
Resource name: multi_priv
System: saturn
    Private Interfaces: eth1 eth2
    Private IPs on eth1: 192.168.12.5
    Aliases for above IPs: saturn-priv
    Private IPs on eth2: 192.168.2.6
    Aliases for above IPs: saturn-priv1
Is this information correct? [y,n,q] (y)
```

- 6 Review and confirm the information.

The installer updates the existing MultiPrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).

- 7 Press **Return** to continue with the other configuration tasks.

Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC Web-based installer

The SF Oracle RAC installer performs the following tasks:

- Creates the Oracle user and groups on the new node
- Configures the private IP addresses and the PrivNIC or MultiPrivNIC resources (if they are configured in the existing cluster).

- If the CFSMount and CVMVolDg resources for OCR and voting disk are configured under the `cvm` service group, the installer brings them online after adding the node to the cluster.
 If the resources are configured in any other service group, make sure that you modify the service group to include the new node and bring the service group online.

Note: If OCR and voting disk are not configured under VCS, manually mount the OCR and voting disk after you finish the steps in the following procedure.

- Starts the CVM group on the new node.

To prepare to install Oracle RAC on the new node using the SF Oracle RAC installer

- 1 After you configure SF Oracle RAC on the new node, the installer displays the following options for configuring Oracle RAC:

```
Create Oracle User and Group
Configure private IP addresses (PrivNIC configuration)
Configure private IP addresses (MultiPrivNIC configuration)
Finish
```

Note: Depending on whether you configured the private IP addresses as PrivNIC or MultiPrivNIC resources in the existing cluster, either option 2 or 3 is displayed.

Select the option **Create Oracle User and Group** from the SF Oracle RAC installer menu. Click **Next**.

See [“Creating Oracle user and groups on the new node”](#) on page 470.

- 2 Configure the private IP addresses and the PrivNIC resource for Oracle Clusterware (only if the IP addresses on the existing cluster are configured as PrivNIC resources).

You must manually update the PrivNIC resource configuration in the following cases:

- If the PrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the PrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {eth1=0, eth2=1}
Device@nebula = {eth1=0, eth2=1}
```

Select the option **Configure private IP addresses (PrivNIC configuration)**. Click **Next**.

See [“Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware”](#) on page 471.

- 3 Configure private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC (only if the IP addresses on the existing cluster are configured as MultiPrivNIC resources).

You must manually update the MultiPrivNIC resource configuration in the following cases:

- If the MultiPrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the MultiPrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {eth1=0, eth2=1, eth3=2}
Device@nebula = {eth1=0, eth2=1, eth3=2}
```

Select the option **Configure private IP addresses (MultiPrivNIC configuration)**. Click **Next**.

See [“Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC”](#) on page 473.

- 4 Select **Finish** to start the `cvm` group on the new node.

Note: The `cssd` resource appears **FAULTED** until the new node is added to Oracle Clusterware.

- 5 If the `cssd` resource is configured as a critical resource, the `cvm` group will be brought offline on the new node. Modify the configuration to make the `cssd` resource non-critical and bring the `cvm` group online.

- On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Bring the `cvm` group online:

```
# hares -online cvm -sys saturn
```

6 Verify that all the GAB ports are up:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen ada401 membership 012
Port b gen ada40d membership 012
Port d gen ada409 membership 012
Port f gen ada41c membership 012
Port h gen ada40f membership 012
Port o gen ada406 membership 012
Port u gen ada416 membership 012
Port v gen ada416 membership 012
Port w gen ada418 membership 012
Port y gen ada419 membership 012
```

7 Complete the following additional preparatory tasks using the instructions in the chapter "Preparing to install Oracle RAC":

- Identify public virtual IP addresses for use by Oracle RAC.
- Set the kernel parameters.
- Verify that the user "nobody" exists.
- Set up Oracle user equivalence for remote shell and remote copy environments.
- Edit the Oracle user profile.
- If the OCR and voting disk resources are not configured under VCS, mount the OCR and voting disk manually.

8 Create Oracle Clusterware and Oracle RAC database home directories manually.

Creating Oracle user and groups on the new node

Perform the steps in the following procedure to create Oracle user and groups on the new node.

Note: Set the password of the Oracle user manually before you configure secure shell or remote shell connection on the node.

To create Oracle user and groups on the new node

- 1 Enter the Oracle user name that is used for Oracle RAC operations on the existing cluster.

Note: If the Oracle user and groups already exist on the new node, make sure that the UID and GID of the Oracle user and groups are the same as that on the current cluster.

Enter Oracle UNIX user name:

Click **Next**. The installer obtains the existing group and identifier information based on the Oracle user name.

- 2 Review and confirm the information.
The installer adds the user and groups to the new node.
- 3 Enter the information for the secondary group, if required.
- 4 Click **Yes** to continue with the other configuration tasks.

Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware

Perform this step only if the private IP addresses are configured as PrivNIC resources in the existing cluster.

Note: Make sure that the network interface names of the private interconnects on the new node are the same as those of the existing cluster. For maximum failover options, all available LLT links are used for PrivNIC configuration.

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

To configure the private IP addresses and PrivNIC resource for Oracle Clusterware

- 1 Enter the name for the PrivNIC resource. The installer displays the names of the existing PrivNIC resources. Specify the name of an existing resource.

Enter the PrivNIC resource name:

Click **Next**.

- 2 Click **Yes** if you want the installer to add the IP address to the `/etc/hosts` file.
 You can also add the IP address to the file manually after the configuration.
 The installer displays the existing PrivNIC resource configuration on the cluster.

```
Resource name: ora_priv
System: galaxy
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.1
      Alias for above IP: galaxy-priv
System: nebula
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.2
      Alias for above IP: nebula-priv
Is this information correct?
```

- 3 Review and confirm the information.
- 4 Enter the private IP address and its private node name for the new node.

```
Enter the private IP for saturn: 192.168.12.5
Enter Hostname alias for the above IP address: saturn-priv
```

Click **Next**.

The installer displays the resource configuration for the new node.

```
Resource name: ora_priv
System: saturn
      Private Interfaces: eth1 eth2
      Private IP address: 192.168.12.5
      Alias for above IP: saturn-priv
Is this information correct?
```

5 Review and confirm the information.

The installer updates the existing PrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).

6 Click **Yes** to continue with the other configuration tasks.

Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

Perform this step only if the private IP addresses are configured as MultiPrivNIC resources in the existing cluster.

Note: Make sure that you configure the same interfaces (as those on the existing cluster) for private interconnects on the new node. For maximum failover options, all available LLT links are used for MultiPrivNIC configuration.

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

To configure the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

- 1 Enter the name for the MultiPrivNIC resource. The installer displays the names of the existing MultiPrivNIC resources. Specify the name of an existing resource.

Enter the MultiPrivNIC resource name:

Click **Yes**.

- 2 Click **Yes** if you want the installer to add the IP address to the `/etc/hosts` file.
 You can also add the IP address to the file manually after the configuration.
 The installer displays the existing MultiPrivNIC resource configuration on the cluster.

```
Resource name: multi_priv
System: galaxy
      Private Interfaces: eth1 eth2
      Private IPs on eth1: 192.168.12.1
      Aliases for above IPs: galaxy-priv
      Private IPs on eth2: 192.168.2.1
      Aliases for above IPs: galaxy-priv1
System: nebula
      Private Interfaces: eth1 eth2
      Private IPs on eth1: 192.168.12.2
      Aliases for above IPs: nebula-priv
      Private IPs on eth2: 192.168.2.2
      Aliases for above IPs: nebula-priv1
Is this information correct?
```

- 3 Review and confirm the information.
- 4 Enter the private IP address and the corresponding private node name for the eth1 interface on the new node.

```
Enter IP addresses for saturn for eth1
separated by space: 192.168.12.5
Enter Hostname aliases for the above IP addresses
separated by space: saturn-priv
```

Click **Next**.

- 5 Enter the private IP address and the corresponding private node name for the eth2 interface on the new node.

```
Enter IP addresses for saturn for eth2
separated by space: 192.168.2.6
Enter Hostname aliases for the above IP addresses
separated by space: saturn-priv1
```

Click **Next**.

The installer displays the resource configuration for the new node.

```
Resource name: multi_priv
System: saturn
      Private Interfaces: eth1 eth2
      Private IPs on eth1: 192.168.12.5
      Aliases for above IPs: saturn-priv
      Private IPs on eth2: 192.168.2.6
      Aliases for above IPs: saturn-priv1
Is this information correct?
```

- 6 Review and confirm the information.

The installer updates the existing MultiPrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).

- 7 Click **Yes** to continue with the other configuration tasks.

Preparing the new node manually for installing Oracle RAC

Complete the following preparatory tasks manually before you install Oracle RAC on the new node.

To prepare to install Oracle RAC on the new node

- 1 Create Oracle user and groups.
- 2 Configure private IP addresses and the PrivNIC resource for Oracle Clusterware.

See [“Configuring private IP address and PrivNIC resource for Oracle Clusterware”](#) on page 483.

- 3 Configure private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC.
 See [“Configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC”](#) on page 484.
- 4 Start VCS on the new node.
 See [“Starting VCS on the new node”](#) on page 485.
- 5 Create the Oracle Clusterware/Grid Infrastructure and Oracle RAC database home directories for installation.
 See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the new node ”](#) on page 486.
- 6 Complete the following additional preparatory tasks using the instructions in the chapter "Preparing to install Oracle RAC":
 - Identify public virtual IP addresses for use by Oracle RAC.
 - Set the kernel parameters.
 - Verify that the user "nobody" exists.
 - Set up Oracle user equivalence for remote shell and remote copy environments.
 - Edit the Oracle user profile.

Configuring private IP address and PrivNIC resource for Oracle Clusterware

This section provides instructions for configuring private IP address and PrivNIC resource for Oracle Clusterware.

Identify a private IP address that you want to use for the new node. Make sure that the IP address is in the same subnet as the existing cluster.

The procedure uses the following IP address for the new node:

On saturn: 192.168.12.5

To configure private IP addresses for Oracle Clusterware

- 1 Make a backup copy of the main.cf file. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

- 2 Add the private IP address to the the `ora_priv` resource on the active node:

```
# haconf -makerw

# hares -modify priv_resname Device -add nic1_node1 0 \
-sys nodenew_name

# hares -modify priv_resname Device -add nic2_node1 1 \
-sys nodenew_name

# hares -modify priv_resname Address "privnic_ip_newnode" \
-sys nodenew_name

# haconf -dump -makero
```

Configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC

This section provides instructions for configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC.

Identify the private IP addresses that you want to use for the new node. Make sure that the IP addresses are in the same subnet as the existing cluster.

The procedure uses the following IP addresses for the new node:

IP addresses for the new node	192.168.12.5
	192.168.2.6

To configure private IP addresses for Oracle Clusterware

- ◆ Use the following commands to add private IP into multi_priv resource on the active node:

```
# haconf -makerw

# hares -modify multipriv_resname Device -add eth1 0 \
-sys nodenew_name

# hares -modify multipriv_resname Device -add eth2 1 \
-sys nodenew_name

# hares -modify multipriv_resname Address -add \
multipriv_ip_newnode 0 -sys nodenew_name

# hares -modify multipriv_resname Address -add \
multipriv_udpip1_newnode 1 -sys nodenew_name

# haconf -dump -makero
```

Starting VCS on the new node

Before you start VCS, create the file `cssd-pretend-offline` on the new node and make the `cssd` resource non-critical. Failing this, the `cssd` resource lapses into an UNKNOWN state until Oracle Clusterware is installed on the new node, thus preventing the `cvm` group from coming online.

Note: The `cssd` resource will remain in FAULTED/OFFLINE state till Oracle Clusterware is installed on the new node.

To start VCS on the new node

- 1 On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- 2 Create the file `cssd-pretend-offline` on the new node:

```
# touch /var/VRTSvcs/lock/cssd-pretend-offline
```

- 3 Start VCS on the new node:

```
# hastart
```

Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the new node

The Oracle Clusterware/Grid Infrastructure and Oracle database home directories must be located on the same storage as that on the existing nodes.

Note: In the case of Oracle RAC 11g Release 1, only the database is supported. References to Oracle RAC 11g Release 1 in the procedure applies to the Oracle database alone.

Depending on the storage in the existing cluster, use one of the following options to create the directories:

Local file system	See “To create the directories on the local file system” on page 486.
Cluster File System	See “To create the file system and directories on cluster file system for Oracle Clusterware and Oracle database” on page 488.

To create the directories on the local file system

- 1 Log in as the root user on the node.
- 2 Create a local file system and mount it using one of the following methods:
 - Using native operating system commands
For instructions, see the operating system documentation.
 - Using Veritas File System (VxFS) commands

As the root user, create a VxVM local diskgroup on each node.	<pre># vxdg init vxvm_dg \ dg_name</pre>
Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure binaries and Oracle binaries.	<pre># vxassist -g vxvm_dg make clus_volname size # vxassist -g vxvm_dg make ora_volname size</pre>
Create the file systems with the volumes.	<pre># mkfs -t vxfs /dev/vx/rdisk/vxvm_dg/clus_volname # mkfs -t vxfs /dev/vx/rdisk/vxvm_dg/ora_volname</pre>
Mount the file system.	<pre># mount -t vxfs /dev/vx/dsk/vxvm_dg/clus_volname \ clus_home # mount -t vxfs /dev/vx/dsk/vxvm_dg/ora_volname \ oracle_home</pre>

3 Create the directories for Oracle RAC.

```
# mkdir -p grid_base
# mkdir -p clus_home
# mkdir -p oracle_base
# mkdir -p oracle_home
```

4 Set appropriate ownership and permissions for the directories.

```
# chown -R grid:oinstall grid_base
# chmod -R 775 grid_base
# chown -R grid:oinstall clus_home
# chmod -R 775 clus_home
# chown -R oracle:oinstall oracle_base
# chmod -R 775 oracle_base
# chown -R oracle:oinstall oracle_home
# chmod -R 775 oracle_home
```

5 Add the resources to the VCS configuration.

See [“To add the storage resources created on VxFS to the VCS configuration”](#) on page 488.

6 Repeat all the steps on each node of the cluster.

To add the storage resources created on VxFS to the VCS configuration

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the VxVM volumes under VCS:

```
# hares -add dg_resname DiskGroup cvm
# hares -modify dg_resname DiskGroup vxvm_dg -sys nodenew_name
# hares -modify dg_resname Enabled 1
```

- 3 Set up the file system under VCS:

```
# hares -add clusbin_mnt_resname Mount cvm

# hares -modify clusbin_mnt_resname MountPoint \
"clus_home"

# hares -modify clusbin_mnt_resname BlockDevice \
"/dev/vx/dsk/vxvm_dg/clus_volname" -sys nodenew_name
# hares -modify clusbin_mnt_resname FSType vxfs
# hares -modify clusbin_mnt_resname FsckOpt "-n"
# hares -modify clusbin_mnt_resname Enabled 1
# hares -add orabin_mnt_resname Mount cvm

# hares -modify orabin_mnt_resname MountPoint \
"oracle_home"

# hares -modify orabin_mnt_resname BlockDevice \
"/dev/vx/dsk/vxvm_dg/ora_volname" -sys nodenew_name
# hares -modify orabin_mnt_resname FSType vxfs
# hares -modify orabin_mnt_resname FsckOpt "-n"
# hares -modify orabin_mnt_resname Enabled 1
```

- 4 Link the parent and child resources:

```
# hares -link clusbin_mnt_resname vxvm_dg
# hares -link orabin_mnt_resname vxvm_dg
```

- 5 Repeat all the steps on each node of the cluster.

To create the file system and directories on cluster file system for Oracle Clusterware and Oracle database

Perform the following steps on the CVM master node in the cluster.

- 1 As the root user, create a VxVM shared disk group:

```
# vxvg -s init cvm_dg dg_name
```

- 2 Create separate volumes for Oracle Clusterware and Oracle database:

```
# vxassist -g cvm_dg make clus_volname size
# vxassist -g cvm_dg make ora_volname size
```

- 3 Create the Oracle base directory, clusterware home directory, and the Oracle home directory.

```
# mkdir -p oracle_base
# mkdir -p oracle_home
# mkdir -p clus_home
# mkdir -p grid_base
```

- 4 Create file systems with the volumes:

```
# mkfs -t vxfs /dev/vx/rdisk/cvm_dg/ora_volname
```

- 5 Mount the file systems. Perform this step on each node.

```
# mount -t vxfs -o cluster /dev/vx/dsk/cvm_dg/ora_volname \
oracle_home
```

- 6 Change the ownership and permissions on all nodes of the cluster.

```
# chown -R grid:oinstall grid_base
# chmod -R 775 grid_base
# chown -R grid:oinstall clus_home
# chmod -R 775 clus_home
# chown -R oracle:oinstall oracle_base
# chmod -R 775 oracle_base
# chown -R oracle:oinstall oracle_home
# chmod -R 775 oracle_home
```

- 7 Add the CVMVolDg and CFSSMount resources to the VCS configuration.

See [“To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI”](#) on page 317.

Adding the new node to Oracle RAC

Install Oracle Clusterware and Oracle RAC database on the node using the Oracle RAC add node procedure.

For instructions, see the Oracle RAC documentation.

After installing Oracle Clusterware and Oracle RAC database, perform the following post-installation tasks:

1. Delete the file `/var/VRTSvcs/lock/cssd-pretend-offline` to bring the CSSD resource online.
2. If CSSD is not configured under the cvm group, add the new node information to the service group containing the CSSD resource.
3. Add new Oracle RAC database instances for the new node.

For instructions, see the Oracle RAC documentation.

4. Update the Oracle RAC database service groups to include the new database instances in the VCS configuration file.
5. For other service groups that are configured under VCS, manually update the service group configuration for the new node.

Adding nodes to a cluster that is using authentication for SFDB tools

To add a node to a cluster that is using authentication for SFDB tools, perform the following steps as the root user

1. Export authentication data from a node in the cluster that has already been authorized, by using the `-o export_broker_config` option of the `sfae_auth_op` command.

Use the `-f` option to provide a file name in which the exported data is to be stored.

```
# /opt/VRTS/bin/sfae_auth_op \  
-o export_broker_config -f exported-data
```

2. Copy the exported file to the new node by using any available copy mechanism such as `scp` or `rcp`.

- 3 Import the authentication data on the new node by using the `-o import_broker_config` option of the `sfae_auth_op` command.

Use the `-f` option to provide the name of the file copied in Step 2.

```
# /opt/VRTS/bin/sfae_auth_op \
-o import_broker_config -f exported-data
Setting up AT
Importing broker configuration
Starting SFAE AT broker
```

- 4 Stop the `vxdbd` daemon on the new node.

```
# /opt/VRTS/bin/vxdbdctrl stop
Stopping Veritas vxdbd
vxdbd stop succeeded
```

- 5 Enable authentication by setting the `AUTHENTICATION` key to `yes` in the `/etc/vx/vxdbed/admin.properties` configuration file.

- 6 Start the `vxdbd` daemon.

```
# /opt/VRTS/bin/vxdbdctrl start
Starting Veritas vxdbd
/opt/VRTSdbed/bin/vxdbd start SUCCESS
```

The new node is now authenticated to interact with the cluster to run SFDB commands.

Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Storage Checkpoints, Database FlashSnap, or SmartTier for Oracle in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

To update the SFDB repository after adding a node

- 1 Copy the `/var/vx/vxdba/rep_loc` file from one of the nodes in the cluster to the new node.
- 2 If the `/var/vx/vxdba/auth/user-authorizations` file exists on the existing cluster nodes, copy it to the new node.

If the `/var/vx/vxdba/auth/user-authorizations` file does not exist on any of the existing cluster nodes, no action is required.

This completes the addition of the new node to the SFDB repository.

For information on using SFDB tools features, refer to *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*.

Sample configuration file for adding a node to the cluster

You may use this sample file as reference information to understand the configuration changes that take place when you add a node to a cluster.

The existing sample configuration before adding the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises two nodes `galaxy` and `nebula` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

The following sample configuration file shows the changes (in **bold**) effected in the configuration after adding a node "saturn" to the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
```

```
cluster rac_cluster101 (
```

```

        UserNames = { admin = bopHo }
        Administrators = { admin }
        UseFence = SCSI3
    )

system galaxy (
)
system nebula (
)
system saturn (
)

```

Note: In the following group oradb1_grp, the saturn node has been added.

```

group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

```

Note: In the following Oracle resource, the saturn node information has been added.

```

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (

```

```

Critical = 0
CVMDiskGroup = oradatadg
CVMVolume = { oradatavol }
CVMActivation = sw
)

```

```

requires group cvm online local firm
oral requires oradata_mnt
oradata_mnt requires oradata_voldg

```

Note: In the following CVM and CVMCluster resources, the saturn node information has been added.

```

group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

```

```

    )

CFSfsckd vxfsckd (
    )

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
    )

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

```

Note: In the following PrivNIC resource, the saturn node information has been added.

```

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { eth1 = 0, eth2 = 1}
    Device@nebula = { eth1 = 0, eth2 = 1}
    Device@saturn = { eth1 = 0, eth2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    Address@saturn = "192.168.12.5"
    NetMask = "255.255.255.0"
    )

cssd requires ocrvote_mnt
cssd requires ora_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

```

Removing a node from SF Oracle RAC clusters

This chapter includes the following topics:

- [About removing a node from a cluster](#)
- [Removing a node from a cluster](#)
- [Modifying the VCS configuration files on existing nodes](#)
- [Removing the node configuration from the CP server](#)
- [Removing security credentials from the leaving node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after removing a node](#)
- [Sample configuration file for removing a node from the cluster](#)

About removing a node from a cluster

You can remove one or more nodes from an SF Oracle RAC cluster.

Overview of tasks for removing a node from a cluster:

- Prepare the node to be removed.
 - Take offline service groups and resources which support Oracle.
 - Stop applications that use File System or Cluster File System mount points not configured under VCS.
 - Remove database instances and software from the node.

See [“Removing a node from a cluster”](#) on page 497.

- Remove the node from the cluster.
 - Stop VCS on the node to be removed.
 - Remove Oracle Clusterware from the node.
 - Unmount the File System and Cluster File System file systems not configured under VCS.
 - Uninstall SF Oracle RAC from the node.
 See [“Removing a node from a cluster”](#) on page 497.
- Modify the VCS configuration files on the existing nodes.
 See [“Modifying the VCS configuration files on existing nodes”](#) on page 498.
- Remove the node configuration from the CP server if it is configured.
 See [“Removing the node configuration from the CP server”](#) on page 501.
- Remove the security credentials from the node if it is part of a secure cluster.
 See [“Removing security credentials from the leaving node ”](#) on page 502.
- Update the SFDB repository if you use SFDB tools.
 See [“Updating the Storage Foundation for Databases \(SFDB\) repository after removing a node”](#) on page 502.

The Veritas product installer does not support removing a node. You must remove a node manually. The example procedures describe how to remove a node from a cluster with three nodes.

Removing a node from a cluster

Perform the following steps to remove a node from a cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

To prepare to remove a node from a cluster

- 1 Take the Oracle RAC service groups offline (if under VCS control) on the node you want to remove.

```
# hagrps -offline oracle_group -sys saturn
```

- 2 Stop the applications that use VxFS/CFS mount points and are not configured under VCS. Use native application commands to stop the applications.
- 3 Remove the Oracle RAC database software from the node.

For instructions, see the Oracle RAC documentation.

To remove a node from a cluster

- 1 Stop VCS on the node:

```
# hastop -local
```

- 2 Remove Oracle Clusterware from the node.

For instructions, see the Oracle RAC document.

- 3 Unmount the VxFS/CFS file systems that are not configured under VCS.

```
# umount mount_point
```

- 4 Uninstall SF Oracle RAC from the node using the SF Oracle RAC installer.

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

```
# ./uninstallsfrac saturn
```

The installer stops all SF Oracle RAC processes and uninstalls the SF Oracle RAC packages.

- 5 Modify the VCS configuration files on the existing nodes to remove references to the deleted node.

See [“Modifying the VCS configuration files on existing nodes”](#) on page 498.

Modifying the VCS configuration files on existing nodes

Modify the configuration files on the remaining nodes of the cluster to remove references to the deleted nodes.

The process involves:

- Editing the /etc/llthosts file
- Editing the /etc/gabtab file
- Modifying the VCS configuration to remove the node

For an example main.cf:

To edit the `/etc/llthosts` file

- ◆ On each of the existing nodes, edit the `/etc/llthosts` file to remove lines that contain references to the removed nodes.

For example, if `saturn` is the node removed from the cluster, remove the line "2 saturn" from the file:

```
0 galaxy
1 nebula
2 saturn
```

Change to:

```
0 galaxy
1 nebula
```

To edit the `/etc/gabtab` file

- ◆ Modify the following command in the `/etc/gabtab` file to reflect the number of systems after the node is removed:

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

where `N` is the number of remaining nodes in the cluster.

For example, with two nodes remaining, the file resembles:

```
/sbin/gabconfig -c -n2
```

Modify the VCS configuration file `main.cf` to remove all references to the deleted node.

Use one of the following methods to modify the configuration:

- Edit the `/etc/VRTSvcs/conf/config/main.cf` file
This method requires application down time.
- Use the command line interface
This method allows the applications to remain online on all remaining nodes.

The following procedure uses the command line interface and modifies the sample VCS configuration to remove references to the deleted node. Run the steps in the procedure from one of the existing nodes in the cluster. The procedure allows you to change the VCS configuration while applications remain online on the remaining nodes.

To modify the VCS configuration using the CLI

- 1 Back up the `/etc/VRTSvcs/conf/config/main.cf` file.

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.3node.bak
```

- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Remove the node from the `AutoStartList` attribute of the service group by specifying the remaining nodes in the desired order:

```
# hagrps -modify cvm AutoStartList galaxy nebula
```

- 4 Remove the node from the `SystemList` attribute of the service group:

```
# hagrps -modify cvm SystemList -delete saturn
```

- 5 Remove the node from the `CVMNodeId` attribute of the service group:

```
# hares -modify cvm_clus CVMNodeId -delete saturn
```

- 6 If you have the other service groups (such as the database service group or the `ClusterService` group) that have the removed node in their configuration, perform step 4 and step 5 for each of them.

- 7 Remove the deleted node from the `NodeList` attribute of all CFS mount resources:

```
# hares -modify CFSMount NodeList -delete saturn
```

- 8 Remove the deleted node from the system list of any other service groups that exist on the cluster. For example, to delete the node saturn:

```
# hagrps -modify crsgrp SystemList -delete saturn
```

- 9 Remove the deleted node from the cluster system list:

```
# hasys -delete saturn
```

- 10 Save the new configuration to disk:

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

- 11 Verify that the node is removed from the VCS configuration.

```
# grep -i saturn /etc/VRTSvcs/conf/config/main.cf
```

If the node is not removed, use the VCS commands as described in this procedure to remove the node.

Removing the node configuration from the CP server

After removing a node from a SF Oracle RAC cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

Note: The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where `cp_server` is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
# cpsadm -s cp_server -a rm_node -h saturn -c rac_cluster101 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
# cpsadm -s cp_server -a list_nodes
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node `saturn`. Perform the following steps.

To remove the security credentials

- 1 Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \
stop
```

- 2 Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Updating the Storage Foundation for Databases (SFDB) repository after removing a node

After removing a node from a cluster, you do not need to perform any steps to update the SFDB repository.

For information on updating the SFDB repository after adding a node to the cluster:

See [“Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node”](#) on page 491.

For information on removing the SFDB repository after removing the product:

See [“Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product”](#) on page 579.

Sample configuration file for removing a node from the cluster

You may use this sample file as reference information to understand the configuration changes involved when you remove a node from a cluster.

The existing sample configuration before removing the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises three nodes `galaxy`, `nebula`, and `saturn` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.

- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Note: The following sample file shows in **bold** the configuration information that is removed when the node "saturn" is removed from the cluster.

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

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

system galaxy (
)
system nebula (
)
system saturn (
)
```

Note: In the following group oradb1_grp, the saturn node must be removed.

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```

Note: In the following Oracle resource, the saturn node information must be removed.

```

Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
oral requires oradata_mnt
oradata_mnt requires oradata_voldg

```

Note: In the following CVM and CVMCluster resources, the saturn node information must be removed.

```

group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
)

```



```

CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
OnlineRetryLimit = 20
)

CFSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

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

```

Note: In the following PrivNIC resource, the saturn node information must be removed.

```
PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { eth1 = 0, eth2 = 1}
    Device@nebula = { eth1 = 0, eth2 = 1}
    Device@saturn = { eth1 = 0, eth2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    Address@saturn = "192.168.12.5"
    NetMask = "255.255.255.0"
)
```

```
cssd requires ocrvote_mnt
cssd requires ora_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```

Configuration of disaster recovery environments

- [Chapter 30. Configuring a campus cluster setup for disaster recovery](#)
- [Chapter 31. Setting up a replicated global cluster](#)
- [Chapter 32. Configuring a global cluster using VVR](#)

Configuring a campus cluster setup for disaster recovery

This chapter includes the following topics:

- [Setting up an SF Oracle RAC campus cluster for disaster recovery](#)
- [Tuning guidelines for campus clusters](#)
- [Best practices for a campus cluster](#)

Setting up an SF Oracle RAC campus cluster for disaster recovery

Perform the following tasks to setup an SF Oracle RAC campus cluster:

- Verify you have met the hardware requirements for SF Oracle RAC:
See [“Hardware requirements”](#) on page 42.
- Verify you have met the license requirements:
 - SF Oracle RAC with HA/DR license
 - FlashSnap license
 - Site awareness licenseWith keyless licensing, SF Oracle RAC enterprise keys enable all of the above features.
- Prepare to set up campus cluster configuration
See [“Preparing to set up a campus cluster”](#) on page 509.

- Configure I/O fencing to prevent data corruption
- Prepare to install Oracle RAC Clusterware and database binaries
- Configure VxVM disk groups for campus cluster
 See [“Configuring VxVM disk groups for Oracle RAC on a campus cluster”](#) on page 510.
- Install Oracle RAC Clusterware and database binaries
- Configure VCS service groups

The sample configuration used to illustrate the configuration procedures includes a four-node SF Oracle RAC campus cluster with two nodes at each site. Each node has SF Oracle RAC 6.0 on Linux.

Table 30-1 Sample setup for campus cluster

Site	Hardware description
Site 1: SFRAC1	Servers: galaxy and nebula Shared LUNs: disk01 disk02 disk03 disk04 (used as coordinator disk) disk5
Site 2: SFRAC2	Servers: mercury and jupiter Shared LUNs: disk06 disk07 disk08 disk09 (used as coordinator disk)
Site 3	Shared LUN disk10 (used as coordinator disk)

Preparing to set up a campus cluster

The following preparation must be completed before configuring the campus cluster.

To prepare to set up a campus cluster

- 1 Configure the physical infrastructure for campus cluster:

- Set up access to the local storage arrays and to remote storage arrays on each node. The storage link will extend to the third site as well.
 - Set up the private heartbeat network
 See the *Veritas Cluster Server Administrator's Guide*
- 2 Install the operating system on all the nodes of the cluster.
 See your operating system documentation.
 - 3 Install and configure SF Oracle RAC on all nodes on both the sites.
 See [“About configuring SF Oracle RAC”](#) on page 114.
 In the sample setup, install and configure SF Oracle RAC 6.0 on all four nodes.

Configuring VxVM disk groups for Oracle RAC on a campus cluster

After configuring I/O fencing for data integrity, you must configure the VxVM disk groups for a campus cluster before installing Oracle RAC by configuring VxVM disk groups for remote mirroring.

To configure VxVM disk groups for Oracle RAC on a campus cluster

- 1 Initialize the disks as CDS disks

```
# vxdisksetup -i disk01 format=cddisk
# vxdisksetup -i disk02 format=cddisk
# vxdisksetup -i disk03 format=cddisk
# vxdisksetup -i disk05 format=cddisk
# vxdisksetup -i disk06 format=cddisk
# vxdisksetup -i disk07 format=cddisk
# vxdisksetup -i disk08 format=cddisk
```

- 2 Set the site name for each host:

```
# vxctl set site=sitename
```

The site name is stored in the /etc/vx/volboot file. To display the site names:

```
# vxctl list | grep siteid
```

For example, for a four node cluster with two nodes at each site, mark the sites as follows:

On the nodes at first site:

```
# vxctl set site= SFRAC1
```

On the nodes at second site:

```
# vxdctl set site= SFRAC2
```

- 3 Set the site name for all the disks in an enclosure.

```
# vxdisk settag site=sitename encl:enclosure
```

- 4 Run the following command if you want to tag only the specific disks:

```
# vxdisk settag site=sitename disk
```

For example:

```
# vxdisk settag site=SFRAC1 disk01
# vxdisk settag site=SFRAC1 disk02
# vxdisk settag site=SFRAC1 disk03
# vxdisk settag site=SFRAC2 disk06
# vxdisk settag site=SFRAC2 disk08
```

- 5 Verify that the disks are registered to a site.

```
# vxdisk listtag
```

For example:

```
# vxdisk listtag
```

DEVICE	NAME	VALUE
disk01	site	SFRAC1
disk02	site	SFRAC1
disk03	site	SFRAC1
disk06	site	SFRAC2
disk08	site	SFRAC2

- 6 Create a disk group for OCR and Vote Disks and another for Oracle data, with disks picked from both the sites. While the example below shows a single disk group, you can create as many as you need.

```
# vxdg -s init ocrvotedg disk05 disk07
# vxdg -s init oradatadg disk01 disk06
```

For the disks used in ocrvotedg, do not set site consistency to **on**.

- 7 Enable site-based allocation on the disk groups for each site.

```
# vxdg -g oradata addsite SFRAC1
# vxdg -g oradata addsite SFRAC2
```

- 8 If you are using an enclosure, set the tag on the enclosure for both sites.

```
# vxdg -o retain -g oradatadg setting encl:3pardata0 site=SFRAC1
# vxdg -o retain -g oradatadg setting encl:3pardata1 site=SFRAC2
```

- 9 Create one or more mirrored volumes in the disk group.

```
# vxassist -g ocrvotedg make ocrvotevol 2048m nmirror=2
# vxassist -g oradatadg make oradatavol 10200m nmirror=2
```

With the Site Awareness license installed on all hosts, the volume created has the following characteristics by default.

- The all sites attribute is set to ON; the volumes have at least one mirror at each site.
- The volumes are automatically mirrored across sites.
- The read policy (rdpol) is set to siteread.
- The volumes inherit the site consistency value that is set on the disk group.

- 10 From the CVM master, start the volumes for all the disk groups.

```
# vxvol -g ocrvotedg startall
# vxvol -g oradatadg startall
```

- 11 Create a file system on each volume and mount the same.

```
# mkfs -t vxfs /dev/vx/rdisk/ocrvotedg/ocrvotevol
# mkfs -t vxfs /dev/vx/rdisk/oradatadg/oradatavol
# mount -t vxfs -o cluster /dev/vx/dsk/ocrvotedg/ocrvotevol /ocr
# mount -t vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol /oradata
```

- 12 Touch two files, one for OCR and another for Vote Disk.

```
# touch /ocr/ocr /ocr/vote
```

Note: One Vote Disk is sufficient since it is already mirrored by VxVM.

Tuning guidelines for campus clusters

An important consideration while tuning an SF Oracle RAC campus cluster is setting the LLT peerinact time. Follow the guidelines below to determine the optimum value of peerinact time:

- Calculate the roundtrip time using `lltping` (1M).
- Evaluate LLT heartbeat time as half of the round trip time.
- Set the LLT peer trouble time as 2-4 times the heartbeat time.
- LLT peerinact time should be set to be more than 4 times the heart beat time.

Best practices for a campus cluster

The following best practices ensure a robust SF Oracle RAC campus cluster:

- Tag all the mirrored volumes in the campus cluster with appropriate site names. VxVM allocates storage from the correct site when creating or resizing a volume and when changing a volume's layout if the volume is tagged with site name.
- All volumes that have data required by the application must be evenly mirrored. Each site must have at least one mirror of all volumes hosting application data, including the FlashSnap log volume.
- Do not enable site consistency on VxVM snapshot volumes.
- Use redundant links for storage and private interconnects. DWDM can be used for storage and heartbeat together. Another redundant DWDM link can be used to prevent single point of failure. Separate switches and multiplexer / de-multiplexer devices should be used.
- Use Coordination Point Server as the third coordination point.
- Use the procedure for online replacement of coordination points, to replace disk based or Coordination Point Server based coordination points.

Setting up a replicated global cluster

This chapter includes the following topics:

- [Replication in the SF Oracle RAC environment](#)
- [About setting up a global cluster in an SF Oracle RAC environment](#)
- [Configuring an SF Oracle RAC global cluster at the primary site](#)
- [Configuring an SF Oracle RAC cluster at the secondary site](#)
- [Configuring replication for SF Oracle RAC clusters at both sites](#)
- [Modifying the ClusterService group for a global SF Oracle RAC cluster](#)
- [Defining the remote SF Oracle RAC cluster and heartbeat objects](#)
- [Configuring the VCS service groups for global SF Oracle RAC clusters](#)

Replication in the SF Oracle RAC environment

You can set up a primary SF Oracle RAC cluster for replication to a secondary SF Oracle RAC cluster by configuring global VCS service groups and using a replication technology. The Oracle RAC cluster at the secondary site can be a single node cluster. For example, you can have a two-node cluster on the primary site and a two-node or single-node cluster on the secondary site.

You can use one of the following replication technologies:

- Veritas Volume Replicator (VVR), which provides host-based volume replication. Using VVR you can replicate data volumes on a shared disk group in SF Oracle RAC.

- Supported hardware-based replication technologies. Using hardware-based replication you can replicate data from a primary array to a secondary array.
- Using SF Oracle RAC with VVR you can run a fire drill to verify the disaster recovery capability of your configuration.
 See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

About setting up a global cluster in an SF Oracle RAC environment

Configuring a global SF Oracle RAC cluster requires the coordination of many component setup tasks. The following procedures provide guidelines.

Tasks required to set up a global cluster:

- Configure a cluster at the primary site.
 See [“Configuring an SF Oracle RAC global cluster at the primary site”](#) on page 517.
- Configure a cluster at the secondary site.
 See [“Configuring an SF Oracle RAC cluster at the secondary site”](#) on page 519.
- Configure replication on clusters at both sites.
 See [“Configuring replication for SF Oracle RAC clusters at both sites”](#) on page 523.
- Configure VCS service groups for replication.
 See [“Modifying the ClusterService group for a global SF Oracle RAC cluster”](#) on page 524.
 See [“Defining the remote SF Oracle RAC cluster and heartbeat objects”](#) on page 526.
 See [“Configuring the VCS service groups for global SF Oracle RAC clusters”](#) on page 529.
- Test the HA/DR configuration.
- Upon successful testing, bring the environment into production.

SF Oracle RAC HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration, as the details of your configuration may differ from the examples given in the procedures. Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

Review the requirements information to make sure your configuration is supported for SF Oracle RAC.

- For product licensing information:

See “[About Veritas SFHA Solutions product licensing](#)” on page 68.

- For supported hardware and software:
Oracle documentation for additional requirements pertaining to your version of Oracle.
- To confirm the compatibility of your hardware, see the current compatibility list in the Symantec Technical Support website:
<http://www.symantec.com/docs/TECH170013>

SF Oracle RAC supports the following replication technologies through the use of Veritas replication agents:

Table 31-1 Supported replication options for SF Oracle RAC global clusters

Replication technology	Supported modes	Supported software
Veritas Volume Replicator (VVR) Supporting agents ■ RVGShared ■ RVGSharedPri ■ RVGLogOwner	■ Asynchronous replication ■ Synchronous replication	Host-based replication
EMC SRDF Supporting agent: SRDF	■ Asynchronous replication ■ Synchronous replication	All versions of Solutions Enabler
Hitachi True Copy Supporting agent: HTC	■ Asynchronous replication ■ Synchronous replication	All versions of the Hitachi CCI
IBM Metro Mirror Supporting agent: MetroMirror	Synchronous replication	All versions of IBM DSCLI. The MetroMirror agent is supported for DS6000 and DS8000 arrays
IBM SVC SVC CopyServices	■ Asynchronous replication ■ Synchronous replication	SSH access to the SVC

Table 31-1 Supported replication options for SF Oracle RAC global clusters
(continued)

Replication technology	Supported modes	Supported software
EMC Mirror View Supporting agent:MirrorView	■ Asynchronous replication ■ Synchronous replication: only individual LUNs may be replicated	All versions of NaviCLI

Note: Check your vendor's compatibility list for the supported software versions. The support listed above only exists if the host, HBA, and array combination is in your vendor's hardware compatibility list. Check your array documentation.

Note: All arrays must support SCSI-3 persistent reservations for SF Oracle RAC.

You can use the Veritas replication agents listed in the table above for global clusters that run SF Oracle RAC. The Veritas replication agents provide application failover and recovery support to your replication configuration. The agents provide this support for environments where data is replicated between clusters.

VCS agents control the direction of replication. They do not monitor the progress or status of replication. The replication agents manage the state of replicated devices that are attached to SF Oracle RAC nodes. The agents make sure that the system which has the resource online also has safe and exclusive access to the configured devices.

For instructions for configuring AT in a global cluster:

See the *Veritas Cluster Server Administrator's Guide*

Configuring an SF Oracle RAC global cluster at the primary site

You can use an existing SF Oracle RAC cluster or you can install a new SF Oracle RAC cluster for your primary site.

For planning information:

See [“Typical configuration of SF Oracle RAC global clusters for disaster recovery”](#) on page 39.

If you are using an existing cluster as the primary and you want to set up a global cluster, skip the steps below and proceed to configure your secondary cluster.

See [“Configuring an SF Oracle RAC cluster at the secondary site”](#) on page 519.

Note: You must have a GCO license enabled for a global cluster. If you are using VVR for replication, you must have a VVR license enabled.

If you do not have an existing cluster and you are setting up two new sites for an SF Oracle RAC global cluster, follow the steps below.

To set up the cluster and database at the primary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option enabled for a global cluster. If you are using VVR for replication, you must have it enabled.

- 4 Install and configure SF Oracle RAC.

For preparation:

For installation:

See [“About installing SF Oracle RAC ”](#) on page 80.

For configuration:

See [“About configuring SF Oracle RAC”](#) on page 114.

For a multi-node cluster, configure I/O fencing.

- 5 Verify the CVM group is online on all nodes in the primary cluster:

```
# hagrps -state cvm
```

- 6 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and database software.

You will need to set up:

- Local storage for Oracle RAC and CRS binaries
- Shared storage for OCR and Vote disk which is not replicated as part of the hardware-based or host-based replication

- Replicated storage for database files
- 7 Install and configure the Oracle RAC binaries:
 See [“About installing Oracle RAC”](#) on page 324.
 See [“Installing the Oracle Clusterware/Grid Infrastructure software”](#) on page 325.
 See [“Installing the Oracle RAC database software”](#) on page 334.

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
```

- 8 Identify the disks that will be replicated, create the required CVM disk group, volume, and file system.
- 9 Create the database on the file system you created in the previous step.
- 10 Configure the VCS service groups for the database.
- 11 Verify that all VCS service groups are online.

Configuring an SF Oracle RAC cluster at 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.
- You must use the same OS user and group IDs for your database for installation and configuration on both the primary and secondary clusters.
- You must use the same directory structure, name, permissions for the CRS/GRID and database binaries.

To set up the cluster on secondary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.

- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.
- 4 Install and configure SF Oracle RAC.
For preparation:
For installation:
See [“About installing SF Oracle RAC”](#) on page 80.
For configuration:
See [“About configuring SF Oracle RAC”](#) on page 114.
For a multi-node cluster, configure I/O fencing.
- 5 For a single-node cluster, do not enable I/O fencing. Fencing will run in disabled mode.
- 6 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and database software.
You will need to set up:
 - Local storage for Oracle RAC and CRS binaries
 - Shared storage for OCR and Vote disk which is not replicated
 - Replicated storage for database files

Note: You must use the same directory structure, name, permissions for the CRS/GRID and database binaries.

7 Install and configure the Oracle RAC binaries:

See [“About installing Oracle RAC”](#) on page 324.

See [“Installing the Oracle Clusterware/Grid Infrastructure software”](#) on page 325.

See [“Installing the Oracle RAC database software”](#) on page 334.

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
```

For Oracle 11gR2, use:

```
$ crsctl stat res -t
```

To set up the database for the secondary site

- 1** Do not create the database. The database will be replicated from the primary site.
 - If you are using hardware-based replication, the database, disk group, and volumes will be replicated from the primary site.
 Create the directory for the CFS mount point which will host the database data and control files.
 - If you are using VVR for replication, create an identical disk group and volumes for the replicated content with the same names and size as listed on the primary site.
 Create the directories for the CFS mount points as they are on the primary site. These will be used to host the database and control files when the failover occurs and the secondary is promoted to become the primary site.
- 2** Copy the init\$ORACLE_SID.ora file from \$ORACLE_HOME/dbs at the primary to \$ORACLE_HOME/dbs at the secondary.

- 3 Create the following subdirectories on the secondary site as you did on the primary site:

For Oracle RAC 10g:

```
$ mkdir -p /$ORACLE_BASE/admin/adump
$ mkdir -p /$ORACLE_BASE/admin/database_name/bdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/cdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/dpdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/hdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/udump
$ mkdir -p /$ORACLE_BASE/admin/database_name/pfile
```

For Oracle 11gR1 and 11gR2:

```
$ ORACLE_BASE/diag
$ ORACLE_BASE/admin
$ ORACLE_BASE/admin/adump
```

For oracle 11gR2 release only, on both the primary and secondary sites, edit the file:

```
$ORACLE_HOME/dbs/init$ORACLE_SID.ora
```

as

```
remote_listener = 'SCAN_NAME:1521'
SPFILE=<SPFILE NAME>
```

- 4 Configure listeners on the secondary site with same name as on primary. You can do this by one of the following methods:
 - Copy the listener.ora and tnsnames.ora files from the primary site and update the names as appropriate for the secondary site.
 - Use Oracle's netca utility to to configure the listener.ora and tnsnames.ora files on the secondary site.
- 5 On the secondary site, register the database using the `srvctl` command as Oracle user.

Registering the database has to be done only once from any node in the secondary cluster.

```
$ srvctl add database -d database_name -o oracle_home -p sp_file
```

- 6 To prevent automatic database instance restart, change the Management policy for the database (automatic, manual) to MANUAL using the `srvctl` command:

```
$ srvctl modify database -d database_name -y manual
```

You need only perform this change once from any node in the cluster.

- 7 Register the instances using `srvctl` command. Execute the following command on each node:

```
$ srvctl add instance -d database_name \  
-i instance_name -n node-name
```

- 8 If the secondary cluster has more than one node, you must create instances for all the nodes.

For example, if the database name is *racdb*, the instance name on *mercury* is *racdb1* and on *jupiter* is *racdb2*.

```
$ srvctl add instance -d racdb -i racdb1 - mercury - mercury
```

```
$ srvctl add instance -d racdb -i racdb2 - jupiter - jupiter
```

Configuring replication for SF Oracle RAC clusters at both sites

You must configure replication for the database files. Once replication is configured, make sure it is functioning correctly by testing before proceeding.

To configure replication at both sites

- 1 If you are using hardware-based replication, make sure that the replication management software for managing replication is installed on all nodes in both clusters.
- 2 At both sites, identify the disks on which the database resides at the primary site and associate them with the corresponding disks at the secondary site.

For VVR:

See “[Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): setting up the primary site](#)” on page 531.

For Hardware-based replication:

See your hardware documentation for details on setting up replication between the two sites.

- 3 Start replication between the sites.

For VVR:

See “[Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): starting replication](#)” on page 539.

See your hardware documentation for the appropriate procedures.

Modifying the ClusterService group for a global SF Oracle RAC cluster

You have configured VCS service groups for SF Oracle RAC on each cluster. Each cluster requires an additional virtual IP address associated with the cluster for cross-cluster communication. The VCS installation and creation of the ClusterService group typically involves defining this IP address.

Configure a global cluster by setting:

- Heartbeat
- Wide area cluster (wac)
- GCO IP (gcoip)
- remote cluster resources

See the *Veritas Cluster Server Administrator's Guide* for complete details on global clustering.

Modifying the global SF Oracle RAC cluster configuration using the wizard

The global clustering wizard completes the following tasks:

- Validates the ability of the current configuration to support a global cluster environment.
- Creates the components that enable the separate clusters, each of which contains a different set of GAB memberships, to connect and operate as a single unit.
- Creates the ClusterService group, or updates an existing ClusterService group.

Run the global clustering configuration wizard on each of the clusters; you must have the global clustering license in place on each node in the cluster.

To modify the ClusterService group for global clusters using the global clustering wizard

- 1 On the primary cluster, start the GCO Configuration wizard:

```
# /opt/VRTSvcs/bin/gcoconfig
```

- 2 The wizard discovers the NIC devices on the local system and prompts you to enter the device to be used for the global cluster. Specify the name of the device and press Enter.
- 3 If you do not have NIC resources in your configuration, the wizard asks you whether the specified NIC will be the public NIC used by all the systems. Enter **y** if it is the public NIC; otherwise enter **n**. If you entered **n**, the wizard prompts you to enter the names of NICs on all systems.
- 4 Enter the virtual IP address for the local cluster.
- 5 If you do not have IP resources in your configuration, the wizard prompts you for the netmask associated with the virtual IP. The wizard detects the netmask; you can accept the suggested value or enter another one.

The wizard starts running commands to create or update the ClusterService group. Various messages indicate the status of these commands. After running these commands, the wizard brings the ClusterService failover group online on any one of the nodes in the cluster.

Defining the remote SF Oracle RAC cluster and heartbeat objects

After configuring global clustering, add the remote cluster object to define the IP address of the cluster on the secondary site, and the heartbeat object to define the cluster-to-cluster heartbeat.

Heartbeats monitor the health of remote clusters. VCS can communicate with the remote cluster only after you set up the heartbeat resource on both clusters.

To define the remote cluster and heartbeat

- 1 On the primary site, enable write access to the configuration:

```
# haconf -makerw
```

- 2 On the primary site, 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 enable you to define the virtual IP address for the remote cluster.

For example:

```
# hahb -modify Icmp ClusterList rac_cluster102
# hahb -modify Icmp Arguments 10.11.10.102 -clus rac_cluster102
```

- 6 Save the configuration and change the access to read-only on the local cluster:

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

- 7 Complete step 4-6 on the secondary site using appropriate values to define the cluster on the primary site and its IP as the remote cluster for the secondary cluster.

- 8 Verify cluster status with the `hastatus -sum` command on both clusters.

```
# hstatus -sum
```

9 Display the global setup by executing `haclus -list` command.

```
# haclus -list
      rac_cluster101
      rac_cluster102
```

Example of heartbeat additions to the `main.cf` file on the primary site:

```
.
.
remotecluster rac_cluster102 (
Cluster Address = "10.11.10.102"
)
heartbeat Icmp (
    ClusterList = { rac_cluster102 }
    Arguments @rac_cluster102 = { "10.11.10.102" }
)

system galaxy (
)

.
.
```

Example heartbeat additions to the `main.cf` file on the secondary site:

```
.
.
remotecluster rac_cluster101 (
    Cluster Address = "10.10.10.101"
)

heartbeat Icmp (
    ClusterList = { rac_cluster101 }
    Arguments @rac_cluster101 = { "10.10.10.101" }
)

system mercury (
)

.
.
```

See the *Veritas Cluster Server Administrator's Guide* for details for configuring the required and optional attributes of the heartbeat object.

Configuring the VCS service groups for global SF Oracle RAC clusters

To configure VCS service groups for global clusters

- 1 Configure and enable global groups for databases and resources.
 - Configure VCS service groups at both sites.
 - Configure the replication agent at both sites.
 - Make the RAC service group a global service group, enabling failover across clusters.
 - For example:
See [“Modifying the VCS Configuration on the Primary Site”](#) on page 546.
- 2 To test real data in an environment where HA/DR has been configured, schedule a planned migration to the secondary site for testing purposes.

For example:

See [“Migrating the role of primary site to the secondary site”](#) on page 557.

- 3 Upon successful testing, bring the environment into production.

For more information about VCS replication agents:

See the *Veritas Cluster Server Bundled Agents Guide*

For complete details on using VVR in a shared disk environment:

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator’s Guide*.

Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring a global SF Oracle RAC cluster using VVR for replication](#)
- [Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): setting up the primary site](#)
- [Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): setting up the secondary site](#)
- [Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): starting replication](#)
- [Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): configuring cluster resources](#)
- [Managing a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\)](#)

About configuring a global SF Oracle RAC cluster using VVR for replication

Before configuring clusters for global clustering, make sure both clusters have product and database software installed and configured.

Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster and VVR enabled.

See [“About Veritas SFHA Solutions product licensing”](#) on page 68.

After setting up two clusters running SF Oracle RAC, you can configure a global cluster environment with VVR. You must modify both cluster configurations to support replication in the global cluster environment.

Configuring SF Oracle RAC for global clusters requires:

- Setting up both clusters as part of a global cluster environment.
See [“About setting up a global cluster in an SF Oracle RAC environment”](#) on page 515.
- Setting up replication for clusters at both sites.
See [“Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): setting up the primary site”](#) on page 531.
See [“Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): setting up the secondary site”](#) on page 534.
- Starting replication of the database.
See [“Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): starting replication”](#) on page 539.
- Configuring VCS for replication on clusters at both sites.
See [“Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator \(VVR\): configuring cluster resources”](#) on page 541.

Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): setting up the primary site

Setting up replication with VVR in a global cluster environment involves the following tasks:

- If you have not already done so, create a disk group to hold data volume, SRL, and RVG on the storage on the primary site. For example:
- Creating the Storage Replicator Log (SRL) in the disk group for the database.
See [“Creating the data and SRL volumes on the primary site”](#) on page 532.
- Creating the Replicated Volume Group (RVG) on the primary site.
See [“Setting up the Replicated Volume Group \(RVG\) on the primary site”](#) on page 533.
- Editing the `/etc/vx/vras/.rdg` files on the primary site after creating the data and SRL volumes on the secondary site.
See [“Editing the `/etc/vx/vras/.rdg` files”](#) on page 535.

Creating the data and SRL volumes on the primary site

Create the data volume if you do not have one already.

- The data volume on the secondary site has the same name and the same size as the data volume on the primary site.
- The data volume and SRL volume should exist in the same disk group.
- Mirror the data volume in the absence of hardware-based mirroring.

To create the data volume on the primary site

- ◆ In the disk group created for the Oracle RAC database, create a data volume of same size as that in primary for data; in this case, the *rac_vol1* volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_vol1 6600M nmirror=2 disk1 disk2
```

Create the SRL. The SRL is a volume in the RVG. The RVG also holds the data volumes for replication.

- The SRL on the secondary site has the same name and the same size as the SRL on the primary site.
- If possible, create SRLs on disks without other volumes.
- Mirror SRLs and in the absence of hardware-based mirroring.

In the example procedure below, *oradatadg* is the disk group and *rac1_vol* is the data volume to be replicated.

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 CVM master or the slave:

```
# vxdctl -c mode
```

- 3 On the CVM master node, issue the following command:

```
# vxassist -g oradatadg make rac1_srl 6800M nmirror=2 disk4 disk5
```

Note: Assume that for the example setup that *disk4* and *disk5* are already added and are part of the same disk group. They are used in this step for mirroring and creation of the SRL.

- 4 If the SRL volume is not already started, start the SRL volume by starting all volumes in the disk group:

```
# vxvol -g oradatadg startall
```

Setting up the Replicated Volume Group (RVG) on the primary site

Before creating the RVG on the primary site, make sure the volumes and CVM group are active and online.

To review the status of replication objects on the primary site

- 1 Verify the volumes you intend to include in the group are active.
- 2 Review the output of the `hagrp -state cvm` command to verify that the CVM group is online.
- 3 On each site, verify `vradmin` is running:

```
# ps -ef |grep vradmin
root 536594 598036 0 12:31:25 0 0:00 grep vradmin
```

If `vradmin` is not running start it:

```
# vxstart_vvr
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradminind: 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/vradminind
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 *.*
```

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

To create the primary RVG

- 1 Determine which node is the CVM master node by entering:

```
# vxdctl -c mode
```

- 2 To create the `rac1_rvg` RVG, you must run the following on the master node:

```
# 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 a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): setting up 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 to hold the data volume, SRL, and RVG on the storage on the secondary site. You must match the names and sizes of these volumes with the names and sizes of the volumes on the primary site.
See [“Creating the data and SRL volumes on the secondary site”](#) on page 535.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.
See [“Editing the /etc/vx/vras/.rdg files”](#) on page 535.
- 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 536.

- Create the replication objects on the secondary site.
See [“Setting up the disk group on secondary site for replication”](#) on page 537.

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 match the sizes and names of the corresponding volumes in the primary site before you create the disk group.
- The disk group must match the size and name of the disk group at the primary site.
- Create the data and SRL volumes on different disks in the disk group. Use the `vxdisk -g diskgroup list` command to list the disks in the disk group.
- Mirror the volumes.

To create the data and SRL volumes on the secondary site

- 1 In the disk group created for the Oracle RAC database, create a data volume of same size as that in primary for data; in this case, the `rac_vol1` volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_vol1 6800M nmirror=2 disk1 disk2
```

- 2 Create the volume for the SRL, using the same name and size of the equivalent volume on the primary site. Create the volume on different disks from the disks for the database volume, but on the same disk group that has the data volume:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk6
```

Editing the `/etc/vx/vras/.rdg` files

Editing the `/etc/vx/vras/.rdg` file on the secondary site enables VVR to replicate the disk group from the primary site to the secondary site. On each node, VVR uses the `/etc/vx/vras/.rdg` file to check the authorization to replicate the RVG on the primary site to the secondary site. The file on each node in the secondary site must contain the primary disk group ID, and likewise, the file on each primary system must contain the secondary disk group ID.

To edit the /etc/vx/vras/.rdg files

- 1 On a node in the primary site, display the primary disk group ID:

```
# vxprint -l diskgroup

.....
```

- 2 On each node in the secondary site, edit the /etc/vx/vras/.rdg file and enter the primary disk group ID on a single line.
- 3 On each cluster node of the primary cluster, edit the /etc/vx/vras/.rdg file and enter the secondary disk group ID on a single line.

Setting up IP addresses for RLINKs on each cluster

Creating objects with the vradm command requires resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.

To set up IP addresses for RLINKs on each cluster

- 1 Using the following command, determine whether a node is the CVM master or the slave:

```
# vxdctl -c mode
```

You must configure rlinks on the CVM master node.

- 2 For each RVG running on each cluster, set up a virtual IP address on one of the nodes of the cluster. These IP addresses are part of the RLINK.

The example assumes for the cluster on the primary site:

- The public network interface is eth0:1
- The virtual IP address is 10.10.9.101
- The net mask is 255.255.255.0
- # ifconfig eth0:1 inet 10.10.9.101 netmask 255.255.255.0 up

- 3 Use the same commands with appropriate values for the interface, IP address, and net mask on the secondary site.

The example assumes for the secondary site:

- The public network interface is eth0:1
- virtual IP address is 10.11.9.102
- net mask is 255.255.255.0

- 4 Define the virtual IP addresses to correspond to a host name in the virtual cluster on the primary site and a host name in the virtual cluster on the secondary site.

Update the `/etc/hosts` file on all the nodes on both the primary and secondary sites.

The examples assume:

- `rac_cluster101` has IP address 10.10.9.101
- `rac_cluster102` has IP address 10.11.9.102

- 5 Use the `ping` command to verify the links are functional.

Setting up the disk group on secondary site for replication

Create the replication objects on the secondary site from the master node of the primary site, using the `vradmin` command.

To set up the disk group on the secondary site for replication

- 1 Issue the command in the following format from the cluster on the primary site:

```
# vradmin -g dg_pri addsec rvg_pri pri_host sec_host
```

where:

- `dg_pri` is the disk group on the primary site that VVR will replicate. For example: `rac1_vol`
- `rvg_pri` is the RVG on the primary site. For example: `rac1_rvg`
- `pri_host` is the virtual IP address or resolvable virtual host name of the cluster on the primary site.
For example: `rac_cluster101_1`
- `sec_host` is the virtual IP address or resolvable virtual host name of the cluster on the secondary site.
For example: `rac_cluster102_1`

For example, the command to add the cluster on the primary site to the Replicated Data Set (RDS) is:

```
vradmin -g oradatadg addsec rac1_rvg \  
rac_cluster101_1  
rac_cluster102_1
```

On the secondary site, the above command performs the following tasks:

- Creates an RVG within the specified disk group using the same name as the one for the primary site
 - Associates the data and SRL volumes that have the same names as the ones on the primary site with the specified RVG
 - Adds a data change map (DCM) for the data volume
 - Creates cluster RLINKS for the primary and secondary sites with the default names; for example, the "primary" RLINK created for this example is *rlk_rac_cluster101_1_rac1_rvg* and the "secondary" RLINK created is *rlk_rac_cluster102_1_rac1_rvg*.
- 2 Verify the list of RVGs in the RDS by executing the following command.

```
# vradmin -g oradatadg -l printrvg
```

For example:

```
Replicated Data Set: rac1_rvg
Primary:
HostName: 10.180.88.187 <localhost>
RvgName: rac1_rvg
DgName: rac1_vol
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_rac_cluster102_1_rac1_rvg, detached=on,
synchronous=off
Secondary:
HostName: 10.190.99.197
RvgName: rac1_rvg
DgName: oradatadg
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_rac_cluster101_1_rac1_rvg, detached=on,
synchronous=off
```

Note: Once the replication is started the value of the detached flag will change the status from ON to OFF.

Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): starting replication

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: `srprot_autodcm`
- Packet size: `packet_size=8400`
- Network protocol: `protocol=UDP`

Method of initial synchronization:

- Automatic synchronization
- Full synchronization with Storage Checkpoint

For guidelines on modifying these settings and information on choosing the method of replication for the initial synchronization:

See the *Veritas Volume Replicator Administrator's Guide*

Starting replication using automatic synchronization

Use the `vradmin` command to start replication or the transfer of data from the primary site to the secondary site over the network. Because the cluster on the secondary site uses only one host name, the command does not require the `sec_host` argument.

To start replication using automatic synchronization

- ◆ From the primary site, use the following command to automatically synchronize the RVG on the secondary site:

```
vradmin -g disk_group -a startrep pri_rvg sec_host
```

where:

- `disk_group` is the disk group on the primary site that VVR will replicate
- `pri_rvg` is the name of the RVG on the primary site
- `sec_host` is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -a startrep rac1_rvg
rac_cluster102
```

Starting replication using full synchronization with Storage Checkpoint

Use the `vradmin` command with the Storage Checkpoint option to start replication using full synchronization with Storage Checkpoint.

To start replication using full synchronization with Storage 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 Storage Checkpoint on the primary site
- `pri_rvg` is the name of the RVG on the primary site
- `sec_host` is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -c rac1_ckpt syncrvg rac1_rvg
rac_cluster102
```

- 2 To start replication after full synchronization, enter the following command:

```
# vradmin -g oradatadg -c rac1_ckpt startrep rac1_rvg
rac_cluster102
```

Verifying replication status

Verify that replication is properly functioning.

To verify replication status

- 1 Check the status of VVR replication:

```
# vradmin -g disk_group_name repstatus rvg_name
```

- 2 Review the `flags` output for the status. The output may appear as `connected` and `consistent`. For example:

```
# vxprint -g oradatadg -l rlk_rac_cluster102_1_rac1_rvg
Rlink: rlk_rac_cluster102_1_rac1_rvg
info: timeout=500 packet_size=8400 rid=0.1078
      latency_high_mark=10000 latency_low_mark=9950
      bandwidth_limit=none
state: state=ACTIVE
      synchronous=off latencyprot=off srlprot=autodcm
.
.
protocol: UDP/IP
checkpoint: rac1_ckpt
flags: write enabled attached consistent connected
asynchronous
```

Configuring a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR): configuring cluster resources

After configuring both clusters for global clustering and setting up the Oracle RAC database for replication, configure VCS to provide high availability for the database. Specifically, configure VCS agents to control the cluster resources, including the replication resources.

To view the sample `main.cf` files on your system:

```
# cd /etc/VRTSvcs/conf/sample_rac
# ls *sfrac* sfrac07_main.cf sfrac08_main.cf
```

The following files include CVM/VVR configuration examples:

```
sfrac_07_main.cf    sfrac_08_main.cf
```

The following sample `main.cf` files illustrate the VCS configuration changes after setting up an existing Oracle RAC database for replication:

See [“sfrac07_main.cf and sample08_main.cf files”](#) on page 644.

Note: The example procedures illustrate the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

About modifying the VCS configuration for replication

The following resources must be configured or modified for replication:

Table 32-1 Cluster resources required for replication

Cluster resources	Configuration required
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 683.

Table 32-1 Cluster resources required for replication (*continued*)

Cluster resources	Configuration required
RVGSharedPri resource	Add the RVGSharedPri resource to the existing Oracle RAC database service group. The CVMVolDg resource must be removed from the existing Oracle RAC database service group.
Oracle RAC database service group	The existing Oracle RAC database service group is a parallel group consisting of the Oracle RAC database resource, CVMVolDg resource, and CFSMount resource (if the database resides in a cluster file system). Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute.

For detailed examples of service group modification:

See [“Configuration examples before and after modification”](#) on page 543.

For more information on service replication resources:

See the *Veritas™ Cluster Server Bundled Agents Reference Guide*.

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 32-1](#)
- Configuration after modification:
[Figure 32-2](#)

All of the dependencies between parent and child groups are online local firm. The CVM service group is the same in all illustrations because its definition requires no changes.


Configuration before modification for replication:

Figure 32-1 Illustration of dependencies before modification for replication



Configuration after modification for replication:

Figure 32-2 Illustration of dependencies after modification for replication for Oracle RAC



Modifying the VCS Configuration on the Primary Site

The following are the procedural highlights required to modify the existing VCS configuration on the primary site:

- Configure two service groups:
 - A log owner group including the RVGLogowner resource.
 - An RVG group including the RVGShared resource replication objects.
- Add the RVGSharedPri resource to the existing Oracle RAC database service group and define this group as a global group by setting the ClusterList and ClusterFailOverPolicy attributes.
- Move the CVMVolDg resource from the existing Oracle RAC database service group to the newly created RVG group.

Note: The example procedure illustrates the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

To modify VCS on the primary site

- 1 Log into one of the nodes on the primary cluster.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while you make changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Review the sample configuration file after the SF Oracle RAC installation.

Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:

- RVGLogowner resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).
- IP resource
- NIC resources

The following are examples of RVGLogowner service group for the different platforms.

```
group rlogowner (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStartList = { galaxy,nebula }
)

IP logowner_ip (
    Device = eth0
    Address = "10.10.9.101"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = eth0
    NetworkType = ether
    NetworkHosts = "10.10.8.1"
)

RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

- 5 Add the RVG service group using the appropriate values for your cluster and nodes.

Example `RVGgroup` service group:

```
group RVGgroup (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy,nebula }
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

CVMVolDg racdata_voldg (
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)

requires group cvm online local firm
racdata_rvg requires racdata_voldg
```

- 6 Modify the Oracle RAC service group using the appropriate values for your cluster and nodes:
 - Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute. See the attribute in bold in the example that follows.

Note: This action must be performed on the primary or secondary site, but not on both.

Note: If you are using the Java Console, the secondary cluster must be configured before you can configure a service group as a global group.

- Add the `ClusterFailOverPolicy` cluster attribute. Symantec recommends using the `Manual` value. See the attribute in bold in the example.
- Add the `RVGSharedPri` resource to the group configuration.
- Remove the `CVMVolDg` resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
- Specify the service group (`online`, `local`, `firm`) to depend on the RVG service group.

- Remove the existing dependency of the Database service group on the CVM service group. Remove the line:

```
requires group CVM online local firm
```

- Remove the existing dependency between the CFSSMount for the database and the CVMVoldg for the Oracle RAC database. Remove the line:

```
oradata_mnt requires oradata_voldg
```

The following is an example of an Oracle RAC database service group configured for replication:

```
group database_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    ClusterList = { rac_cluster101 = 0, rac_cluster102 = 1 }
    Parallel = 1
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { galaxy,nebula }
)

CFSSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

Oracle rac_db (
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Owner = Oracle
    Home = "/oracle/orahome"

    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
rac_db requires oradata_mnt
```

- 7 Save and close the main.cf file.
- 8 Use the following command to verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 9 Stop and restart VCS.

```
# hstop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hstart
```

Modifying the VCS Configuration on the Secondary Site

The following are highlights of the procedure to modify the existing VCS configuration on the secondary site:

- Add the log owner and RVG service groups.
- Add a service group to manage the Oracle RAC database and the supporting resources.
- Define the replication objects and agents, such that the cluster at the secondary site can function as a companion to the primary cluster.

The following steps are similar to those performed on the primary site.

Note: The example procedure illustrates the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

To modify VCS on the secondary site

- 1 Log into one of the nodes on the secondary site as root.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while making changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Edit the CVM group on the secondary site.

Review the sample configuration file after the SF Oracle RAC installation to see the CVM configuration.

In our example, the secondary site has rac_cluster102 consisting of the nodes mercury and jupiter. To modify the CVM service group on the secondary site, use the CVM group on the primary site as your guide.

- 5 Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:

- RVGLogowner resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).
- IP resource
- NIC resources

Example RVGLogowner service group:

```
group rlogowner (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoStartList = { mercury, jupiter }
)

IP logowner_ip (
    Device = eth0
    Address = "10.11.9.102"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = eth0
    NetworkHosts = { "10.10.8.1" }
    NetworkType = ether
)

RVGLogowner logowner (
    RVG = rac1_rvg
```

```

        DiskGroup = oradatadg
    )

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic

```

- 6 Add the RVG service group using the appropriate values for your cluster and nodes.

The following is an example RVGgroup service group:

```

group RVGgroup (
    SystemList = { mercury = 0, jupiter = 1 }
    Parallel = 1
    AutoStartList = { mercury, jupiter }
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

CVMVolDg racdata_voldg (
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)

requires group cvm online local firm
racdata_rvg requires racdata_voldg

```

- 7 Add an Oracle RAC service group. Use the Oracle RAC service group on the primary site as a model for the Oracle RAC service group on the secondary site.
 - Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute.

Note: This action must be performed on the primary or secondary site, but not on both.

- Assign this global group the same name as the group on the primary site; for example, *database_grp*.

- Include the ClusterList and ClusterFailOverPolicy cluster attributes. Symantec recommends using the Manual value.
- Add the RVGSharedPri resource to the group configuration.
- Remove the CVMVolDg resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
- Specify the service group to depend (online, local, firm) on the RVG service group.

Example of the Oracle RAC group on the secondary site:

```
group database_grp (
    SystemList = { mercury = 0, jupiter = 1 }
    ClusterList = { rac_cluster102 = 0, rac_cluster101 = 1 }
    Parallel = 1
    OnlineRetryInterval = 300
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { mercury, jupiter }
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"
    Critical = 0
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

Oracle rac_db (
    Sid @mercury = vrts1
    Sid @jupiter = vrts2
    Owner = Oracle
    Home = "/oracle/orahome"
    Pfile @mercury = "/oracle/orahome/dbs/initvrts1.ora"
```

```
Pfile @jupiter = "/oracle/orahome/dbs/initvrts2.ora"
StartUpOpt = SRVCTLSTART
ShutDownOpt = SRVCTLSTOP
)
```

```
requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
rac_db requires oradata_mnt
```

- 8 Save and close the `main.cf` file.
- 9 Use the following command to verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 10 Stop and restart VCS.

```
# hstop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hstart
```

- 11 Verify that VCS brings all resources online. On one node, enter the following command:

```
# hagrps -display
```

The Oracle RAC, RVG, and CVM groups are online on both nodes of the primary site. The RVGLogOwner and ClusterService groups are online on one node of the cluster. If either the RVG group or the RVGLogOwner group is partially online, manually bring the groups online using the `hagrps -online` command. This information applies to the secondary site, except for the Oracle RAC group which must be offline.

- 12 Verify the service groups and their resources that are brought online. On one node, enter the following command:

```
# hagrps -display
```

The Oracle RAC service group is offline on the secondary site, but the ClusterService, CVM, RVG log owner, and RVG groups are online.

This completes the setup for an SF Oracle RAC global cluster using VVR for replication. Symantec recommends testing a global cluster before putting it into production.

Managing a global SF Oracle RAC cluster using Veritas Volume Replicator (VVR)

For information on using the VCS commands for global clusters:

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

If you have two SF Oracle RAC clusters configured to use VVR for replication, the following administrative functions are available:

Table 32-2

Migration of the role of the primary site to the remote site	Migration of the role of the primary site to the remote site is a planned transfer of the role of primary replication host from one cluster to a remote cluster. This transfer enables the application on the remote cluster to actively use the replicated data. The former primary cluster becomes free for maintenance or other activity.

Table 32-2 (continued)

Takeover of the primary site role by the secondary site	Takeover of the primary site role by the secondary site occurs when an unplanned event (such as a disaster) causes a failure, making it necessary for the applications using the replicated data to be brought online on the remote cluster.

VCS agents manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

Table 32-3

VCS replication agents bundled with VVR	Description
DNS agent	The DNS agent updates the canonical name-mapping in the domain name server after a wide-area failover. See the Veritas Cluster Server Bundled Agents Reference Guide for more information about the agent.
RVG agent	The RVG agent manages the Replicated Volume Group (RVG). Specifically, it brings the RVG online, monitors read-write access to the RVG, and takes the RVG offline. Use this agent when using VVR for replication. RVGPrimary agent The RVGPrimary agent attempts to migrate or take over a Secondary to a Primary following an application failover. The agent has no actions associated with the offline and monitor routines.

VCS provides agents for other array-based or application-based solutions. For more information about the RVG and RVGPrimary agents.

See the *Veritas™ Cluster Server Bundled Agents Reference Guide*.

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

Migrating the role of primary site to the secondary site

After configuring the replication objects within VCS, you can use VCS commands to migrate the role of the cluster on the primary site to the remote cluster. In the procedure below, VCS takes the replicated database service group, *database_grp*, offline on the primary site and brings it online on the secondary site; the secondary site now assumes the role of the primary site.

Note: The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

To migrate the role of primary site to the remote site

- 1 From the primary site, use the following command to take the Oracle service group offline on all nodes.

```
# hagrp -offline database_grp -any
```

Wait for VCS to take all Oracle service groups offline on the primary site.

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
# vxlink -g data_disk_group status rlk_rac_cluster102_priv_rac1_rvg
```

Where `rlk_rac_cluster101_priv_rac1_rvg` is the RLINK.

- 3 On the secondary site, which is now the new primary site, bring the Oracle service group online on all nodes:

```
# hagrp -online database_grp -any
```

After migrating the role of the primary site to the secondary site, you can use VCS commands to migrate the role of the cluster on the new primary site to the original primary site. In the procedure below, VCS takes the replicated database service group, *database_grp*, offline on the new primary (former secondary) site and brings it online on the original primary site; the original primary site now resumes the role of the primary site.

Note: The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

To migrate the role of new primary site back to the original primary site

- 1 Make sure that all CRS resources are online, and switch back the group *database_grp* to the original primary site.

Issue the following command on the remote site:

```
# hagrps -offline database_grp -any
```

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the current primary cluster.

For example:

```
# vxlink -g data_disk_group status rlk_rac_cluster101_priv_rac1_rvg
```

Where *rlk_rac_cluster101_priv_rac1_rvg* is the RLINK.

- 3 Make sure that *database_grp* is offline on the new primary site. Then, execute the following command on the original primary site to bring the *database_grp* online:

```
# hagrps -online database_grp -any
```

Taking over the primary role by the remote cluster

Takeover occurs when the remote cluster on the secondary site starts the application that uses replicated data. This situation may occur if the secondary site perceives the primary site as dead, or when the primary site becomes inaccessible (perhaps for a known reason). For a more detailed description of concepts of taking over the primary role:

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

Before enabling the secondary site to take over the primary role, the administrator on the secondary site must "declare" the type of failure at the remote (primary, in this case) site and designate the failure type using one of the options for the `haclus` command.

Table 32-4 Options for the remote cluster to take over the primary role

Takeover options	Description
Disaster	When the cluster on the primary site is inaccessible and appears dead, the administrator declares the failure type as "disaster." For example, fire may destroy a data center, including the primary site and all data in the volumes. After making this declaration, the administrator can bring the service group online on the secondary site, which now has the role as "primary" site.
Outage	When the administrator of a secondary site knows the primary site is inaccessible for a known reason, such as a temporary power outage, the administrator may declare the failure as an "outage." Typically, an administrator expects the primary site to return to its original state. After the declaration for an outage occurs, the RVGSharedPri agent enables DCM logging while the secondary site maintains the primary replication role. After the original primary site becomes alive and returns to its original state, DCM logging makes it possible to use fast fail back resynchronization when data is resynchronized to the original cluster. Before attempting to resynchronize the data using the fast fail back option from the current primary site to the original primary site, take the precaution at the original primary site of making a snapshot of the original data. This action provides a valid copy of data at the original primary site for use in the case the current primary site fails before the resynchronization is complete.
Disconnect	When both clusters are functioning properly and the heartbeat link between the clusters fails, a split-brain condition exists. In this case, the administrator can declare the failure as "disconnect," which means no attempt will occur to take over the role of the primary site at the secondary site. This declaration is merely advisory, generating a message in the VCS log indicating the failure results from a network outage rather than a server outage.
Replica	In the rare case where the current primary site becomes inaccessible while data is resynchronized from that site to the original primary site using the fast fail back method, the administrator at the original primary site may resort to using a data snapshot (if it exists) taken before the start of the fast fail back operation. In this case, the failure type is designated as "replica".

The examples illustrate the steps required for an outage takeover and resynchronization.

To take over after an outage

- 1 From any node of the secondary site, issue the `haclus` command:

```
# haclus -declare outage -clus rac_cluster101
```

- 2 After declaring the state of the remote cluster, bring the `database_grp` service group online on the secondary site. For example:

```
# hagrps -online -force database_grp -any
```

To resynchronize after an outage

- 1 On the original primary site, create a snapshot of the RVG before resynchronizing it in case the current primary site fails during the resynchronization. Assuming the disk group is `data_disk_group` and the RVG is `rac1_rvg`, type:

```
# vxrvrg -g data_disk_group -F snapshot rac1_rvg
```

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for details on RVG snapshots.

- 2 Resynchronize the RVG. From any node of the current primary site, issue the `hares` command and the `-action` option with the `fbssync` action token to resynchronize the RVGSharedPri resource. For example:

```
# hares -action ora_vvr_shpri fbssync -sys mercury
```

- 3 Perform one of the following commands, depending on whether the resynchronization of data from the current primary site to the original primary site is successful:

- If the resynchronization of data is successful, use the `vxrvrg` command with the `snapback` option to reattach the snapshot volumes on the original primary site to the original volumes in the specified RVG:

```
# vxrvrg -g data_disk_group snapback rac1_rvg
```

- A failed attempt at the resynchronization of data (for example, a disaster hits the primary RVG when resynchronization is in progress) could generate inconsistent data.

You can restore the contents of the RVG data volumes from the snapshot taken in step 1:

```
# vxrvg -g data_disk_group snaprestore rac1_rvg
```

Troubleshooting Veritas Volume Replicator (VVR) components of SF Oracle RAC

If the rlink is not up to date, use the `hares -action` command with the `resync` action token to synchronize the RVG.

To update the rlink

- ◆ Enter the following command example on any node in the primary cluster, specifying the RVGSharedPri resource:

```
# hares -action ora_vvr_shpri resync -sys galaxy
```

Uninstallation of SF Oracle RAC

- [Chapter 33. Preparing to uninstall SF Oracle RAC from a cluster](#)
- [Chapter 34. Uninstalling SF Oracle RAC from a cluster](#)

Preparing to uninstall SF Oracle RAC from a cluster

This chapter includes the following topics:

- [About uninstalling SF Oracle RAC from a cluster](#)
- [Options for uninstalling SF Oracle RAC](#)
- [Preparing to uninstall SF Oracle RAC from a cluster](#)


About uninstalling SF Oracle RAC from a cluster

You can uninstall SF Oracle RAC using the `uninstallsrac` program.

Note: After you uninstall SF Oracle RAC, you cannot access the Oracle database as Veritas Volume Manager and Veritas File System are uninstalled from the cluster. Make sure that you back up the Oracle database before you uninstall SF Oracle RAC.

[Figure 33-1](#) illustrates the steps that are required to uninstall SF Oracle RAC from a cluster.

Figure 33-1 SF Oracle RAC uninstallation



Options for uninstalling SF Oracle RAC

Table 33-1 lists the available options for uninstalling SF Oracle RAC:

Table 33-1 Options for uninstalling SF Oracle RAC

Options	Description
SF Oracle RAC uninstallation program	Use the <code>uninstallsrac</code> program to uninstall SF Oracle RAC.

Table 33-1 Options for uninstalling SF Oracle RAC (*continued*)

Options	Description
Response file	Use a response file to automate or perform an unattended installation of SF Oracle RAC. See “Uninstalling SF Oracle RAC using a response file” on page 581.

Preparing to uninstall SF Oracle RAC from a cluster

Perform the steps in the following procedure before you uninstall SF Oracle RAC from a cluster.

To prepare to uninstall SF Oracle RAC from a cluster

- Stop Oracle instances.
See [“Stopping Oracle instances”](#) on page 566.
- Back up the Oracle database.
See [“Backing up the Oracle database”](#) on page 567.
- Unlink the SF Oracle RAC libraries from Oracle RAC libraries (optional).
See [“Unlinking the SF Oracle RAC libraries from Oracle RAC”](#) on page 567.
- Uninstalling Oracle RAC (optional)
See [“Uninstalling Oracle RAC \(optional\)”](#) on page 567.
- Remove root disk encapsulation.
See [“Removing root disk encapsulation”](#) on page 568.
- Stop the applications that use CFS (outside of VCS control).
See [“Stopping the applications that use CVM or CFS \(outside of VCS control\)”](#) on page 569.
- Unmount CFS file systems (outside of VCS control).
See [“Unmounting CFS file systems \(outside of VCS control\)”](#) on page 569.
- Stop VCS.
See [“Stopping VCS”](#) on page 569.

- 9 Stop the applications that use VxFS (outside of VCS control).
See [“Stopping the applications that use VxVM or VxFS \(outside of VCS control\)”](#) on page 570.
- 10 Unmount VxFS file systems (outside of VCS control).
See [“Unmounting VxFS file systems \(outside of VCS control\)”](#) on page 570.

Stopping Oracle instances

You need to stop Oracle Clusterware and the Oracle instances on the cluster nodes where you want to uninstall SF Oracle RAC. Before you stop the Oracle instances, stop the applications that are dependent on the service groups that contain Oracle.

The procedure in this section provides instructions to stop the instances on a two-node cluster; the nodes are galaxy and nebula. Depending on the VCS configuration, the procedure to stop Oracle instances may vary.

To stop Oracle instances

- 1 Log in as the superuser on one of the nodes in the cluster.
- 2 On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

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

For example:

```
# hagrps -offline oral -sys galaxy
```

```
# hagrps -offline oral -sys nebula
```

These commands stop the Oracle resources under VCS control.

If the database is not managed by VCS, stop the Oracle database as follows:

```
$ srvctl stop database -d db_name
```

- 3 Verify that the state of the Oracle (if the database is managed by VCS) and CVM service groups are offline and online respectively.

```
# hagrps -state
```

Group	Attribute	System	Value
oral	State	galaxy	OFFLINE
oral	State	nebula	OFFLINE
cvm	State	galaxy	ONLINE
cvm	State	nebula	ONLINE

Backing up the Oracle database

If you plan to retain the Oracle database, you must back up the Oracle database.

For instructions on backing up the Oracle database, see the Oracle documentation.

Unlinking the SF Oracle RAC libraries from Oracle RAC

If you want to reuse the Oracle RAC binaries, unlink the SF Oracle RAC libraries from Oracle RAC.

Note: The Oracle RAC binaries that are located on Veritas File System will not be accessible after the uninstallation.

To unlink the SF Oracle RAC libraries from Oracle RAC libraries

- 1 Stop Oracle's Clusterware:

```
# hares -offline cssd_resname -sys node_name
```

- 2 Unlink the Veritas VCSMM library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

```
$ cd $GRID_HOME/lib
```

```
$ ln -sf $GRID_HOME/lib/libskgxns.so libskgxns2.so
```

- 3 Unlink the Veritas ODM library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

```
$ cd $ORACLE_HOME/lib
```

```
$ ln -sf $ORACLE_HOME/lib/libodm11.so libodm11.so
```

Uninstalling Oracle RAC (optional)

You cannot use the Oracle database after you uninstall SF Oracle RAC. However, you can continue using the Oracle RAC software, provided the Oracle Clusterware and database binaries are not installed on Veritas file system. To continue using

Oracle RAC, unlink the SF Oracle RAC libraries from Oracle RAC as described in the preceding topic.

For instructions on uninstalling Oracle RAC, see the Oracle documentation.

Removing root disk encapsulation

Perform this step only if you plan to remove the VxVM and VVR RPMs.

If you have VxVM and VVR installed, you need to indicate to the installer whether or not you want to remove the VxVM RPMs from all nodes in the cluster. If you want to remove these RPMs, you need to ensure that the root disk is not encapsulated. The uninstallation fails if you choose to remove these RPMs while the root disk is encapsulated.

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following procedure.

To remove root disk encapsulation

- 1 Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- 2 Convert all the encapsulated volumes in the root disk to make them accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

- 3 To check if the root disk is unencapsulated:

```
# df -v /
```

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

Stopping the applications that use CVM or CFS (outside of VCS control)

You need to stop the applications that use CVM volumes or CFS mount points not controlled by VCS.

To stop the applications that use CVM or CFS (outside of VCS control)

- 1 Using native application commands, stop the applications that use CVM or a CFS.
- 2 Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

Unmounting CFS file systems (outside of VCS control)

You need to unmount CFS file systems that are not under VCS control on all nodes.

To unmount CFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the mount command. The command lists all the mounted clustered file systems. Consult the main.cf file for identifying the files that are under VCS control.

```
# mount | grep vxfs | grep cluster
```

- 2 Unmount each file system that is not controlled by VCS:

```
# umount mount_point
```

Stopping VCS

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

The process also stops replication as the replication service group is also taken offline.

To stop VCS

- 1 Log in as the superuser on one of the cluster nodes.
- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped.

In this command output, the VCS engine or high availability daemon (HAD) port `h` is not displayed. This output indicates that VCS has been stopped.

```
# gabconfig -a
```

```
GAB Port Memberships
```

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

Stopping the applications that use VxVM or VxFS (outside of VCS control)

You need to stop all applications that use VxVM volumes or VxFS mount points not under VCS control.

To stop the applications that use VxVM or VxFS (outside of VCS control)

- 1 Using native application commands, stop the applications that use VxVM or VxFS.
- 2 Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

Unmounting VxFS file systems (outside of VCS control)

You need to unmount VxFS file systems that are not under VCS control on all nodes.

Note: To avoid issues on rebooting, you must remove all entries of VxFS from the `/etc/fstab` file.

To unmount VxFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the mount command. The command lists all the mounted file systems.

```
# mount | grep vxfs
```

- 2 Unmount each file system that is not under VCS control:

```
# umount mount_point
```

Uninstalling SF Oracle RAC from a cluster

This chapter includes the following topics:

- [Uninstalling SF Oracle RAC from a cluster](#)
- [Uninstalling SF Oracle RAC with the Veritas Web-based installer](#)
- [Uninstalling SF Oracle RAC using a response file](#)

Uninstalling SF Oracle RAC from a cluster

You can remove the SF Oracle RAC RPMs from all nodes in the SF Oracle RAC cluster using the `uninstallsrac` program. The `uninstallsrac` program can be accessed from the product disc or from the `/opt/VRTS/install` directory.

Perform the steps in the following procedure to remove SF Oracle RAC from a cluster.

To remove SF Oracle RAC from a cluster

- 1 Remove the CP server configuration using the removal script.
See [“Removing the CP server configuration using the removal script”](#) on page 573.
- 2 Remove the SF Oracle RAC RPMs. You can remove the RPMs using the uninstallation program or using the response file.

Using the uninstallation program:

See [“Removing the SF Oracle RAC RPMs”](#) on page 577.

Using the response file:

See [“Uninstalling SF Oracle RAC using a response file”](#) on page 581.

- 3 Remove other configuration files (optional).

See [“Removing other configuration files \(optional\)”](#) on page 579.

- 4 Remove the Storage Foundation for Databases (SFDB) repository after removing the product.

See [“Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product”](#) on page 579.

- 5 Reboot the nodes.

```
# shutdown -r now
```

Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or a cluster that hosts the CP server.

Warning: Ensure that no SF Oracle RAC cluster (application cluster) uses the CP server that you want to unconfigure.

You can use the CP server configuration utility (`configure_cps.pl`) to remove the CP server configuration. This utility performs the following tasks when you choose to unconfigure the CP server:

- Removes all CP server configuration files
- Removes the VCS configuration for CP server

After you run this utility, you can uninstall VCS from the node or the cluster.

Note: You must run the configuration utility only once per CP server (which can be on a single-node VCS cluster or an SFHA cluster), when you want to remove the CP server configuration.

To remove the CP server configuration

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 Select option 3 from the menu to unconfigure the CP server.

```
VERITAS COORDINATION POINT SERVER CONFIGURATION UTILITY
=====
```

Select one of the following:

```
[1] Configure Coordination Point Server on single node VCS system
```

```
[2] Configure Coordination Point Server on SFHA cluster
```

```
[3] Unconfigure Coordination Point Server
```

- 3 Review the warning message and confirm that you want to unconfigure the CP server.

```
WARNING: Unconfiguring Coordination Point Server stops the
vxcpsserv process. VCS clusters using this server for
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)
(Default:n) :y
```

- 4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration

- 5 Answer **y** to delete the CP server database.

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 6 Answer **y** at the prompt to confirm the deletion of the CP server database.

```
Warning: This database won't be available if CP server
is reconfigured on the cluster. Are you sure you want to
proceed with the deletion of database? (y/n) (Default:n) :
```

- 7 Answer **y** to delete the CP server configuration file and log files.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the `hagrp -state` command to ensure that the CPSSG service group has been removed from the node. For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG
```

```
VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```

Manually remove the CP server fencing configuration

The following procedure describes how to manually remove the CP server fencing configuration from the CP server. This procedure is performed as part of the process to stop and remove server-based IO fencing.

Note: This procedure must be performed after the SF Oracle RAC cluster has been stopped, but before the SF Oracle RAC cluster software is uninstalled.

This procedure is required so that the CP server database can be reused in the future for configuring server-based fencing on the same SF Oracle RAC cluster(s).

Perform the steps in the following procedure to manually remove the CP server fencing configuration.

Note: The `cpsadm` command is used in the following procedure. For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

To manually remove the CP server fencing configuration

- 1 Unregister all SF Oracle RAC cluster nodes from all CP servers using the following command:

```
# cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
```

- 2 Remove the SF Oracle RAC cluster from all CP servers using the following command:

```
# cpsadm -s cp_server -a rm_clus -u uuid
```

- 3 Remove all the SF Oracle RAC cluster users communicating to CP servers from all the CP servers using the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

- 4 Proceed to uninstall the SF Oracle RAC cluster software.

Manually deleting cluster details from a CP server

You can manually delete the cluster details from a coordination point server (CP server) using the following procedure.

To manually delete cluster details from a CP server

- 1 List the nodes in the CP server cluster:

```
# cpsadm -s mycps1 -a list_nodes
```

ClusterName	UUID	Hostname(Node ID)	Registered
=====	=====	=====	=====
cluster1	{3719a60a-1dd2-11b2-b8dc-197f8305ffc0}	node0 (0)	1

- 2 List the CP server users:

```
# cpsadm -s mycps1 -a list_users
```

Username/Domain Type	Cluster Name/UUID	Role
=====	=====	=====
cpsclient@hostname/vx	cluster1/{3719a60a-1dd2-11b2-b8dc-197f8305ffc0}	Operator

- 3 Remove the privileges for each user of the cluster that is listed in step 2 from the CP server cluster. For example:

```
# cpsadm -s mycps1 -a rm_clus_from_user
-c cluster1 -e cpsclient@hostname -g vx -f cps_operator
Cluster successfully deleted from user cpsclient@hostname privileges.
```

- 4 Remove each user of the cluster that is listed in step 2. For example:

```
# cpsadm -s mycps1 -a rm_user -e cpsclient@hostname -g vx
User cpsclient@hostname successfully deleted
```

- 5 Unregister each node that is registered to the CP server cluster. See the output of step 1 for registered nodes. For example:

```
# cpsadm -s mycps1 -a unreg_node -c cluster1 -n 0
Node 0 (node0) successfully unregistered
```

- 6 Remove each node from the CP server cluster. For example:

```
# cpsadm -s mycps1 -a rm_node -c cluster1 -n 0
Node 0 (node0) successfully deleted
```

- 7 Remove the cluster.

```
# cpsadm -s mycps1 -a rm_clus -c cluster1
Cluster cluster1 deleted successfully
```

- 8 Verify that the cluster details are removed successfully.

```
# cpsadm -s mycps1 -a list_nodes
```

```
ClusterName      UUID              Hostname(Node ID) Registered
=====
=====
```

```
# cpsadm -s mycps1 -a list_users
```

```
Username/Domain Type Cluster Name/UUID      Role
=====
=====
```

Removing the SF Oracle RAC RPMs

The `uninstallsrac` program can remove these RPMs only if there are no open volumes.

The installer performs the following tasks:

- Removes the SF Oracle RAC RPMs.
- Removes the language RPMs, if installed.

Note: The following directories remain after uninstallation: `/opt/VRTS`, `/opt/VRTSperl`, `/etc/VRTSvcs`, `/var/VRTSvcs`, `/var/VRTSat_lhc`, `/var/VRTSat`. They contain logs and configuration information for future reference. You may or may not remove them.

To remove the SF Oracle RAC RPMs

- 1 Log in as the superuser on any node in the cluster.
- 2 Navigate to the directory that contains the `uninstallsfrac` program:

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

- 3 Start the `uninstallsfrac` program:

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

The program displays the directory where the logs are created and the copyright message.

- 4 Confirm to uninstall SF Oracle RAC.

```
All SF Oracle RAC processes that are currently running
must be stopped.
Do you want to stop SF Oracle RAC processes now?
[y,n,q,?] (y) y
```

Note: If you have not already unmounted the VxFS file systems, the installer will display a message asking that the file systems be unmounted. Make sure that you unmount the file systems before you proceed.

The program performs the following tasks:

- Stops the agents and performs verifications on each node to proceed with uninstallation
- Stops the SF Oracle RAC processes and uninstalls the SF Oracle RAC RPMs
- Displays the location of the uninstallation summary, response file, and log files for reference.

Removing other configuration files (optional)

You can remove the Veritas configuration files and the RPMs that are left after running the `uninstallsrac` program.

To remove residual Veritas configuration files (optional)

- 1 List all VRTS RPMs that can be removed.

```
# rpm -qa |grep -i vrts
```

- 2 Run the `rpm -e rpm_name` command to remove the remaining VRTS RPMs.

- 3 Move the residual Veritas configuration files to the `vrts.bkp` directory:

```
# cd /var
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# mv vx vrts.bkp
# cd /var/opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /etc
# mkdir vrts.bkp
# mv vx *llt* *fen* *gab* *vcs* vcsmmtab vrts.bkp
```

You can remove the `vrts.bkp` directories at a later time.

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

- 1 Identify the SFDB repositories created on the host.

```
# cat /var/vx/vxdba/rep_loc

{
  "sfae_rept_version" : 1,
  "oracle" : {
    "SFAEDB" : {
      "location" : "/data/sfaedb/.sfae",
      "old_location" : "",
      "alias" : [
        "sfaedb"
      ]
    }
  }
}
```

- 2 Remove the directory identified by the `location` key.

```
# rm -rf /data/sfaedb/.sfae
```

- 3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Uninstalling SF Oracle RAC with the Veritas Web-based installer

This section describes how to uninstall using the Veritas Web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout Version in SF Oracle RAC 6.0 with a previous version of SF Oracle RAC.

To uninstall SF Oracle RAC

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
 See [“Starting the Veritas Web-based installer”](#) on page 86.
- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 4 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list, and click **Next**.
- 5 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.
- 6 After the validation completes successfully, click **Next** to uninstall SF Oracle RAC on the selected system.
- 7 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 8 After the installer stops the processes, the installer removes the products from the specified system.
 Click **Next**.
- 9 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 10 Click **Finish**.
 You see a prompt recommending that you reboot the system, and then return to the Web page to complete additional tasks.

Uninstalling SF Oracle RAC using a response file

Perform the steps in the following procedure to uninstall SF Oracle RAC using a response file.

To uninstall SF Oracle RAC using a response file

- 1 Make sure that you have completed the pre-uninstallation tasks.

- 2 Create a response file using one of the available options.

For information on various options available for creating a response file:

See [“About response files”](#) on page 373.

Note: You must replace the host names in the response file with that of the systems from which you want to uninstall SF Oracle RAC.

For a sample response file:

See [“Sample response file for uninstalling SF Oracle RAC”](#) on page 583.

- 3 Navigate to the directory containing the SF Oracle RAC uninstallation program:

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

- 4 Start the uninstallation:

```
# ./uninstallsfrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

- 5 Remove other Veritas configuration files and RPMs that may be present.

- 6 Reboot the nodes:

```
# shutdown -r now
```

- 7 Optionally, remove residual configuration files, if any.

See [“Removing other configuration files \(optional\)”](#) on page 579.

Response file variables to uninstall SF Oracle RAC

[Table 34-1](#) lists the response file variables that you can define to uninstall SF Oracle RAC.

Table 34-1 Response file variables specific to uninstalling SF Oracle RAC

Variable	List or Scalar	Description
CFG{opt}{uninstall}	Scalar	Uninstalls SF Oracle RAC RPMs. (Required)
CFG{systems}	List	List of systems on which the product is to be uninstalled. (Required)

Table 34-1 Response file variables specific to uninstalling SF Oracle RAC
(continued)

Variable	List or Scalar	Description
CFG{prod}	Scalar	Defines the product to be uninstalled. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is <i>/opt/VRTS/install/logs</i> . Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)

Sample response file for uninstalling SF Oracle RAC

The following sample response file uninstalls SF Oracle RAC from nodes, galaxy and nebula.

```
our %CFG;  
  
$CFG{opt}{uninstall}=1;  
$CFG{prod}="SFRAC60";  
$CFG{systems}=[ qw(galaxy nebula) ];  
  
1;
```

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Tunable files for installation](#)
- [Appendix C. SF Oracle RAC installation and configuration sample values](#)
- [Appendix D. Sample configuration files](#)
- [Appendix E. Setting up inter-system communication](#)
- [Appendix F. Automatic Storage Management](#)
- [Appendix G. Creating a test database](#)
- [Appendix H. High availability agent information](#)
- [Appendix I. SF Oracle RAC deployment scenarios](#)
- [Appendix J. Compatibility issues when installing Veritas Storage Foundation for Oracle RAC with other products](#)

Installation scripts

This appendix includes the following topics:

- [About Veritas Storage Foundation and High Availability Solutions installation scripts](#)
- [Installation script options](#)
- [About using the postcheck option](#)

About Veritas Storage Foundation and High Availability Solutions installation scripts

Veritas Storage Foundation and High Availability Solutions products 6.0 provides several installation scripts. You can find these scripts at the root of the product media in the scripts directory.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from the Symantec download site, which does not include the installer, use the appropriate product installation script.

The following product installation scripts are available:

Table A-1 Product installation scripts

Product installation script	Veritas product name
<code>installvcs</code>	Veritas Cluster Server (VCS)
<code>installsf</code>	Veritas Storage Foundation (SF)
<code>installsfha</code>	Veritas Storage Foundation and High Availability (SFHA)

Table A-1 Product installation scripts (*continued*)

Product installation script	Veritas product name
<code>installsfcfsha</code>	Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
<code>installsffrac</code>	Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
<code>installsfbasece</code>	Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)
<code>installvm</code>	Veritas Volume Manager
<code>installfs</code>	Veritas File System
<code>installdmp</code>	Veritas Dynamic Multi-Pathing
<code>installsvs</code>	Symantec VirtualStore

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Installation script options

[Table A-2](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 585.

Table A-2 Available command line options

Command Line Option	Function
<code>system1 system2...</code>	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>-addnode</code>	Adds a node to a high availability cluster.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-allpkgs</code>	Displays all RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
<code>-comcleanup</code>	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
<code>-configure</code>	Configures the product after installation.
<code>-fencing</code>	Configures I/O fencing in a running cluster.
<code>-hostfile full_path_to_file</code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-install</code>	The <code>-install</code> option is used to install products on systems.
<code>-installallpkgs</code>	Specifies that all RPMs are installed.
<code>-installminpkgs</code>	Specifies that the minimum RPM set is installed.
<code>-installrecpkgs</code>	Specifies that the required RPM set is installed.
<code>-keyfile ssh_key_file</code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
<code>-kickstart dir_path</code>	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The <code>dir_path</code> indicates the path to the directory in which to create the file.
<code>-license</code>	Registers or updates product licenses on the specified systems.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-logpath log_path</code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-makeresponsefile</code>	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
<code>-minpkgs</code>	Displays the minimal RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-nolic</code>	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
<code>-pkginfo</code>	Displays a list of RPMs and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS RPMs.
<code>-pkgpath package_path</code>	Designates the path of a directory that contains all RPMs to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
<code>-pkgset</code>	Discovers and displays the RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's RPMs in correct installation order by group.
<code>-postcheck</code>	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
<code>-recpkgs</code>	Displays the recommended RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required RPMs and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-rolling_upgrade</code>	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-rsh</code>	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
<code>-serial</code>	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
<code>-settunables</code>	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
<code>-start</code>	Starts the daemons and processes for the specified product.
<code>-stop</code>	Stops the daemons and processes for the specified product.
<code>-tmppath <i>tmp_path</i></code>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.
<code>-uninstall</code>	The <code>-uninstall</code> option is used to uninstall products from systems.
<code>-tunablesfile</code>	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
<code>-upgrade</code>	Specifies that an existing version of the product exists and you plan to upgrade it.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>--version</code>	Checks and reports the installed products and their versions. Identifies the installed and missing RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
<code>-yumgroupxml</code>	The <code>-yumgroupxml</code> option is used to generate a yum group definition XML file. The <code>createrepo</code> command can use the file on Redhat Linux to create a yum group for automated installation of all rpms for a product. An available location to store the XML file should be specified as a complete path. The <code>-yumgroupxml</code> option is supported on Redhat Linux only.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists.
- The `VRTSIlt` pkg version is not consistent on the nodes.
- The `Ilt-linkinstall` value is incorrect.
- The `Ilthosts(4)` or `Ilttab(4)` configuration is incorrect.
- the `/etc/gabtab` file is incorrect.

- The incorrect GAB linkinstall value exists.
- The VRTSgab pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The VRTSvcs pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg` file is invalid.
- The vxfen link-install value is incorrect.
- The VRTSvxfen pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `Volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the `Autostartlist` value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because Vxfen is not started.
- Cluster Volume Manager cannot start because gab is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

See [“Performing a postcheck on a node”](#) on page 170.

Tunable files for installation

This appendix includes the following topics:

- [About setting tunable parameters using the installer or a response file](#)
- [Setting tunables for an installation, configuration, or upgrade](#)
- [Setting tunables with no other installer-related operations](#)
- [Setting tunables with an un-integrated response file](#)
- [Preparing the tunables file](#)
- [Setting parameters for the tunables file](#)
- [Tunables value parameter definitions](#)

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

```
# ./installer -tunablesfile tunables_file_name
```

See [“Setting tunables for an installation, configuration, or upgrade”](#) on page 594.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [  
system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 595.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

See [“Setting tunables with an un-integrated response file”](#) on page 596.

For more information on response files, see the *chapter: About response files*.

You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 598.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 598.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 597.
- 2 Make sure the systems where you want to install SF Oracle RAC meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 598.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 597.
- 2 Make sure the systems where you want to install SF Oracle RAC meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 598.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

- 1 Make sure the systems where you want to install SF Oracle RAC meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 597.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-settunables` option.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name -settunables
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

For more information on response files, see the *chapter: About response files*.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

Format the tunable parameter as follows:

```
$TUN{"tunable_name"}{"system_name"|"*"}=value_of_tunable;
```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#  
# Tunable Parameter Values:  
#  
our %TUN;  
  
$TUN{"tunable1"}{"*"}=1024;  
$TUN{"tunable3"}{"sys123"}="SHA256";  
  
1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See [“Tunables value parameter definitions”](#) on page 598.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the `dmp_daemon_count` value from its default of 10 to 16. You can use the wildcard symbol `"**"` for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"**"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table B-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table B-1 Supported tunable parameters

Tunable	Description
<code>dmp_cache_open</code>	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
<code>dmp_daemon_count</code>	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
<code>dmp_delayq_interval</code>	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
<code>dmp_fast_recovery</code>	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
<code>dmp_health_time</code>	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
<code>dmp_log_level</code>	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer sets only the system default value of read_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer sets only the system default value of read_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires system reboot to take effect.

Table B-1 Supported tunable parameters (*continued*)

Tunable	Description
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

SF Oracle RAC installation and configuration sample values

This appendix includes the following topics:

- [About the installation and configuration worksheets](#)
- [SF Oracle RAC worksheet](#)
- [Oracle RAC worksheet](#)
- [Replicated cluster using VVR worksheet](#)
- [Replicated cluster using SRDF worksheet](#)
- [Required installation information for Oracle Clusterware/Grid Infrastructure](#)
- [Required installation information for Oracle database](#)

About the installation and configuration worksheets

The installation programs prompt you for information during the installation and configuration of SF Oracle RAC and Oracle RAC. The installation and configuration worksheets provide sample values that you can use as examples of the information required when you run the Veritas installation programs or the Oracle Universal Installer.

Symantec recommends using the worksheets to record values for your systems before you begin the installation and configuration process.

SF Oracle RAC worksheet

This section provides worksheets for installing and configuring SF Oracle RAC, its component products, and features.

[Table C-1](#) contains the sample values that may be used when you install and configure SF Oracle RAC. Enter the SF Oracle RAC values for your systems in the following table:

Table C-1 SF Oracle RAC worksheet

Installation information	Sample value	Assigned value
Number of nodes in the cluster	2	
Host names for Primary cluster	galaxy and nebula Note: Do not use the underscore character in host names. Host names that use the underscore character are not compliant with RFC standards and cause issues.	
Host names for added or removed node	saturn Note: Do not use the underscore character in host names. Host names that use the underscore character are not compliant with RFC standards and cause issues.	
License keys	License keys can be one of the following types: - Valid license keys for each system in the cluster - Valid site license key - Valid demo license key If you want to configure global clusters and enable disaster recovery, you must enter appropriate license keys. Note: You can choose between SF Oracle RAC and SF Oracle RAC Disaster Recovery and High Availability options for license keys.	

Table C-1 SF Oracle RAC worksheet (*continued*)

Installation information	Sample value	Assigned value
SF Oracle RAC RPMs to be installed on the cluster	Select one of the following options: ■ Install minimal required Veritas Storage Foundation for Oracle RAC RPMs ■ Install recommended Veritas Storage Foundation for Oracle RAC RPMs ■ Install all Veritas Storage Foundation for Oracle RAC RPMs Install only the required RPMs if you do not want to configure any optional components or features. Default option is to install the recommended RPMs.	
Primary cluster name	rac_cluster101	
Primary cluster ID number	101	
Private network links	eth1, eth2	
Cluster Manager NIC (Primary NIC)	eth0	
Cluster Manager IP	10.10.12.1, 10.10.12.2	
Netmask for the virtual IP address	255.255.240.0	
VCS user name (not required if you configure your cluster in secure mode)	VCS usernames must not exceed 1024 characters. Example: smith	
VCS user password	VCS passwords must not exceed 512 characters.	
VCS user privileges	Users have three levels of privileges: A=Administrator, O=Operator, or G=Guest. Example: A	
Domain-based address of SMTP server	smtp.symantecexample.com	

Table C-1 SF Oracle RAC worksheet (*continued*)

Installation information	Sample value	Assigned value
Email address of SMTP notification recipients	admin@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: I The severity levels are defined as follows: ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption	
SNMP trap daemon port number the console	162	
System name for the SNMP console	system2	

Table C-1 SF Oracle RAC worksheet (*continued*)

Installation information	Sample value	Assigned value
Minimum severity of events for SNMP trap notification	Events have four levels of severity: ■ I=Information ■ W=Warning ■ E=Error ■ S=SevereError Example: I The severity levels are defined as follows: ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption	
Vxfen disks These values are required if SCSI-3 disks are used as coordination points for your configuration.	/dev/sdc, /dev/sdd, /dev/sde	
Vxfen disk group	vxfencoordg	

Veritas Cluster Server component information

[Table C-2](#) displays the information that is required to configure the Veritas Cluster Server component.

Table C-2 Veritas Cluster Server component information

Information	Example	Assigned values
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 that contains 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 that are associated with each NIC for each network link must be the same on all nodes. For example: - eth1 - eth2 Do not use the network interface card that is used for the public network, which is typically eth0.	

I/O fencing information

[Table C-3](#) displays the information that is required to configure I/O fencing.

Table C-3 I/O fencing information

Information	Sample values	Assigned values
The name of three disks that form the coordinator disk group.	The following are examples of disk names: ■ /dev/sdc ■ /dev/sdd ■ /dev/sde	
The names for each disk in the coordinator disk group (if using DMP).	The following are examples: ■ /dev/vx/rdmp/sdc3 ■ /dev/vx/rdmp/sdd3 ■ /dev/vx/rdmp/sde3	

SF Oracle RAC add user information

[Table C-4](#) displays the information that is required to add VCS users. If you configure SF Oracle RAC cluster in secure mode, you need to add VCS users.

Note: Adding VCS users is optional.

Table C-4 SF Oracle RAC add user information

Information	Examples	Assigned values
User name	smith	
User password	*****	

Table C-4 SF Oracle RAC add user information (*continued*)

Information	Examples	Assigned values
User privilege	Users have three levels of privileges: ■ A=Administrator ■ O=Operator ■ G=Guest Example: A VCS privilege levels include: ■ Administrators— Can perform all operations, including configuration options on the cluster, service groups, systems, resources, and users. ■ Operators—Can perform specific operations on a cluster or a service group. ■ Guests—Can view specified objects.	

Global cluster information

[Table C-5](#) displays the information that is required to configure global clusters.

Note: Global clusters are an optional feature that requires a license.

Table C-5 Global cluster information

Information	Example	Assigned values
Name of the public NIC	You must specify appropriate values for the NIC when you are prompted. Example: <code>eth0</code>	
Virtual IP address of the NIC	You must specify appropriate values for the virtual IP address when you are prompted. Example: <code>10.10.12.1</code>	

Table C-5 Global cluster information (*continued*)

Information	Example	Assigned values
Netmask for the virtual IP address	You must specify appropriate values for the netmask when you are prompted. Example: 255.255.255.0	

Oracle RAC worksheet

This section provides a worksheet with sample values for installing and configuring Oracle RAC.

Note: *Italicized* text in parenthesis indicates the corresponding variable that is used in procedures. When you perform the steps in the procedures, ensure that you replace the variables with the values that are assigned to them in this worksheet.

[Table C-6](#) displays the sample worksheets that may be used as reference when you perform the corresponding tasks.

Table C-6 Required information for Oracle RAC

Sample value sheet	Go to
Oracle user and group	See Table C-7 on page 614.
Public IP addresses and host names	See Table C-8 on page 615.
PrivNIC and MultiPrivNIC	See Table C-9 on page 615.
Oracle RAC home directories	See Table C-10 on page 621.
OCR and voting disk	See Table C-11 on page 623.
CSSD and Oracle database configuration	See Table C-12 on page 624.

[Table C-7](#) displays sample values that may be used when you create Oracle users and groups.

Table C-7 Sample value sheet - Oracle user and group

Information	Sample value	Assigned value
Oracle user name (<i>user_name</i>)	- For Oracle Clusterware oracle - For Oracle Grid Infrastructure grid	
Oracle user ID (<i>user_id</i>)	1000	
Oracle group name - Primary group (<i>grp_name</i>)	oinstall (for inventory group as primary group)	
Oracle group name - Secondary group (<i>grp_name_sec</i>)	dba (for dba group as secondary group)	
Oracle group ID - Primary group ID (<i>grp_id</i>)	1000 (for inventory group as primary group)	
Oracle group ID - Secondary group ID (<i>grp_id_sec</i>)	1001 (for dba group as secondary group)	
Oracle user home directory (<i>usr_home_ora</i>)	/home/oracle	
Grid user home directory (<i>usr_home_grid</i>)	/home/grid	

Table C-8 displays sample values for public IP addresses and host names.

Table C-8 Sample value sheet - Public IP addresses and host names

Information	Sample value	Assigned value
Name of a node in the cluster (<i>node_name1</i>) (<i>node_name2</i>)	For a two-node cluster: ▪ galaxy ▪ nebula	
Name of a new node added to the cluster (<i>nodenew_name</i>)	saturn	
Virtual IP address	10.10.10.10	
Virtual IP alias	galaxy-vip	
SCAN IP addresses (only for Oracle RAC 11g Release 2)	A minimum of three addresses is recommended. 10.10.10.20 10.10.10.21 10.10.10.22	
SCAN name (only for Oracle RAC 11g Release 2)	rac_cluster101_scan	

[Table C-9](#) displays sample values that may be used when you configure PrivNIC or MultiPrivNIC.

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC

Information	Sample value	Assigned value
Private IP address for first node in the cluster - PrivNIC (<i>privnic_ip_node1</i>)	192.168.12.1 (on galaxy) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula.	

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC (*continued*)

Information	Sample value	Assigned value
Private IP address for second node in the cluster - PrivNIC <i>(privnic_ip_node2)</i>	192.168.12.2 (on nebula) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula.	
Private IP address when you add a node in the cluster - PrivNIC <i>(privnic_ip_newnode)</i>	192.168.12.5 (on nebula) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula.	

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC (*continued*)

Information	Sample value	Assigned value
NIC address for network on first node in the cluster (<i>nic1_node1</i>) (<i>nic2_node1</i>)	You have to choose an LLT device as a device for the Oracle Clusterware heartbeat. The interfaces specified should be exactly the same in name and total number as those which have been used for LLT configuration. For example, if the LLT devices on galaxy are eth1, eth2: eth1, eth2 Then the PrivNIC or MultiPrivNIC device names will be as follows: Device@galaxy= { eth1 = 0, eth2 = 1 } If aggregated device names are configured under LLT, then the aggregated names must be used in PrivNIC or MultiPrivNIC agent. Note: If you configured aggregated interfaces for LLT, then you must set the Device attribute value to use the same aggregated interface names that you configured for LLT. For example, if LLT device name on galaxy is: aggr1 Then the Device Attribute for the MultiPrivNIC agent would be as follows: Device@galaxy = { aggr1 = 0 }	

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC (*continued*)

Information	Sample value	Assigned value
NIC address for network on second node in the cluster (<i>nic1_node2</i>) (<i>nic2_node2</i>)	You have to choose an LLT device as a device for the Oracle Clusterware heartbeat. The interfaces specified should be exactly the same in name and total number as those which have been used for LLT configuration. For example, if the LLT devices on nebula are eth1, eth2: eth1, eth2 (on nebula) Then the PrivNIC or MultiPrivNIC device names will be as follows: Device@nebula= { eth1 = 0, eth2 = 1 } If aggregated device names are configured under LLT, then the aggregated names must be used in PrivNIC or MultiPrivNIC agent. Note: If you configured aggregated interfaces for LLT, then you must set the Device attribute value to use the same aggregated interface names that you configured for LLT. For example, if LLT device name on nebula is: aggr1 (on nebula) Then the Device Attribute for the MultiPrivNIC agent would be as follows: Device@nebula = { aggr1 = 0 }	
VCS resource name for PrivNIC (<i>priv_resname</i>)	ora_priv	
VCS resource name for MultiPrivNIC (<i>multipriv_resname</i>)	multi_priv	

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC (*continued*)

Information	Sample value	Assigned value
Private IP addresses for first node in the cluster - MultiPrivNIC <i>(multipriv_ip1_node1)</i> <i>(multipriv_ip2_node1)</i> <i>(multipriv_ip3_node1)</i>	On galaxy: ■ 192.168.12.1 (Oracle Clusterware) ■ 192.168.2.1 (UDP to use for the CLUSTER_INTERCONNECT parameter) ■ 192.168.3.1 (Second UDP to use for the CLUSTER_INTERCONNECT parameter) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X.0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula. Symantec recommends that all Oracle Clusterware and UDP cache-fusion links be LLT links.	
Private IP addresses for second node in the cluster - MultiPrivNIC <i>(multipriv_ip1_node2)</i> <i>(multipriv_ip2_node2)</i> <i>(multipriv_ip3_node2)</i>	On nebula: ■ 192.168.12.2 (Oracle Clusterware) ■ 192.168.2.2 (UDP to use for the CLUSTER_INTERCONNECT parameter) ■ 192.168.3.2 (Second UDP to use for the CLUSTER_INTERCONNECT parameter) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X.0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula. Symantec recommends that all Oracle Clusterware and UDP cache-fusion links be LLT links.	

Table C-9 Sample value sheet - PrivNIC and MultiPrivNIC (*continued*)

Information	Sample value	Assigned value
Private IP addresses when you add a new node in the cluster - MultiPrivNIC <i>(multipriv_ip1_newnode)</i> <i>(multipriv_ip2_newnode)</i>	On nebula: ■ 192.168.12.5 (Oracle Clusterware) ■ 192.168.2.6 (UDP to use for the CLUSTER_INTERCONNECT parameter) Note: The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X.0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula. Symantec recommends that all Oracle Clusterware and UDP cache-fusion links be LLT links.	
Private hostnames (set in /etc/hosts) for first node in the cluster <i>(ip_node1)</i>	On galaxy: ■ galaxy-priv ■ galaxy-priv1	
Private hostnames (set in /etc/hosts) for second node in the cluster <i>(ip_node2)</i>	On nebula: ■ nebula-priv ■ nebula-priv1	
Netmask for cluster <i>(netmask_ip)</i>	255.255.255.0	

Table C-10 displays sample values that may be used when you create the Oracle RAC home directories.

Table C-10 Sample value sheet - Oracle RAC home directories

Information	Sample value	Assigned value
Disk for each node that contains the Oracle Clusterware and database binaries <i>disk_name</i>	Disk_1	
VxVM local disk group name <i>dg_name</i>	■ bindg_galaxy (on galaxy) ■ bindg_nebula (on nebula)	
CVM disk group name <i>cvm_dg</i>	bindg	
Volume name for Oracle Clusterware binaries <i>clus_volname</i>	■ For Oracle Clusterware crsbinvol ■ For Oracle Grid Infrastructure gridbinvol	
Oracle Clusterware home directory <i>clus_home</i>	■ For Oracle Clusterware, path to <i>crs_home</i> /u01/app/oracle/product/10.2.0/crshome ■ For Oracle Grid Infrastructure, path to <i>grid_home</i> /u01/app/grid/product/11.2.0/gridhome	
Volume name for Oracle Database binaries <i>ora_volname</i>	orabinvol	
VCS resource name for disk groups containing the Oracle Clusterware/Grid Infrastructure and Oracle Database home directories <i>dg_resname</i>	bin_dg	

Table C-10 Sample value sheet - Oracle RAC home directories (*continued*)

Information	Sample value	Assigned value
Oracle Clusterware binary mount point resource name <i>clusbin_mnt_resname</i>	- For Oracle Clusterware crsbin_mnt - For Oracle Grid Infrastructure gridbin_mnt	
Oracle Database binary mount point resource name <i>orabin_mnt_resname</i>	orabin_mnt	
Oracle base directory for Oracle Clusterware <i>oracle_base</i>	/u01/app/oracle	
Oracle base directory for Oracle Grid Infrastructure <i>grid_base</i>	/u01/app/grid	
Oracle base directory for Database <i>oracle_base</i>	/u01/app/oracle	
Oracle Database home directory <i>db_home</i>	/u02/app/oracle/product/11.2.0/dbhome_1	
Oracle Grid Infrastructure home directory <i>grid_home</i>	/u01/app/11.2.0/grid	

Table C-11 displays sample values that may be used when you create the storage for OCR and voting disk.

Table C-11 Sample value sheet - OCR and voting disk

Information	Sample value	Assigned value
Disks for creating a shared disk group for OCR and voting disk <i>disk_name2</i> <i>disk_name3</i>	Disk_2 Disk_3	
Shared disk group for OCR and voting disk <i>ocrvote_dgname</i>	ocrvotedg	
OCR volume on CVM raw volumes <i>ocr_volname</i>	ocrvol	
Voting disk volume on CVM raw volumes <i>vote_volname</i>	votevol	
OCR and voting disk volume on CFS <i>ocrvote_volname</i>	ocrvotevol	
Volume options for OCR and voting disk	nmirror=2	
File system on shared volume (CFS)	/dev/vx/rdsk/ocrvotedg/ocrvotevol	
Mount point for shared file system <i>ocrvote_mnt</i>	/ocrvote	
CVMVolDg resource name for OCR and voting disk <i>ocrvotevol_resname</i>	ocrvote_voldg_ocrvotedg	
CFSMount resource name for OCR and voting disk <i>ocrvotemnt_resname</i>	ocrvote_mnt_ocrvotedg	

Table C-11 Sample value sheet - OCR and voting disk (*continued*)

Information	Sample value	Assigned value
Mount point for archive logs	/oradata	
CVM group name <i>cvm_grpname</i>	cvm	

[Table C-12](#) displays sample values that may be used when you configure the CSSD agent and the Oracle database.

Table C-12 Sample value sheet - CSSD and Oracle database configuration

Information	Sample value	Assigned value
CVM group name <i>cvm_grpname</i>	cvm	
CSSD group name <i>cssd_grpname</i>	cssd	
VCS service group for Oracle Database <i>oradb_grpname</i>	oradb_grp	
VCS resource name for Oracle <i>db_resname</i>	oradb	
Oracle database name <i>db_name</i>	db	
Global Database Name	db.symantecexample.com (database name.domain)	
Disk group name for Oracle Database <i>oradb_dgname</i>	oradatadg (Shared)	
Volume name for Oracle Database <i>oradb_volname</i>	oradatavol	

Table C-12 Sample value sheet - CSSD and Oracle database configuration
(continued)

Information	Sample value	Assigned value
Mount point for Oracle Database <i>oradb_mnt</i>	/oradata (Shared)	
CVMVolDg resource for Oracle Database <i>orabdg_resname</i>	oradata_voldg	
CFSMount resource for Oracle Database <i>orabmnt_resname</i>	oradata_mnt	
SID prefix for policy-managed databases <i>oradb_sid_prefix</i>	db	
SID on first node in the cluster <i>oradb_sid_node1</i>	db1	
SID on second node in the cluster <i>oradb_sid_node2</i>	db2	

Replicated cluster using VVR worksheet

[Table C-13](#) contains the sample values that may be used when you install and configure CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the following table:

Table C-13 Replicated cluster using VVR worksheet

Installation information	Sample value	Assigned value
Host names for Secondary Cluster	mercury, jupiter	
Secondary Cluster Name	rac_cluster102	

Table C-13 Replicated cluster using VVR worksheet (*continued*)

Installation information	Sample value	Assigned value
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	rac1_srl	
Resolvable Virtual Hostname of the cluster on Primary Site (for VVR)	rac_clus101_priv	
Resolvable Virtual Hostname of the cluster on Secondary Site (for VVR)	rac_clus102_priv	
Private IP addresses for Secondary Cluster	192.168.12.3 - 192.168.12.4	

Replicated cluster using SRDF worksheet

Table C-14 contains the sample values that may be used when you install and configure CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the following table:

Table C-14 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	rac1_srl	
Resolvable Virtual Hostname of the cluster on Primary Site (for VVR)	rac_clus101_priv	
Resolvable Virtual Hostname of the cluster on Secondary Site (for VVR)	rac_clus102_priv	
Private IP addresses for Secondary Cluster	192.168.12.3 - 192.168.12.4	

Required installation information for Oracle Clusterware/Grid Infrastructure

This section describes the information required by the Oracle Universal Installer for installing the Oracle Clusterware/Grid Infrastructure software.

[Table C-15](#) lists the information required by the Oracle Universal Installer when you install Oracle Grid Infrastructure.

Table C-15 Required installation information for Oracle Grid Infrastructure - Oracle RAC 11g Release 2

OUI menu	Description
Select installation option	Select the option Install and Configure Grid Infrastructure for a Cluster .
Select installation type	Select the option Advanced Installation .
Specify cluster configuration	Enter the SCAN name for the cluster that will be used by the database clients to connect to databases within the cluster.
Grid Plug and Play information	Provide the following information: - Name of the cluster - SCAN name The SCAN address on the domain name server (DNS) must resolve to three addresses (recommended) or at least one address. - SCAN port
Specify network interface usage	Identify the planned use for each interface: Public, Private, or Do Not use. Note: Make sure that the same private interfaces that you specified at the time of configuring PrivNIC and MultiPrivNIC are listed on the screen. Note: Mark the interfaces for the subnet containing the private IP addresses managed by the PrivNIC/MultiPrivNIC agents as 'Private'. The interfaces that are Private are stored in GPnP profile as a 'cluster_interconnect' for Oracle Grid Infrastructure communication and database cache fusion traffic.
Storage option information	Select the option Shared File System .

Table C-15 Required installation information for Oracle Grid Infrastructure - Oracle RAC 11g Release 2 (*continued*)

OUI menu	Description
OCR storage option	Enter the full path of the location where you want to store the OCR information. For example, if you are storing the OCR information on CFS, enter: <code>/ocrvote/ocr</code>. Note: Select the option External Redundancy when you install Oracle Clusterware/Grid Infrastructure. Mirroring is performed by CVM. If you are storing the OCR information on an ASM disk group that uses CVM raw volumes, enter: <code>/dev/vx/rdisk/ocrvotedg/ocrvol</code> Note: Select the option External when you create the ASM disk group and instances. Mirroring is performed by CVM.
Voting Disk storage option	Enter the full path of the location where you want to store the voting disk information. For example, if you are storing the voting disk information on CFS, enter: <code>/ocrvote/vote</code> Note: Select the option External Redundancy when you install Oracle Clusterware/Grid Infrastructure. Mirroring is performed by CVM. If you are storing the voting disk information on an ASM disk group that uses CVM raw volumes, enter: <code>/dev/vx/rdisk/ocrvotedg/votevol</code> Note: Select the option External when you create the ASM disk group and instances. Mirroring is performed by CVM.
Specify installation location	Enter the full path to the Oracle base directory and the Oracle Grid Infrastructure home directory.
Create inventory	Enter the full path to the Oracle inventory directory where you want to store the installation files.

Required installation information for Oracle database

This section describes the information required by the Oracle Universal Installer for installing the Oracle Clusterware/Grid Infrastructure software.

Table C-16 lists the information required by the Oracle Universal Installer when you install Oracle database software.

Table C-16 Required installation information for Oracle database - Oracle RAC 11g Release 2

OUI menu	Description
Select installation option	Select the option Install database software only .
Node selection	Select Real Application Clusters database installation . Select the nodes on which the Oracle RAC database software must be installed.
Select database edition	Select Enterprise Edition .
Specify installation location	Review or enter the ORACLE_BASE and ORACLE_HOME directory paths. The Oracle Universal Installer runs product-specific prerequisite checks. Any items that are flagged must be manually checked and configured.

Sample configuration files

This appendix includes the following topics:

- [About VCS configuration file](#)
- [About the LLT and GAB configuration files](#)
- [About I/O fencing configuration files](#)
- [Sample configuration files](#)

About VCS configuration file

This section provides a high-level overview of the contents of the VCS configuration file after the SF Oracle RAC installation. Review the configuration file after the SF Oracle RAC installation and before the Oracle installation.

The configuration file includes the following information:

- The "include" statements list the various VCS types files (`.conf`) for SF Oracle RAC.
The files are located in the `/etc/VRTSvcs/conf/config` directory. These files define the agents that control the resources in the cluster.
- The cluster definition, with the cluster name provided during installation (for example, `rac_cluster101`), includes the names of users and administrators of the cluster. The `UseFence = SCSI3` attribute is not automatically present; you must manually add it after the installation.
See ["Setting up disk-based I/O fencing manually"](#) on page 149.
- The `main.cf` includes the `cvm` service group. The service group includes definitions for monitoring the CFS and the CVM resources. The `CVMCluster` agent resource definition indicates that the nodes use GAB for messaging operations.

The cvm group has the Parallel attribute set to 1. This value enables the resources to run in parallel on each node in the system list.

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

[Table D-1](#) lists the LLT configuration files and the information that these files contain.

Table D-1 LLT configuration files

File	Description
/etc/sysconfig/llt	This file stores the start and stop environment variables for LLT: ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: 1—Indicates that LLT is enabled to start up. 0—Indicates that LLT is disabled to start up. ■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.
/etc/llthosts	The file <code>llthosts</code> is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster. For example, the file <code>/etc/llthosts</code> contains the entries that resemble: <pre>0 galaxy 1 nebula</pre>

Table D-1 LLT configuration files (*continued*)

File	Description
/etc/llttab	The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the LLT network links that correspond to the specific system. <pre> set-node galaxy set-cluster 2 link eth1 eth1 - ether - - link eth2 eth2 - ether - - </pre> For example, the file <code>/etc/llttab</code> contains the entries that resemble: <pre> set-node galaxy set-cluster 2 link eth1 eth-00:04:23:AC:12:C4 - ether - - link eth2 eth-00:04:23:AC:12:C5 - ether - - </pre> If you use aggregated interfaces, then the file contains the aggregated interface name instead of the <code>eth-MAC_address</code>. The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the <code>link</code> command. These lines identify the two network cards that the LLT protocol uses. If you configured a low priority link under LLT, the file also includes a "link-lowpri" line. Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.

[Table D-2](#) lists the GAB configuration files and the information that these files contain.

Table D-2 GAB configuration files

File	Description
/etc/sysconfig/gab	This file stores the start and stop environment variables for GAB: ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: 1—Indicates that GAB is enabled to start up. 0—Indicates that GAB is disabled to start up. ■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.
/etc/gabtab	After you install SF Oracle RAC, the file /etc/gabtab contains a <code>gabconfig(1)</code> command that configures the GAB driver for use. The file /etc/gabtab contains a line that resembles: <pre>/sbin/gabconfig -c -nN</pre> The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster. Note: Symantec does not recommend the use of the <code>-c -x</code> option for <code>/sbin/gabconfig</code>. Using <code>-c -x</code> can lead to a split-brain condition.

About I/O fencing configuration files

Table D-3 lists the I/O fencing configuration files.

Table D-3 I/O fencing configuration files

File	Description
/etc/sysconfig/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing.

Table D-3 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	This file contains the following parameters: ■ vxfen_mode ■ scsi3—For disk-based fencing ■ customized—For server-based fencing ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ vxfen_mechanism This parameter is applicable only for server-based fencing. Set the value as cps. ■ scsi3_disk_policy ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. ■ raw—Configure the vxfen module to use the underlying raw character devices Note: You must use the same SCSI-3 disk policy on all the nodes. ■ security This parameter is applicable only for server-based fencing. 1—Indicates that communication with the CP server is in secure mode. This setting is the default. 0—Indicates that communication with the CP server is in non-secure mode. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in a server-based fencing can include coordinator disks, CP servers, or a mix of both. If you use coordinator disks, you must create a coordinator disk group with the coordinator disk names. Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ single_cp This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk. ■ autoseed_gab_timeout This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature requires that I/O fencing is enabled. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table D-3 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points. Note: The /etc/vxfentab file is a generated file; do not modify this file. For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows: ■ Raw disk: /dev/sdx /dev/sdy /dev/sdz ■ DMP disk: /dev/vx/rdmp/sdx3 /dev/vx/rdmp/sdy3 /dev/vx/rdmp/sdz3 For server-based fencing, the /etc/vxfentab file also includes the security settings information. For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.

Sample configuration files

SF Oracle RAC provides several sample configuration files illustrating various scenarios. You may use the sample files as a guideline for setting up your cluster environment. These sample files are located at /etc/VRTSvcs/conf/sample_rac/.

This section briefly describes each of the sample files and illustrates the service group configuration for each of them. The section does not include a copy of the main.cf files.

The following sample files are discussed in this section:

- [sfrac02_main.cf](#) file
- [sfrac03_main.cf](#) file
- [sfrac04_main.cf](#) file
- [sfrac05_main.cf](#) file

- [sfrac06_main.cf](#) file
- [sfrac07_main.cf](#) and [sample08_main.cf](#) files
- [sfrac09_main.cf](#) and [sample10_main.cf](#) files
- [sfrac11_main.cf](#) file
- [sfrac12_main.cf](#) and [sample13_main.cf](#) files
- [sfrac14_main.cf](#) file


sfrac02_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.
The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CVM raw volumes.

[Figure D-1](#) illustrates the configuration.

Figure D-1 Service group configuration for sfrac02_main.cf file




sfrac03_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- One IP address (on eth1) is shared by Oracle Clusterware and one of the databases for cache fusion.
The second IP address (on eth2) is used by the second database for cache fusion.
The private IP addresses are managed by the MultiPrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure D-2 illustrates the configuration.

Figure D-2 Service group configuration for sfrac03_main.cf file



sfrac04_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- Database is not managed by VCS. Oracle Clusterware starts, stops, and monitors the databases.

The CRSResource agent monitors the status of the database, the VIP resource, and the listener resource configured under Oracle Clusterware.

Note: The CFSMount and CVMVolDg resources for Oracle database can not be set as critical resources in the group.

The CRSResource agent appears FAULTED until Oracle Clusterware brings up the database.

- The database uses the Oracle UDP IPC for database cache fusion.
- One IP address (on eth1) is shared by Oracle Clusterware and one of the databases for cache fusion.
The second IP address (on eth2) is used by the second database for cache fusion.
The private IP addresses are managed by the MultiPrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

[Figure D-3](#) illustrates the configuration.

Figure D-3 Service group configuration for sfrac04_main.cf file



sfrac05_main.cf file


This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on ASM.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.

- A common IP address is used by Oracle Clusterware and database cache fusion. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure D-4 illustrates the configuration.

Figure D-4 Service group configuration for sfrac05_main.cf file



sfrac06_main.cf file

This sample file describes the following configuration:


- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle. The agent starts, stops, and monitors the database.
- The database uses Oracle UDP IPC for cache fusion. A dedicated link is used for each database .
- One private IP address (on eth1) is used by Oracle Clusterware.

The private IP address on eth2 is used by one of the databases for cache fusion. The private IP address on NIC3 is used by the other database for cache fusion. The private IP addresses are managed by the MultiPrivNIC agent for high availability.

- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure D-5 illustrates the configuration.

Figure D-5 Service group configuration for sfrac06_main.cf file



sfrac07_main.cf and sample08_main.cf files

The sample configuration, `sfrac07_main.cf`, describes a disaster recovery configuration for the primary site. The sample configuration, `sample08_main.cf`, describes a disaster recovery configuration for the secondary site. The configuration uses VVR for replicating data between the sites.


This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.

- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Veritas Volume Replicator (VVR) is used to replicate data between the sites.
- The shared volumes replicated across the sites are configured under the RVG group.
- The replication link used by VVR for communicating log information between sites are configured under the `rlogowner` group. This is a failover group that will be online on only one of the nodes in the cluster at each site.
- The database group will be online on the primary cluster. The `RVGSharedPri` resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.

Figure D-6 illustrates the configuration. The service group configuration is the same on the primary and secondary site. The availability of groups (online/offline) differ between the sites.

Figure D-6 Service group configuration for sfrac07_main.cf and sfrac08_main.cf files



sfrac09_main.cf and sample10_main.cf files

The sample configuration, `sfrac09_main.cf`, describes a disaster recovery configuration for the primary site. The sample configuration, `sample10_main.cf`, describes a disaster recovery configuration for the secondary site. The sample configuration uses EMC SRDF technology for replicating data between the sites.

Note: You can use other supported hardware-based replication technologies with this configuration.

This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.
- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- EMC SRDF is used to replicate data between the sites.
- The SRDF disk groups that are replicated across the sites using SRDF technology and the replication mode are specified under the SRDF resource in the database group. The CVM disk group that comprises the SRDF disk group must be configured under the `CVMVolDg` resource in the database group.
- The database group will be online on the primary cluster. The SRDF resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.

[Figure D-7](#) illustrates the configuration on the primary site.

Figure D-7 Service group configuration for sfrac09_main.cf file





Figure D-8 illustrates the configuration on the secondary site.

Figure D-8 Service group configuration for sfrac10_main.cf file



sfrac11_main.cf file

This sample file describes the following configuration:


- An SF Oracle RAC campus cluster with four nodes hosted across two sites.
- Each site comprises two nodes of the cluster hosting a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- The IP address on eth1 is used by Oracle Clusterware.
The second IP address on NIC2 is used for Oracle database cache fusion.

The private IP addresses are managed by the MultiPrivNIC agent for high availability.

- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Group the hosts at each physical site into separate logical system zones using the SystemZones attribute.

Figure D-9 illustrates the configuration.

Figure D-9 Service group configuration for sfrac11_main.cf file



sfrac12_main.cf and sample13_main.cf files

The sample configuration, sfrac12_main.cf, describes a disaster recovery configuration for the primary site. The sample configuration, sample13_main.cf, describes a disaster recovery configuration for the secondary site with fire-drill capability. The sample configuration uses Hitachi True Copy technology for replicating data between the sites.

Note: You can use other supported hardware-based replication technologies with this configuration.

This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.
- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Hitachi True Copy is used to replicate data between the sites.
- The HTC disk groups that are replicated across the sites using HTC technology and the replication mode are specified under the HTC resource in the database group. The CVM disk group that comprises the HTC disk group must be configured under the CVMVolDg resource in the database group.
- The database group will be online on the primary cluster. The HTC resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.
- The database group `oradb_grp_fd` on the secondary is configured for fire drill.
- When the group `oradb_grp_fd` is brought online, the HTCSnap creates a snapshot of the disk group configured under the HTC resource in the database group `oradg_grp`.
Further, the Oracle database and the associated volumes and mount points configured under the service group `oradb_grp_fd` are brought online using the snapshots created by HTCSnap.

Figure D-10 illustrates the configuration on the primary site.

Figure D-10 Service group configuration for sfrac12_main.cf file





Figure D-11 illustrates the configuration on the secondary site.

Figure D-11 Service group configuration for sfrac13_main.cf file




sfrac14_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.
The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure D-12 illustrates the configuration.

Figure D-12 Service group configuration for sfrac14_main.cf file

Sample configuration files for CP server

The `/etc/vxcps.conf` file determines the configuration of the coordination point server (CP server.)

See [“Sample CP server configuration \(/etc/vxcps.conf\) file output”](#) on page 660.

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:
See [“Sample main.cf file for CP server hosted on a single node that runs VCS”](#) on page 655.
- The main.cf file for a CP server that is hosted on an SFHA cluster:
See [“Sample main.cf file for CP server hosted on a two-node SFHA cluster”](#) on page 657.

Note: The CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with SF Oracle RAC clusters (application clusters). The example main.cf files use IPv4 addresses.

Sample main.cf file for CP server hosted on a single node that runs VCS

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: mycps1

```
include "types.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster name:  cps1
// CP server: mycps1

cluster cps1 (
    UserNames = { admin = bMNFmHmJNiNNlVnHMK, haris = fopKojNvpHouNn,
                  "mycps1.symantecexample.com@root@vx" = aj,
                  "root@mycps1.symantecexample.com" = hq }
    Administrators = { admin, haris,
                       "mycps1.symantecexample.com@root@vx",
                       "root@mycps1.symantecexample.com" }
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

system mycps1 (
)

group CPSSG (
    SystemList = { mycps1 = 0 }
    AutoStartList = { mycps1 }
)

IP cpsvip1 (
    Critical = 0
    Device @mycps1 = eth0
    Address = "10.209.3.1"
    NetMask = "255.255.252.0"
```

```
    )

IP cpsvip2 (
    Critical = 0
    Device @mycps1 = eth1
    Address = "10.209.3.2"
    NetMask = "255.255.252.0"
)

NIC cpsnic1 (
    Critical = 0
    Device @mycps1 = eth0
    PingOptimize = 0
    NetworkHosts @mycps1 = { "10.209.3.10" }
)

NIC cpsnic2 (
    Critical = 0
    Device @mycps1 = eth1
    PingOptimize = 0
)

Process vxcperv (
    PathName = "/opt/VRTScps/bin/vxcperv"
    ConfInterval = 30
    RestartLimit = 3
)

Quorum quorum (
    QuorumResources = { cpsvip1, cpsvip2 }
)

cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
vxcperv requires quorum

// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//     {
```



```
//      NIC cpsnic1
//      }
// IP cpsvip2
//      {
//      NIC cpsnic2
//      }
// Process vxcperv
//      {
//      Quorum quorum
//      }
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: mycps1, mycps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster: cps1
// CP servers:
// mycps1
// mycps2

cluster cps1 (
    UserNames = { admin = ajkCjeJgkFkkIskEjh,
                  "mycps1.symantecexample.com@root@vx" = JK,
                  "mycps2.symantecexample.com@root@vx" = dl }
    Administrators = { admin, "mycps1.symantecexample.com@root@vx",
                        "mycps2.symantecexample.com@root@vx" }
    SecureClus = 1
)

system mycps1 (
```

```
    )

system mycps2 (
)

group CPSSG (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    AutoStartList = { mycps1, mycps2 } )

DiskGroup cpsdg (
    DiskGroup = cps_dg
)

IP cpsvip1 (
    Critical = 0
    Device @mycps1 = eth0
    Device @mycps2 = eth0
    Address = "10.209.81.88"
    NetMask = "255.255.252.0"
)

IP cpsvip2 (
    Critical = 0
    Device @mycps1 = eth1
    Device @mycps2 = eth1
    Address = "10.209.81.89"
    NetMask = "255.255.252.0"
)

Mount cpsmount (
    MountPoint = "/etc/VRTScps/db"
    BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
    FSType = vxfs
    FsckOpt = "-y"
)

NIC cpsnic1 (
    Critical = 0
    Device @mycps1 = eth0
    Device @mycps2 = eth0
    PingOptimize = 0
    NetworkHosts @mycps1 = { "10.209.81.10" }
)
```

```
NIC cpsnic2 (
    Critical = 0
    Device @mycps1 = eth1
    Device @mycps2 = eth1
    PingOptimize = 0
)

Process vxcpserve (
    PathName = "/opt/VRTScps/bin/vxcpserve"
)

Quorum quorum (
    QuorumResources = { cpsvip1, cpsvip2 }
)

Volume cpsvol (
    Volume = cps_volume
    DiskGroup = cps_dg
)

cpismount requires cpsvol
cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
cpsvol requires cpsdg
vxcpserve requires cpismount
vxcpserve requires quorum

// resource dependency tree
//
// group CPSSG
// {
//   IP cpsvip1
//   {
//     NIC cpsnic1
//   }
//   IP cpsvip2
//   {
//     NIC cpsnic2
//   }
//   Process vxcpserve
//   {
```

```
//      Quorum quorum
//      Mount cpsmount
//      {
//          Volume cpsvol
//          {
//              DiskGroup cpsdg
//          }
//      }
//  }
// }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
## The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=mycps1
vip=[10.209.81.88]
vip=[10.209.81.89]:56789
port=14250
security=1
db=/etc/VRTScps/db
```

Setting up inter-system communication

This appendix includes the following topics:

- [About using ssh or rsh with the Veritas installer](#)
- [Setting up inter-system communication](#)

About using ssh or rsh with the Veritas installer

The installer uses passwordless secure shell (ssh) or remote shell (rsh) communications among systems. The installer uses the ssh or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh or rsh communication among the systems is removed when the installation process completes, unless the installation abruptly terminates. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh or rsh configuration from the systems.

See [“Installation script options”](#) on page 586.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure ssh or rsh on the target systems. In the following scenarios, you need to set up ssh or rsh manually:

- When you add new nodes to an existing cluster.
- When the nodes are in a subcluster during a phased upgrade.
- When you perform installer sessions using a response file.

Setting up inter-system communication

If you manually need to set up a communication mode, refer to these procedures. You must have root privilege to issue ssh or rsh commands on all systems in the cluster. If ssh is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases. Similarly, rsh must be configured in such a way to not prompt for passwords.

If system communication is not possible between systems using ssh or rsh, contact Symantec Support. See <http://support.symantec.com>.

Warning: The rsh and ssh commands to the remote systems, where SF Oracle RAC is to be installed, must not print any extraneous characters.

Setting up ssh on cluster systems

Use the Secure Shell (ssh) to install SF Oracle RAC on all systems in a cluster from a system outside of the cluster. Before you start the installation process, verify that ssh is configured correctly.

Use Secure Shell (ssh) to do the following:

- Log on to another system over a network
- Execute commands on a remote system
- Copy files from one system to another

The ssh shell provides strong authentication and secure communications over channels. It is intended to replace rlogin, rsh, and rcp.

The Remote Shell (rsh) is disabled by default to provide better security. Use ssh for remote command execution.

Configuring ssh

The procedure to configure ssh uses OpenSSH example file names and commands.

Note: You can configure ssh in other ways. Regardless of how ssh is configured, complete the last step in the example to verify the configuration.

To configure ssh

- 1 Log on to the system from which you want to install SF Oracle RAC.
- 2 Generate a DSA key pair on this system by running the following command:

```
# ssh-keygen -t dsa
```

- 3 Accept the default location of `~/.ssh/id_dsa`.
- 4 When the command prompts, enter a passphrase and confirm it.
- 5 Change the permissions of the `.ssh` directory by typing:

```
# chmod 755 ~/.ssh
```

- 6 The file `~/.ssh/id_dsa.pub` contains a line that begins with `ssh_dss` and ends with the name of the system on which it was created. Copy this line to the `/root/.ssh/authorized_keys2` file on all systems where you plan to install SF Oracle RAC.

If the local system is part of the cluster, make sure to edit the `authorized_keys2` file on that system.

- 7 Run the following commands on the system where you are installing:

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

This step is shell-specific and is valid for the duration the shell is alive.

- 8 When the command prompts, enter your DSA passphrase.
You are ready to install VCS on several systems in one of the following ways:

- Run the `installvcs` program on any one of the systems
- Run the `installvcs` program on an independent system outside the cluster

To avoid running the `ssh-agent` on each shell, run the X-Window system and configure it so that you are not prompted for the passphrase. Refer to the documentation for your operating system for more information.

- 9 To verify that you can connect to the systems where you plan to install SF Oracle RAC, type:

```
# ssh -x -l root north ls
# ssh -x -l root south date
```

The commands should execute on the remote system without having to enter a passphrase or password.

Automatic Storage Management

This appendix includes the following topics:

- [About ASM in SF Oracle RAC environments](#)
- [ASM configuration with SF Oracle RAC](#)
- [Configuring ASM in SF Oracle RAC environments](#)
- [Sample configuration file Veritas CVM and ASM main.cf file](#)

About ASM in SF Oracle RAC environments

ASM is an integrated storage management solution from Oracle RAC that combines file system and volume management capabilities.

ASM can be configured with Cluster Volume Manager (CVM) for better performance and availability. CVM mirrored volumes with dynamic multi-pathing improves data access performance and offers continuous data availability in large heterogeneous SAN environments. You can create CVM disk groups and volumes for use as ASM disks groups and configure the ASM disk groups to be managed by the Veritas ASMDG agent. The ASMDG agent mounts, unmounts, and monitors the ASM disk groups.

ASM provides storage for Oracle Cluster Registry devices (OCR), voting disk, data files, control files, online redo logs, archive log files, and backup files.

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

ASM configuration with SF Oracle RAC

Configure ASM disk groups over CVM volumes in SF Oracle RAC environments. The CVM volumes are mirrored for high availability of data and leverage Dynamic Multi-Pathing to access the shared storage.

Figure F-1 illustrates the configuration of ASM disk groups over CVM.

Figure F-1 ASM disk groups over CVM





Figure F-2 illustrates the following supported configuration for Oracle RAC 11g Release 2:

- ASM disk groups configured over CVM volumes
- Oracle Clusterware and database binaries stored locally
- Oracle database files stored on ASM configured over CVM
The Oracle databases are managed by Oracle Clusterware.
- Oracle Cluster Registry and voting disk stored on ASM over CVM volumes

Figure F-2 Supported Oracle RAC configuration (Oracle RAC 11g Release 2)

Configuring ASM in SF Oracle RAC environments

Before you configure ASM, review the planning guidelines for ASM:

See [“Planning for Oracle ASM over CVM”](#) on page 57.

Note: Make sure that you have installed SF Oracle RAC and Oracle RAC before configuring ASM.

The ASM home directory is the same as the Oracle Grid Infrastructure home directory (GRID_HOME).

To configure ASM in SF Oracle RAC installations

- 1 Unlink the Veritas ODM library from the ORACLE_HOME directory.
See [“Unlinking the SF Oracle RAC ODM library”](#) on page 667.
- 2 Create the required database storage on ASM by creating CVM disk groups and volumes for use as ASM disks.
See [“Creating database storage on ASM”](#) on page 667.
- 3 Create ASM disk groups and instances.
See [“Creating ASM disk groups and instances”](#) on page 668.
- 4 Verify the ASM setup.
See [“Verifying the ASM setup”](#) on page 669.

- 5 Create the Oracle database. For instructions, see the Oracle RAC documentation.
- 6 Configure VCS service groups for the Oracle database.
See [“Configuring VCS service groups for database instances on ASM”](#) on page 670.

Unlinking the SF Oracle RAC ODM library

ODM is a disk management interface for data files that reside on the Veritas File system.

You need to unlink the ODM library from the ORACLE_HOME directory for ASM before you configure the database storage on ASM with SF Oracle RAC.

For Oracle RAC 11g Release 2, the Oracle home directory for ASM is the same as the GRID_HOME directory.

Perform the steps in the procedure on each node in the cluster.

To unlink the Veritas ODM library

- 1 Log in as the Oracle user.
- 2 Shut down the ASM instance if it is running:

```
$ srvctl stop asm -n galaxy
```

```
$ srvctl stop asm -n nebula
```
- 3 Change to the Oracle home directory for ASM:

```
$ cd $ORACLE_HOME/lib
```
- 4 Unlink the Veritas ODM libraries:

```
$ ln -sf libodmdm11.so libodm11.so
```

Creating database storage on ASM

This step creates the database storage on ASM using CVM volumes. To create the storage for Oracle databases on ASM, first create the required CVM disk groups and volumes. Then, use these CVM volumes to create ASM disk groups for storing the database files.

To create database storage on ASM

- 1 Log in as the root user to the CVM master:

To determine the CVM master:

```
# vxdctl -c mode
```

- 2 Initialize the disks as VxVM disks:

```
# vxdisksetup -i sdx
```

- 3 Create the CVM disk group and volume:

```
# vxdg -s init ora_asm_dg sdx
```

```
# vxassist -g ora_asm_dg make ora_asm_vol 2000M
```

- 4 Set the permissions for the Oracle user on the volume:

```
# vxedit -g ora_asm_dg \  
set group=dba user=oracle mode=660 ora_asm_vol
```

Configure ASM using either ASMCA or OEM.

Creating ASM disk groups and instances

You can create ASM disk groups using ASM Configuration Assistant (ASMCA) or manually.

The following tasks are performed:

- The ASM disk group and instance is created.
- The ASM instance is started and the disk group is mounted on all nodes in the cluster. The default ASM instance name is +ASM n where n is the instance number depending on the number of ASM instances.

Note: For ASM instances that use a pfile, if you restart a node, make sure that the underlying volumes are available before the ASM disk group is mounted. Then, update the ASM init.ora parameter with the full path that contains the VxVM volumes created for ASM.

For ASM instances that use an spfile, the parameter is updated automatically by the Oracle Configuration Assistant.

In either case, dependencies must be managed manually. Symantec recommends the use of the ASMDG agent for easier management of ASM disk groups.

SF Oracle RAC requires the following settings when you create ASM disk groups and instances:

Type of parameter file for the ASM instance	server parameter file (spfile)
Disk Discovery Path	Enter the full path that contains the VxVM volumes created for ASM, for example /dev/vx/rdisk/ora_asm_dg/ora_asm_vol.
Redundancy	Select the option External . Mirroring is performed by CVM.
Select member disks	Select the VxVM disks you want to use; You may need to select the Force checkbox next to the ASM disks you want to use if the disk group creation fails.

For detailed instructions, see the Oracle documentation.

Verifying the ASM setup

Verify that the database services for ASM are up and running after the installation.

To verify the ASM installation

- 1 Change to the Oracle Grid Infrastructure home directory:

```
# cd $GRID_HOME/bin
```

- 2 Verify the status of ASM on all the nodes in the cluster:

```
# ./srvctl status asm
ASM instance +ASM1 is running on node galaxy
```

The sample output shows that there is one ASM instance running on the local node.

Configuring VCS service groups for database instances on ASM

This section describes how to configure the Oracle service group using the CLI for databases on ASM.

For sample service group illustration:

See “[sfrac05_main.cf file](#)” on page 642.

For a sample configuration file describing the configuration, see the file `sfrac05_main.cf` in the directory `/etc/VRTSvcs/conf/sample_rac/`.

The following procedure assumes that you have created the database.

To configure the Oracle service group using the CLI

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the service group to the VCS configuration:

```
# hagrps -add dbgrp
```

- 3 Modify the attributes of the service group:

```
# hagrps -modify dbgrp Parallel 1
```

```
# hagrps -modify dbgrp SystemList galaxy 0 nebula 1
```

```
# hagrps -modify dbgrp AutoStartList galaxy nebula
```

- 4 Add the CVMVolDg resource for the service group:

```
# hares -add oradata_voldg CVMVolDg dbgrp
```

- 5 Modify the attributes of the CVMVolDg resource for the service group:

```
# hares -modify oradata_voldg CVMDiskGroup ora_asm_dg
```

```
# hares -modify oradata_voldg CVMActivation sw
```

```
# hares -modify oradata_voldg CVMVolume ora_asm_vol
```

- 6 Add the ASMDG resource for the service group:

```
# hares -add asmdg ASMDG dbgrp
```

- 7 Modify the attributes of the ASMDG resource for the service group.

Note: The \$ASM_HOME variable refers to the ASM home directory.

For Oracle RAC 11g Release 2, the ASM home directory is the same as the GRID_HOME directory.

```
# hares -modify asmdg DiskGroups ASM_RAC_DG
# hares -modify asmdg Home "$ASM_HOME"
# hares -local asmdg Sid
# hares -modify asmdg Sid "+ASM1" -sys galaxy
# hares -modify asmdg Sid "+ASM2" -sys nebula
# hares -modify asmdg Owner grid
```

- 8 Add the Oracle RAC database instance to the service group:

```
# hares -add asmdb Oracle dbgrp
```

- 9 Modify the attributes of the Oracle resource for the service group:

```
# hares -modify asmdb Owner oracle
# hares -local asmdb Sid
# hares -modify asmdb Sid oradb1 -sys galaxy
# hares -modify asmdb Sid oradb2 -sys nebula
# hares -modify asmdb Home "$ORACLE_HOME"
# hares -modify asmdb StartUpOpt SRVCTLSTART
# hares -modify asmdb ShutDownOpt SRVCTLSTOP
```

- 10 Set the dependencies between the ASMDG resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link asmdg oradata_voldg
```

- 11 Set the dependencies between the Oracle resource and the ASMDG resource for the Oracle service group:

```
# hares -link asmdb asmdg
```

- 12 Create an online local firm dependency between the dbgrp service group and the cvm service group:

```
# hagrps -link dbgrp cvm online local firm
```

13 Enable the Oracle service group:

```
# hagrps -enableresources dbgrp
```

14 Change the cluster configuration to the read-only mode:

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

15 Bring the Oracle service group online on all the nodes:

```
# hagrps -online dbgrp -any
```

Sample configuration file Veritas CVM and ASM main.cf file

A sample configuration file (sfrac05_main.cf) illustrating ASM configuration in SF Oracle RAC is available at `/etc/VRTSvcs/conf/sample_rac/`.

For an illustration of the service group configuration:

See [“sfrac05_main.cf file”](#) on page 642.

Creating a test database

This appendix includes the following topics:

- [About creating a test database](#)
- [Creating a database for Oracle](#)

About creating a test database

A test database can be created and used for both testing and troubleshooting purposes.

The following optional procedures describe the methods for creating a test database.

Creating a database for Oracle

Before you begin to create the database, ensure that the following prerequisites are met:

- The CRS daemons must be running.
To verify the status of Oracle Clusterware, type the following command:

```
# $CRS_HOME/bin/crs_stat
```
- All private IP addresses on each node must be up.
Use the `ping` command to verify that all private IP addresses on each node are up.

Refer to your Oracle documentation for instructions on how to install the Oracle database.

You can create the database on one of the following types of storage:

- Shared raw volume

See [“Creating the database storage on raw volumes”](#) on page 674.

- Cluster File System (CFS)
 See [“Creating the database storage on CFS”](#) on page 675.
- ASM
 See [“Creating database storage on ASM”](#) on page 667.

Creating the database storage on raw volumes

You can create the database storage on shared raw volume.

To create the database storage on shared raw volumes

- 1 Log in as root user.
- 2 On the master node, create a shared disk group:

```
# vxdg -s init oradatadg Disk_1
```

- 3 Create a volume in the shared group for each of the required tablespaces.
 Refer to the Oracle documentation to determine the tablespace requirements.
 For example, type:

```
# vxassist -g oradatadg make VRT_galaxy 1000M
# vxassist -g oradatadg make VRT_sfile1 10M
.
.
```

- 4 Define the access mode and permissions for the volumes storing the Oracle data.

For each volume listed in \$ORACLE_HOME/raw_config, use the `vxedit(1M)` command:

```
# vxedit -g disk_group set group=group user=user mode=660 volume
```

For example:

```
# vxedit -g oradatadg set group=dba user=oracle mode=660 \
VRT_galaxy
```

In this example, VRT_galaxy is the name of one of the volumes. Repeat the command to define access mode and permissions for each volume in the oradatadg.

- 5 Create the database using the Oracle documentation.

Creating the database storage on CFS

If you plan to use a cluster file system to store the Oracle database, use the following procedure to create the file system.

To create the database storage on CFS

- 1 Create a disk group (for example, oradatadg):

```
# vxvg -s init oradatadg Disk_1
```

- 2 Create a single shared volume (for example, oradatavol) that is large enough to contain a file system for all tablespaces.

For example, assuming 6.8 GB is required for database storage, type:

```
# vxassist -g oradatadg make oradatavol 6800M
```

- 3 Start the volume in the disk group:

```
# vxvol -g oradatadg startall
```

- 4 Create a VxFS file system in this volume. From one node, type:

```
# mkfs -t 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 -t vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol \
/oradata
```

- 7 Set the "Oracle" user as the owner of the file system, and set "755" as the permissions:

```
# chown oracle:oinstall /oradata
# chmod 755 /oradata
```

- 8 On the other node(s), complete step 5 and step 6.
- 9 Create the database using the Oracle documentation.

High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount agent](#)
- [CFSfsckd agent](#)
- [PrivNIC agent](#)
- [MultiPrivNIC agent](#)
- [CSSD agent](#)
- [VCS agents for Oracle](#)
- [CRSResource agent](#)

About agents

An agent is defined as a process that starts, stops, and monitors all configured resources of a type, and reports their status to Veritas Cluster Server (VCS). Agents have both entry points and attributes. Entry points are also known as agent functions and are referred to as "agent functions" throughout the document.

Attributes contain data about the agent. An attribute has a definition and a value. You change attribute values to configure resources, which are defined as the individual components that work together to provide application services to the public network. For example, a resource may be a physical component such as a disk or a network interface card, a software component such as Oracle or a Web server, or a configuration component such as an IP address or mounted file system.

Attributes are either optional or required, although sometimes the attributes that are optional in one configuration may be required in other configurations. Many optional attributes have predefined or default values, which you should change as required. A variety of internal use only attributes also exist. Do not modify these attributes—modifying them can lead to significant problems for your clusters. Attributes have type and dimension. Some attribute values can accept numbers, others can accept alphanumeric values or groups of alphanumeric values, while others are simple boolean on/off values.

The entry points and attributes for each SF Oracle RAC agent are described in this appendix.

VCS agents included within SF Oracle RAC

SF Oracle RAC includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDg agent
- CFSSMount agent

An SF Oracle RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

You must configure the CVMVolDg agent for each shared disk group. If the database uses cluster file systems, configure the CFSSMount agent for each volume in the disk group.

Use the information in this appendix about the entry points and attributes of the listed agents to make necessary configuration changes. For information on how to modify the VCS configuration:

See the *Veritas Cluster Server Administrator's Guide*

VCS agents for Oracle included within SF Oracle RAC

SF Oracle RAC includes the following VCS agents for Oracle:

- Oracle agent

The Oracle agent monitors the database processes.

- **Netlsnr agent**
The Netlsnr agent brings the listener services online, monitors their status, and takes them offline.
- **PrivNIC agent**
The PrivNIC agent provides high availability to a single private IP address across LLT Ethernet interfaces for a system.
- **MultiPrivNIC agent**
The MultiPrivNIC agent provides high availability to multiple private IP addresses across LLT Ethernet interfaces for a system.
- **CSSD agent**
The CSSD (Cluster Synchronization Services daemon) agent provides the resources to monitor Oracle Clusterware. The agent ensures that the dependency of cssd on the OCR and the VOTE resources and the PrivNIC (optional) resource are satisfied.
- **ASMDG agent**
The ASMDG agent mounts and unmounts the ASM disk groups onto an ASM instance.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide*

CVMCluster agent

The CVMCluster agent controls system membership on the cluster port that is associated with Veritas Volume Manager (VxVM).

The CVMCluster agent performs the following functions:

- Joins a node to the CVM cluster port.
- Removes a node from the CVM cluster port.
- Monitors the node's cluster membership state.

Entry points for CVMCluster agent

[Table H-1](#) describes the entry points used by the CVMCluster agent.

Table H-1 CVMCluster agent entry points

Entry Point	Description
Online	Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by automatically importing the shared disk groups.

Table H-1 CVMCluster agent entry points (*continued*)

Entry Point	Description
Offline	Removes a node from the CVM cluster port.
Monitor	Monitors the node's CVM cluster membership state.

Attribute definition for CVMCluster agent

[Table H-2](#) describes the user-modifiable attributes of the CVMCluster resource type.

Table H-2 CVMCluster agent attributes

Attribute	Description
CVMClustName	Name of the cluster. ■ Type and dimension: string-scalar
CVMNodeAddr	List of host names and IP addresses. ■ Type and dimension: string-association
CVMNodeId	Associative list. The first part names the system; the second part contains the LLT ID number for the system. ■ Type and dimension: string-association
CVMTransport	Specifies the cluster messaging mechanism. ■ Type and dimension: string-scalar ■ Default = gab Note: Do not change this value.
PortConfigd	The port number that is used by CVM for vxconfigd-level communication. ■ Type and dimension: integer-scalar
PortKmsgd	The port number that is used by CVM for kernel-level communication. ■ Type and dimension: integer-scalar
CVMTimeout	Timeout in seconds used for CVM cluster reconfiguration. ■ Type and dimension: integer-scalar ■ Default = 200

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```
type CVMCluster (
    static int InfoTimeout = 0
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName,
        CVMNodeAddr, CVMNodeId, PortConfigd, PortKmsgd,
        CVMTimeout }
    NameRule = ""
    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
    int PortConfigd
    int PortKmsgd
    int CVMTimeout
)
```

Note: The attributes `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` are not used in an SF Oracle RAC environment. GAB, the required cluster communication messaging mechanism, does not use them.

CVMCluster agent sample configuration

The following is an example definition for the CVMCluster service group:

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

CVMVxconfigd agent

The CVMVxconfigd agent starts and monitors the vxconfigd daemon. The vxconfigd daemon maintains disk and disk group configurations, communicates configuration

changes to the kernel, and modifies the configuration information that is stored on disks. CVMVxconfigd must be present in the CVM service group.

The CVMVxconfigd agent is an OnOnly agent; the agent starts the resource when the cluster starts up and VCS restarts the resource when necessary. The Operations attribute specifies these default aspects of startup.

Symantec recommends starting the vxconfigd daemon with the `syslog` option, which enables logging of debug messages. Note that the SF Oracle RAC installation configures the `syslog` option for the CVMVxconfigd agent.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CVMVxconfigd agent

[Table H-3](#) describes the entry points for the CVMVxconfigd agent.

Table H-3 CVMVxconfigd entry points

Entry Point	Description
Online	Starts the <code>vxconfigd</code> daemon
Offline	N/A
Monitor	Monitors whether <code>vxconfigd</code> daemon is running
imf_init	Initializes the agent to interface with the AMF kernel module. This function runs when the agent starts up.
imf_getnotification	Gets notification about the <code>vxconfigd</code> process state. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification. If the <code>vxconfigd</code> process fails, the function initiates a traditional CVMVxconfigd monitor entry point.
imf_register	Registers or unregisters the <code>vxconfigd</code> process id (pid) with the AMF kernel module. This function runs after the resource goes into steady online state.

Attribute definition for CVMVxconfigd agent

[Table H-4](#) describes the modifiable attributes of the CVMVxconfigd resource type.

Table H-4 CVMVxconfigd agent attribute

Attribute	Description
CVMVxconfigdArgs	List of the arguments that are sent to the <code>online</code> entry point. Symantec recommends always specifying the <code>syslog</code> option. ■ Type and dimension: keylist
IMF	This resource-type level attribute determines whether the CVMVxconfigd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association For more details of IMF attribute for the agent type, refer to the <i>Veritas Cluster Server Administrator's Guide</i>.

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 }
)
```

CVMVoIDg agent

The CVMVoIDg agent manages the CVM disk groups and CVM volumes and volume sets within the disk groups by performing the following functions:

- Imports the shared disk group from the CVM master node
- Starts the volumes and volume sets in the disk group
- Monitors the disk group, volumes, and volume sets
- Optionally, deports the disk group when the dependent applications are taken offline. The agent deports the disk group only if the appropriate attribute is set.

Configure the CVMVoIDg agent for each disk group used by a Oracle service group. A disk group must be configured to only one Oracle service group. If cluster file systems are used for the database, configure the CFSMount agent for each volume or volume set in the disk group.

Entry points for CVMVoIDg agent

[Table H-5](#) describes the entry points used by the CVMVoIDg agent.

Table H-5 CVMVoIDg agent entry points

Entry Point	Description
Online	Imports the shared disk group from the CVM master node, if the disk group is not already imported. Starts all volumes and volume sets in the shared disk group specified by the CVMVolume attribute. 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.
Offline	Removes the temporary files created by the online entry point. If the <code>CVMDeportOnOffline</code> attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.
Monitor	Determines whether the disk group, the volumes, and the volume sets are online. The agent takes a volume set offline if the file system metadata volume of a volume set is discovered to be offline in a monitor cycle. Note: If the CFSMount resource goes offline and the file system on the volume set is unmounted, the agent retains the online state of the volume set even if the file system metadata volume in the volume set is offline. This is because the CVMVoIDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.
Clean	Removes the temporary files created by the online entry point.

Attribute definition for CVMVoIDg agent

[Table H-6](#) describes the user-modifiable attributes of the CVMVoIDg resource type.

Table H-6 CVMVoIDg agent attributes

Attribute	Description
CVMDiskGroup (required)	Shared disk group name. Type and dimension: string-scalar

Table H-6 CVMVolDg agent attributes (*continued*)

Attribute	Description
CVMVolume (required)	Name of shared volumes or volume sets. This list is used to check that the volumes or volume sets are in the correct state before allowing the resource to come online, and that the volumes remain in an enabled state. ■ Type and dimension: string-keylist
CVMActivation (required)	Activation mode for the disk group. ■ Type and dimension: string-scalar ■ Default = <code>sw</code> (<code>shared-write</code>) This is a localized attribute.
CVMVolumeIoTest(optional)	List of volumes and volume sets that will be periodically polled to test availability. The polling is in the form of 4 KB reads every monitor cycle to a maximum of 10 of the volumes or volume sets in the list. For volume sets, reads are done on a maximum of 10 component volumes in each volume set. ■ Type and dimension: string-keylist
CVMDeportOnOffline (optional)	Indicates whether or not the shared disk group must be deported when the last online CVMVolDg resource for a disk group is taken offline. The value 1 indicates that the agent will deport the shared disk group from the CVM master node, if not already deported, when the last online CVMVolDg resource for the disk group is taken offline. The value 0 indicates that the agent will not deport the shared disk group when the CVMVolDg resource is taken offline. ■ Type and dimension: integer-scalar ■ Default = 0 Note: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the attribute to either 1 or 0 for all of the resources. The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the <code>CVMDeportOnOffline</code> attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported. The deport operation fails if the shared disk group contains open volumes.

CVMVolDg agent type definition

The CVMTypes.cf file includes the CVMVolDg type definition:

```

type CVMVolDg (
    static keylist RegList = { CVMActivation, CVMVolume }
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMDiskGroup, CVMVolume, CVMActivation,
        CVMVolumeIoTest, CVMDGAction, CVMDeportOnOffline }
    str CVMDiskGroup
    str CVMDGAction
    keylist CVMVolume
    str CVMActivation
    keylist CVMVolumeIoTest
    int CVMDeportOnOffline
    temp int voldg_stat
)

```

CVMVolDg agent sample configuration

Each Oracle service group requires a CVMVolDg resource type to be defined. The following is a sample configuration:

```

CVMVolDg ora_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradata1, oradata2 }
    CVMActivation = sw
)

```

CFSMount agent

The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point.

The agent executable is located in /opt/VRTSvcs/bin/CFSMount/CFSMountAgent.

The CFSMount type definition is described in the /etc/VRTSvcs/conf/config/CFSTypes.cf file.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CFSMount agent

[Table H-7](#) provides the entry points for the CFSMount agent.

Table H-7 CFSMount agent entry points

Entry Point	Description
Online	Mounts a block device in cluster mode.
Offline	Unmounts the file system, forcing unmount if necessary, and sets primary to secondary if necessary.
Monitor	Determines if the file system is mounted. Checks mount status using the <code>fsclustadm</code> command.
Clean	Generates a null operation for a cluster file system mount.
imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.
imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSMount agent

[Table H-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

Table H-8 CFSMount Agent attributes

Attribute	Description
MountPoint	Directory for the mount point. ■ Type and dimension: string-scalar
BlockDevice	Block device for the mount point. ■ Type and dimension: string-scalar
NodeList	List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list. ■ Type and dimension: string-keylist

Table H-8 CFSMount Agent attributes (*continued*)

Attribute	Description
IMF	Resource-type level attribute that determines whether the CFSMount agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the <code>RegisterRetyLimit</code> key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the <code>Mode</code> key changes. Default: 3. ■ Type and dimension: integer-association

Table H-8 CFSMount Agent attributes (*continued*)

Attribute	Description
MountOpt (optional)	Options for the mount command. To create a valid MountOpt attribute string: ■ Use the VxFS type-specific options only. ■ Do not use the <code>-o</code> flag to specify the VxFS-specific options. ■ Do not use the <code>-t vxfs</code> file system type option. ■ Be aware the cluster option is not required. ■ Specify options in comma-separated list: <pre>ro ro,cluster blkclear,mincache=closesync</pre> ■ Type and dimension: string-scalar
Policy (optional)	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. ■ Type and dimension: string-scalar

CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition:

```
type CFSMount (
    static keylist RegList = { MountOpt, Policy, NodeList, ForceOff, SetPrimary }
    static keylist SupportedActions = { primary }
    static int FaultOnMonitorTimeouts = 1
    static int OnlineWaitLimit = 1
    static str ArgList[] = { MountPoint, BlockDevice, MountOpt, Primary, AMFMountType }
    str MountPoint
    str MountType
    str BlockDevice
    str MountOpt
    keylist NodeList
    keylist Policy
    temp str Primary
    str SetPrimary
    temp str RemountRes
    temp str AMFMountType
    str ForceOff
)
```

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;  
)
```

CFSfsckd agent

The CFSfsckd agent starts, stops, and monitors the `vxfsckd` process. The CFSfsckd agent executable is `/opt/VRTSvcs/bin/CFSfsckd/CFSfsckdAgent`. The type definition is in the `/etc/VRTSvcs/conf/config/CFSTypes.cf` file. The configuration is added to the `main.cf` file after running the `cfsccluster config` command.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CFSfsckd agent

[Table H-9](#) describes the CFSfsckd agent entry points.

Table H-9 CFSfsckd agent entry points

Entry Points	Description
Online	Starts the vxfsckd process.
Offline	Kills the vxfsckd process.
Monitor	Checks whether the vxfsckd process is running.
Clean	A null operation for a cluster file system mount.
imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.
imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.

Table H-9 CFSfsckd agent entry points *(continued)*

Entry Points	Description
imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSfsckd agent

[Table H-10](#) lists user-modifiable attributes of the CFSfsckd Agent resource type.

Table H-10 CFSfsckd Agent attributes

Attribute	Description
IMF	Resource-type level attribute that determines whether the CFSfsckd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association

CFSfsckd agent type definition

The CFSfsckd type definition:

```
type CFSfsckd (  
    static int RestartLimit = 1  
    str ActivationMode{}  
)
```

CFSfsckd agent sample configuration

This is a sample of CFSfsckd configuration:

```
CFSfsckd vxfsckd (  
)
```

PrivNIC agent

The PrivNIC agent provides high availability to a single private IP address across LLT Ethernet interfaces for a system. Private IP addresses are required by Oracle Clusterware and the Oracle database to provide communication between the cluster nodes.

Note: The PrivNIC agent operates over LLT links. To use the agent, the Oracle Clusterware interconnects and the Oracle RAC database communication links must be configured as LLT links.

The PrivNIC agent relies on LLT to monitor the LLT Ethernet interfaces. It queries LLT for the number of visible nodes on each of the LLT Ethernet interfaces.

The PrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide increased bandwidth using multiple network interfaces. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to the connected or available LLT link. If the preferred link becomes available, the IP address is failed back to the preferred link.

Note: The PrivNIC agent is not supported with Oracle RAC 11.2.0.2. For more information, see <http://www.symantec.com/business/support/index?page=content&id=TECH145261>

Functions of the PrivNIC agent

[Table H-11](#) describes the PrivNIC agent's monitor entry point.

Note: Because the resource is persistent, only the monitor entry point is required.

Table H-11 PrivNIC agent entry point

Entry Point	Description
Monitor	The PrivNIC agent queries LLT to create a list of nodes visible on every LLT network interface. The PrivNIC agent then applies various filters to this list to arrive at a most desired failover decision and calculates a "winner" device on which to configure the IP address. The "winner" device is compared to the currently active device where the IP address is currently configured. If the active and "winner" devices are different, the PrivNIC agent initiates a failover to the "winner" device.

Attributes of the PrivNIC agent

[Table H-12](#) describes the user-modifiable attributes of the PrivNIC agent.

Table H-12 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: <pre>Device@galaxy = {eth1=0, eth2=1, eth3=2} Device@nebula = {eth1=0, eth2=1, eth3=2}</pre>

Table H-12 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. When configuring private IPv4 addresses for Oracle Clusterware, make sure that there are no leading zeroes in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1 or X.0X.X.1 or 0X.X.X.1. Ensure that the IPv4 addresses have the format as "X.X.X.1". 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 = "255.255.255.0"</pre>

Optional attributes of the PrivNIC agent

Table H-13 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. For example: in the common case, <code>/etc/llttab</code> contains: <pre>link eth0 /dev/eth:0 - ether - - link eth1 /dev/eth:1 - ether - - link-lowpri eth0 /dev/eth:0 - ether - -</pre> In the above case, DeviceTag does not need to be specified. However, if <code>/etc/llttab</code> contains: <pre>link link1 /dev/eth:0 - ether - - link link2 /dev/eth:1 - ether - - link-lowpri spare /dev/eth:0 - ether - -</pre> And, <pre>Device@galaxy = { eth0=0, eth1=1, eth2=2 }</pre> DeviceTag needs to be specified as: <pre>DeviceTag@galaxy = { spare=2 }</pre>
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>
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.

Table H-13 Optional attributes for PrivNIC agent (*continued*)

Attribute	Dimension	Description
UseSystemList	integer-scalar	The value 1 specifies that the agent use the SystemList of the service group to filter the node list. Default = 0.
ExcludeNode	integer-vector	List of nodes to permanently exclude from calculation.

States of the PrivNIC agent


[Table H-14](#) lists the states of the PrivNIC agent.

Table H-14 States of the PrivNIC agent

State	Description
Online	Indicates that the private IP address is available.
Unknown	Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues.

Sample service group configuration with the PrivNIC agent

[Figure H-1](#) illustrates a basic service group configuration with the PrivNIC agent.

Figure H-1 Basic service group configuration with the PrivNIC agent

This configuration shows the dependency of the CSSD resource on the private IP address configured as a PrivNIC resource along with the OCR and voting disk volume and mount point dependencies.

For sample deployment scenarios, see the appendix *SF Oracle RAC deployment scenarios*.

Type definition of the PrivNIC resource

The following extract shows the type definition of the PrivNIC resource in the PrivNIC.cf file:

```
type PrivNIC (
    static str ArgList[] = { Device, DeviceTag, Address,
                           NetMask, UseVirtualIP, GabPort, UseSystemList,
                           ExcludeNode }
    static int OfflineMonitorInterval = 60
    static int MonitorTimeout = 300
    static str Operations = None

    str Device{}
    str DeviceTag{}
    str Address = ""
    str NetMask = ""
    int UseVirtualIP = 0
    str GabPort = "o"
    int UseSystemList = 0
    int ExcludeNode[]
)
```

Sample configuration of the PrivNIC resource

The following extract from the configuration file illustrates the configuration of a PrivNIC resource.

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

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { eth0 = 0, eth1 = 1 }
    Device@nebula = { eth0 = 0, eth1 = 1 }
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
```

```
NetMask = "255.255.255.0"  
)
```

For more examples, see the sample configuration files located at
`/etc/VRTSvcs/conf/sample_rac/`.

MultiPrivNIC agent

The MultiPrivNIC agent provides high availability to multiple private IP addresses across LLT Ethernet interfaces for a system. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to one of the available LLT links. To use the agent, the Oracle Clusterware interconnects and the Oracle RAC database communication links must be configured as LLT links.

The MultiPrivNIC agent is a reliable alternative in scenarios where operating system limitations prevent you from using NIC bonding to provide increased bandwidth and high availability using multiple network interfaces. Even if link aggregation solutions in the form of bonded NICs are implemented, the MultiPrivNIC agent can be used to provide additional protection against the failure of aggregated links by failing over the IP addresses to the available alternate links. These alternate links can be simple NIC interfaces or bonded NICs.

Note: The MultiPrivNIC agent is not supported with Oracle RAC 11.2.0.2. For more information, see

<http://www.symantec.com/business/support/index?page=content&id=TECH145261>

Managing high availability of private interconnects

The MultiPrivNIC agent operates over LLT links and relies on LLT to monitor the cluster interfaces. It queries LLT to count and report the number of visible nodes on each of the LLT interfaces. In the event that a preferred link goes down, the IP address is failed over to the private link on which maximum number of peer nodes are visible. If multiple links see maximum nodes and if load-balancing is enabled, the agent considers the current traffic on all devices and calculates a "winner" device with lower traffic. If load balancing is not enabled, the IP address is failed over to the link with the lower network-id.

The failover decision for an IP address is made only when the link hosting the IP address fails. If the preferred link becomes available, the IP address is failed back to the preferred link regardless of whether load-balancing is enabled or disabled.

Functions of the MultiPrivNIC agent

[Table H-15](#) describes the MultiPrivNIC agent's monitor entry point.

Note: Because the resource is persistent, only the monitor entry point is required.

Table H-15 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 agent applies various filters to this list and calculates a winner device on which to configure the IP address. If the active device on which the IP address is configured does not see the same number of nodes as the winner device, the agent fails over the IP address to the winner device.

Attributes of the MultiPrivNIC agent

[Table H-16](#) below describes the user-modifiable attributes of the MultiPrivNIC resource type.

Table H-16 MultiPrivNIC agent attribute definitions

Attribute	Dimension	Description
Device	string-association	The device attribute specifies the network interface 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 interfaces specified in the MultiPrivNIC configuration should be exactly the same in name and total number as those which have been used for LLT configuration. An example of the device attribute is as follows: <pre>Device@galaxy = {eth1=0, eth2=1, eth3=2} Device@nebula = {eth1=0, eth2=1, eth3=2}</pre>

Table H-16 MultiPrivNIC agent attribute definitions (*continued*)

Attribute	Dimension	Description
Address	string-association	The numerical private IP address and its preferred device. The agent verifies that the IP addresses are valid addresses. When you configure private IP addresses, ensure that the addresses do not have a leading 0 in any of the octets that comprise the IP address. The IP address must be in the format "X.X.X.1". In the following example, 0 and 1 indicates the ID of the device on which the IP address is hosted. If the device is unavailable, the agent automatically reconfigures the IP address on one of the available devices. The available device is determined based on the UseLoadBalance attribute setting. An example of the address attribute is as follows: <pre>Address @galaxy = { "192.168.12.1" =0, "192.168.2.1" =0, "192.168.3.1" =1 } Address @nebula = { "192.168.12.2" =0, "192.168.2.2" =0, "192.168.3.2" =1 }</pre>
Netmask	string-association	The netmask attribute is the numerical netmask for the private IP address. For example, Address = "255.255.255.0".
UseLoadBalance		A boolean value 0 or 1. In the event that the preferred device is unavailable and multiple links see maximum nodes during failover: ■ Setting the attribute to 1 fails over the IP address to the device with lower traffic. ■ Setting the attribute to 0 fails over the IP address to the device with lower network ID. Note: If the preferred device becomes available, the IP address is failed back to the preferred device regardless of the value of UseLoadBalance.

States of the MultiPrivNIC agent

[Table H-17](#) lists the states of the MultiPrivNIC agent.


Table H-17 States of the MultiPrivNIC agent

State	Description
Online	Indicates that the private IP addresses are available.
Unknown	Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues.

Sample service group configuration with the MultiPrivNIC agent

[Figure H-2](#) illustrates a sample service group configuration with the MultiPrivNIC agent.

Figure H-2 Basic service group configuration with MultiPrivNIC agent



The illustrated configuration shows the dependency of the CSSD resource on the private IP address configured as a MultiPrivNIC resource along with the OCR and voting disk volume and mount point dependencies.

For sample deployment scenarios, see the appendix *SF Oracle RAC deployment scenarios*.

Type definition of the MultiPrivNIC resource

The following extract shows the type definition of the MultiPrivNIC resource in the `MultiPrivNIC.cf` file:

```

type MultiPrivNIC (
    static int MonitorTimeout = 300
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { Device, DeviceTag, Address, NetMask,

```

```

        UseVirtualIP, GabPort, UseSystemList,
        ExcludeNode, UseLoadBalance }

static str Operations = None
str Device{}
str DeviceTag{}
str Address{}
str NetMask
int UseVirtualIP
str GabPort = o
int UseSystemList
int ExcludeNode[]
int UseLoadBalance
)

```

Sample configuration of the MultiPrivNIC resource

The following extract from the configuration file illustrates the configuration of a MultiPrivNIC resource.

```

MultiPrivNIC multi_priv (
    Critical = 0
    Device @galaxy = { eth1 = 0, eth2 = 1 }
    Device @nebula = { eth1 = 0, eth2 = 1 }
    Address @galaxy = { "192.168.12.1" =0, "192.168.2.1" =0,
                        "192.168.3.1" =1 }
    Address @nebula = { "192.168.12.2" =0, "192.168.2.2" =0,
                        "192.168.3.2" =1 }
    NetMask = "255.255.255.0"
    UseLoadBalance = 1
)

```

For more examples, see the sample configuration files located at
/etc/VRTSvcs/conf/sample_rac/.

In the above example, the interface next to the IP address is the preferred device for that particular IP address. If the same number of nodes are visible on each of the LLT interfaces, the IP addresses are configured on the preferred interfaces.

CSSD agent

The CSSD agent starts, stops, and monitors Oracle Clusterware/Grid Infrastructure. It ensures that the OCR, the voting disk and the private IP address resources required by Oracle Clusterware/Grid Infrastructure are online before Oracle Clusterware/Grid Infrastructure starts. For this purpose, the cssd resource must be

configured as a parent of the resources that manage the OCR, the voting disk, and the private IP address used by Oracle Clusterware/Grid Infrastructure. Using the CSSD agent in SF Oracle RAC installations ensures adequate handling of inter-dependencies and thus prevents the premature startup of Oracle Clusterware/Grid Infrastructure, which causes cluster failures.

For Oracle RAC 11g Release 2, the automatic startup of Oracle Clusterware/Grid Infrastructure must be disabled when the system starts. This prevents the premature startup of Oracle Clusterware/Grid Infrastructure during system startup. Thus, VCS ensures that Oracle Clusterware/Grid Infrastructure is started using the CSSD agent only when all the dependent resources of the cssd resource configured under VCS are online.

During system shutdown, the agent stops Oracle Clusterware/Grid Infrastructure before the OCR and voting disk resources are taken offline by VCS. This ensures that Oracle Clusterware/Grid Infrastructure does not panic the nodes in the cluster due to unavailability of the required resources.

Note: It is mandatory to use CSSD agent in SF Oracle RAC installations. You must configure the CSSD agent after installing Oracle Clusterware/Grid Infrastructure.

Functions of the CSSD agent

[Table H-18](#) describes the functions of the CSSD agent.

Table H-18 CSSD agent functions

Function	Description
Online	Starts Oracle Clusterware.
Offline	Stops Oracle Clusterware.
Monitor	Checks the status of Oracle Clusterware.

Attributes of the CSSD agent

[Table H-19](#) lists the required attributes of the CSSD agent:

Table H-19 Required attributes for CSSD resource

Attribute Name	Required Value
Critical	0
StartProgram	/opt/VRTSvcs/rac/bin/cssd-online

Table H-19 Required attributes for CSSD resource (*continued*)

Attribute Name	Required Value
StopProgram	/opt/VRTSvcs/rac/bin/cssd-offline
CleanProgram	/opt/VRTSvcs/rac/bin/cssd-clean
MonitorProgram	/opt/VRTSvcs/rac/bin/cssd-monitor

States of the CSSD agent


[Table H-20](#) describes the states of the CSSD agent.

Table H-20 CSSD agent states

State	Description
Online	Indicates that Oracle Clusterware is running.
Offline	Indicates that Oracle Clusterware is not running.
Unknown	Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues.

Sample service group configurations with the CSSD agent

[Figure H-3](#) illustrates a basic service group configuration with the CSSD agent.

Figure H-3 Basic service group configuration with the CSSD agent

In this basic configuration, the OCR and voting disk volumes/mount points and IP addresses are configured under the CSSD resource. This ensures that these resources are online before the CSSD agent starts Oracle Clusterware.

Note: Depending on whether the Oracle database is started by Oracle Clusterware or by the VCS agent for Oracle, you must configure the Oracle database mounts such that they are online before the database starts.

For more sample configurations:

See [“Sample configuration files”](#) on page 637.

Type definition of the CSSD resource

The CSSD agent is an application agent. You can determine the name of the CSSD resource.

The following extract shows the type definition of the CSSD resource in the `types.cf` file.

```
type Application (
    static keylist SupportedActions =
    { "program.vfd", "user.vfd", "cksum.vfd", getcksum }
    static str ArgList[] =
    { User, StartProgram, StopProgram, CleanProgram,
      MonitorProgram, PidFiles, MonitorProcesses }
    static int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }
    str User
    str StartProgram
    str StopProgram
    str CleanProgram
    str MonitorProgram
    str PidFiles[]
    str MonitorProcesses[]
)
```

Sample configuration of the CSSD resource

The following extract from the `main.cf` file illustrates a sample CSSD agent configuration:

```
Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
)
```

VCS agents for Oracle

The VCS agents for Oracle include the following agents that work together to make Oracle highly available:

- The Oracle agent monitors the database processes.
See [“Oracle agent functions”](#) on page 707.
See [“Resource type definition for the Oracle agent”](#) on page 713.
- The Netlsnr agent monitors the listener process.
See [“Netlsnr agent functions”](#) on page 720.
- The ASMDG agent monitors the Oracle ASM disk groups.
See [“ASMDG agent functions”](#) on page 726.
See [“Resource type definition for the ASMDG agent”](#) on page 726.

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 H-21](#) lists the Oracle agent functions.

Table H-21 Oracle agent functions

Agent operation	Description
Online	Starts the Oracle database. The agent uses the default startup option <code>STARTUP_FORCE</code>. For RAC clusters, you must manually change the value of the <code>StartUpOpt</code> attribute to <code>SRVCTLSTART</code>. If you set the option to <code>SRVCTLSTART</code>, the agent uses the following <code>srvctl</code> command to start the Oracle database: <pre>srvctl start database -d database_name</pre>

Table H-21 Oracle agent functions (*continued*)

Agent operation	Description
Offline	Stops the Oracle database. The agent uses the default shutdown option IMMEDIATE. For RAC clusters, you must manually change the value of the ShutDownOpt attribute to SRVCTLSTOP. If you set the option to SRVCTLSTOP, the agent uses the following <code>srvctl</code> command to stop the Oracle database: <pre>srvctl stop database -d database_name</pre>
Monitor	Verifies the status of the Oracle processes. The Oracle agent provides two levels of monitoring: basic and detail. See “Monitor options for the Oracle agent” on page 709.
oracle_imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Oracle agent. This function runs when the agent starts up.
oracle_imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
oracle_imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).
Clean	Forcibly stops the Oracle database. If you set the shutdown option to SRVCTLSTOP, the agent uses the following <code>srvctl</code> command: <pre>srvctl stop database -d database_name</pre> If the process does not respond to the <code>srvctl</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

Table H-21 Oracle agent functions (*continued*)

Agent operation	Description
Info	Provides the static and dynamic information about the state of the database. See “Info entry point for SF Oracle RAC agent for Oracle” on page 710.
Action	Performs the predefined actions on a resource. See “Action entry point for SF Oracle RAC agent for Oracle” on page 711.

Monitor options for the Oracle agent

The Oracle agent provides two levels of monitoring: basic and detail. By default, the agent does a basic monitoring.

The basic monitoring mode has the following options:

- Process check
- Health check

The MonitorOption attribute of the Oracle resource determines whether the agent must perform basic monitoring in Process check or Health check mode.

[Table H-22](#) describes the basic monitoring options.

Table H-22 Basic monitoring options

Option	Description
0 (Default)	Process check The agent scans the process table for the ora_dbw0, ora_smon, ora_pmon, ora_lmon, and ora_lgwr processes to verify that Oracle is running. In this mode, the agent also supports intelligent resource monitoring.
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. The agent does not support intelligent resource monitoring in this mode.

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.

If Oracle processes are online and you attempt to bring online the Oracle database, then basic monitoring detects and reports concurrency violation by the Oracle database on the standby node. However, detail monitoring does not detect concurrency violation because Oracle processes are online but the Oracle database is reported offline. Detail monitoring only reports concurrency violation when both processes and database are online.

In the detail monitoring mode, the agent performs a transaction on a test table in the database to ensure that Oracle database functions properly. The agent uses this test table for internal purposes. Symantec recommends that you do not perform any other transaction on the test table.

Info entry point for SF Oracle RAC agent for Oracle

The supports the Info entry point, which provides static and dynamic information about the state of the database.

To invoke the Info entry point, type the following command:

```
# hares -value resource ResourceInfo [system] \  
[-clus cluster | -localclus]
```

The entry point retrieves the following static information:

■ Version	■ InstanceNo	■ InstanceName
■ DatabaseName	■ HostName	■ StartupTime
■ Parallel	■ Thread	■ InstanceRole

The entry point retrieves the following dynamic information:

- InstanceStatus ■ Logins ■ OpenMode
- LogMode ■ ShutdownPending ■ DatabaseStatus
- Shared Pool Percent free ■ Buffer Hits Percent

You can add additional attributes by adding sql statements to the file `/opt/VRTSagents/ha/bin/Oracle/resinfo.sql`. For example:

```
select 'static:HostName:'||host_name from v$instance;
select 'dynamic:ShutdownPending:'||shutdown_pending from
v$instance;
```

The format of the selected record must be as follows:

```
attribute_type:userkey_name:userkey_value
```

The variable *attribute_type* can take the value static and/or dynamic.

Action entry point for SF Oracle RAC agent for Oracle

The 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 H-24](#) on page 712. describes the agent's predefined virtual fire drill actions.

[Table H-23](#) describes the agent's predefined actions.

Table H-23 Predefined agent actions

Action	Description
VRTS_GetInstanceName	Retrieves the name of the configured instance. You can use this option for the Oracle and the Netlsnr resources.
VRTS_GetRunningServices	Retrieves the list of processes that the agent monitors. You can use this option for the Oracle and the Netlsnr resources.

Table H-23 Predefined agent actions (*continued*)

Action	Description
DBRestrict	Changes the database session to enable the RESTRICTED mode.
DBUndoRestrict	Changes the database session to disable the RESTRICTED mode.
DBSuspend	Suspends a database.
DBResume	Resumes a suspended database.
DBTbspBackup	Backs up a tablespace; <code>actionargs</code> contains name of the tablespace to be backed up.

[Table H-24](#) lists the virtual fire drill actions of the lets you run infrastructure checks and fix specific errors.

Table H-24 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.

Table H-24 Predefined virtual fire drill actions (*continued*)

Virtual fire drill action	Description
tnsadmin.vfd (Netlsnr agent)	Checks if listener.ora file is present. If the listener.ora file is not present, it checks if ORACLE_HOME is mounted and displays appropriate messages.

Resource type definition for the Oracle agent

The Oracle agent of the is represented by the Oracle resource type in SF Oracle RAC.

```

type Oracle (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Oracle"
    static keylist SupportedActions = { VRTS_GetInstanceName,
        VRTS_GetRunningServices, DBRestrict, DBUndoRestrict,
        DBResume, DBSuspend, DBTbspBackup, "home.vfd", "owner.vfd",
        "getid", "pfile.vfd" }
    static str ArgList[] = { Sid, Owner, Home, Pfile, StartUpOpt,
        ShutDownOpt, DBAUser, DBAPword, EnvFile, AutoEndBkup,
        User, Pword, Table, MonScript, AgentDebug,
        Encoding, MonitorOption }
    static str IMFRegList[] = { Home, Owner, Sid, MonitorOption }
    str Sid
    str Owner
    str Home
    str Pfile
    str StartUpOpt = STARTUP_FORCE
    str ShutDownOpt = IMMEDIATE
    str DBAUser
    str DBAPword
    str EnvFile
    boolean AutoEndBkup = 1
    str MonScript = "./bin/Oracle/SqlTest.pl"
    str User
    str Pword
    str Table
    boolean AgentDebug = 0
    str Encoding
    int MonitorOption = 0
    static boolean IntentionalOffline = 0
)

```

Attribute definition for the Oracle agent

Review the description of the Oracle agent attributes. The agent attributes are classified as required, optional, and internal.

[Table H-25](#) lists the required attributes. You must assign values to the required attributes.

Table H-25 Required attributes for Oracle agent

Required attributes	Type and dimension	Definition
Sid	string-scalar	The variable \$ORACLE_SID that represents the Oracle instance. The Sid is considered case-sensitive by the Oracle agent and by the Oracle database server. For a policy managed database, the Sid attribute should be set to Sid prefix. See “About the Sid attribute in a policy managed database” on page 719.
Owner	string-scalar	The Oracle user who has privileges to start or stop the database instance. The agent also supports LDAP users as Oracle user.
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path.

[Table H-26](#) lists the optional attributes for Oracle agent. You can configure the optional attributes if necessary.

Table H-26 Optional attributes for Oracle agent

Optional Attributes	Type and Dimension	Definition
DBAUser	string-scalar	The database user who has sysdba privileges to start or stop the database.
DBAPword	string-scalar	Encrypted password for DBAUser. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.

Table H-26 Optional attributes for Oracle agent (*continued*)

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 ■ SRVCTLSTART_RO ■ CUSTOM Default is STARTUP_FORCE.
ShutDownOpt	string-scalar	Shut down options for the Oracle instance. This attribute can take the following values: ■ IMMEDIATE ■ TRANSACTIONAL ■ SRVCTLSTOP ■ SRVCTLSTOP_TRANSACT ■ SRVCTLSTOP_ABORT ■ SRVCTLSTOP_IMMEDIATE ■ CUSTOM Default is IMMEDIATE.
EnvFile	string-scalar	The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on. 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.

Table H-26 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
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 The agent supports intelligent resource monitoring only when this attribute value is set to 0. Default = 0 See "Monitor options for the Oracle agent" on page 709.

Table H-26 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
IMF	integer-association	This resource-type level attribute determines whether the Oracle agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 3 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 5 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>oracle_imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3

Table H-26 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
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.
LevelTwoMonitorFreq	integer-scalar	Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. You can also override the value of this attribute at resource-level. The value indicates the number of monitor cycles after which the agent will monitor Oracle in detail. For example, the value 5 indicates that the agent will monitor Oracle in detail every five online monitor intervals. If you manually upgraded to the SF Oracle RAC 6.0 agent, and if you had enabled detail monitoring in the previous version, then do the following: Set the value of the LevelTwoMonitorFreq attribute to the same value of that of the DetailMonitor attribute. Note: If you set the AutoEndBkup attribute value to 0, then make sure that the LevelTwoMonitorFreq attribute value is 1 for detail monitoring. Default = 0
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 User/Pword.
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".

Table H-26 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
IntentionalOffline		This resource-type level attribute 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. See the <i>Veritas Cluster Server Administrator's Guide</i>.
DBName	string-scalar	Set this attribute only when the database is a policy managed RAC database. The value of this attribute must be set to the database unique name.
ManagedBy	string-scalar	Default value for this attribute is ADMIN. In a policy managed RAC database this attribute must be set to POLICY.

[Table H-27](#) lists the internal attribute for Oracle agent. This attribute is for internal use only. recommends not to modify the value of this attribute.

Table H-27 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.

About the Sid attribute in a policy managed database

The SID attribute is a required attribute. This section provides information to define the SID attribute in a policy managed database.

The SID prefix comprises of the first 8 alphanumeric characters of the database unique name. It can be a combination of letters a-z; uppercase and lowercase and numbers 0-9.

The SID prefix cannot have operating system special characters. Therefore, avoid the use of special characters in the first 8 characters of the database unique name. Special characters are omitted if used in the first 8 characters. There is a single SID prefix for every database. The SID prefix for a database must be unique within the cluster.

For an Oracle RAC database, each instance has a unique identifier, `ORACLE_SID`, which consists of the SID prefix and an instance number. The `ORACLE_SID` for Oracle RAC database instances is generated differently, depending on how you choose to manage the database. If you select a policy-managed database, then Oracle generates the SID in the format `name_#`, where `name` is the first eight alphanumeric characters of `DB_UNIQUE_NAME`, and `#` is the instance number. If you select an admin-managed database, then DBCA generates the SID for the instance names in advance, and the SID is in the format `name#`.

To find the Sid prefix name, run the following command:

```
# ${GRID_HOME}/bin/crsctl status resource ora.${DBName}.db -f | grep  
GEN_USR_ORA_INST_NAME@ | tail -1 | sed 's/.*=/' | sed 's/_[0-9]$//',
```

where `GRID_HOME` is grid home path and `DBName` is the database unique name.

Note: When a policy managed database is created, the Sid prefix is displayed on the confirmation page of the installation procedure.

See [“Attribute definition for the Oracle agent”](#) on page 714.

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.

The Netlsnr agent is IMF-aware.

[Table H-28](#) lists the Netlsnr agent functions.

Table H-28 Netlsnr agent functions

Agent operation	Description
Online	Starts the listener process by using the following command: <code>lsnrctl start \$LISTENER</code>

Table H-28 Netlsnr agent functions (*continued*)

Agent operation	Description
Offline	Stops the listener process by using the following command: <pre>lsnrctl stop \$LISTENER</pre> 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)
netlsnr_imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Netlsnr agent. This function runs when the agent starts up.
netlsnr_imf_getnotification	Gets notification about resource state change. This function runs after the agent registers with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
netlsnr_imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).
Clean	Scans the process table for <code>tnslsnr \$LISTENER</code> and kills it.
Action	Performs the predefined actions on a resource. See “Action entry point for SF Oracle RAC agent for Oracle” on page 711.

Resource type definition for the Netlsnr agent

The Netlsnr agent of the 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, Encoding }  
static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }  
static str IMFRegList[] = { Home, Owner, Listener }  
str Owner  
str Home  
str TnsAdmin  
str Listener  
str EnvFile  
str MonScript = "./bin/Netlsnr/LsnrTest.pl"  
str LsnrPwd  
str Encoding  
static boolean IntentionalOffline = 0  
)
```

Attribute definition for the Netlsnr agent

Review the description of the Netlsnr agent attributes. The agent attributes are classified as required, optional, and internal.

[Table H-29](#) lists the required attributes for Netlsnr agent. You must assign values to the required attributes.

Table H-29 Required attributes for Netlsnr agent

Required attributes	Type and dimension	Definition
Owner	string-scalar	The Oracle user who has privileges to start or stop the listener process. The agent also supports LDAP users as Oracle user.
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Do not append a slash (/) at the end of the path.

[Table H-30](#) lists the optional attributes for Netlsnr agent. You can configure the optional attributes if necessary.

Table H-30 Optional attributes for Netlsnr agent

Optional attributes	Type and dimension	Definition
TnsAdmin	string-scalar	The \$TNS_ADMIN path to directory in which the Listener configuration file resides (listener.ora). 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.
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.

Table H-30 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
IMF	integer-association	This resource-type level attribute determines whether the Netlsnr agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 3 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 5 You can set this attribute to a non-zero value in some cases where the agent requires to perform poll-based resource monitoring in addition to the intelligent resource monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the netlsnr_imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3

Table H-30 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
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 <code>/opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl</code>. MonScript also accepts a pathname relative to <code>/opt/VRTSagents/ha</code>. A relative pathname should start with <code>"/</code>, as in the path <code>./bin/Netlsnr/LsnrTest.pl</code>.
LevelTwoMonitorFreq	integer-scalar	Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. If you enabled detail monitoring, then set the value of the LevelTwoMonitorFreq attribute. Default = 0
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is <code>""</code>.
IntentionalOffline		For future use. Do not change the value of this attribute. Default = 0

[Table H-31](#) lists the internal attribute for Netlsnr agent. This attribute is for internal use only. recommends not to modify the value of this attribute.

Table H-31 Internal attributes for Netlsnr agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	Specifies the location of binaries, scripts, and other files related to the Netlsnr agent. Default is <code>/opt/VRTSagents/ha/bin/Netlsnr</code>.

ASMDG agent functions

The ASMDG agent mounts the ASM disk groups that the Oracle databases use, monitors the status, unmounts the ASM disk groups.

You must have specified the disk group names in the DiskGroups attribute of the ASMDG agent.

[Table H-32](#) lists the ASMDG agent operations.

Table H-32 ASMDG agent operations

Agent operation	Description
Online	Mounts the specified Oracle ASM disk groups to an ASM instance by using the following SQL command: <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>

Resource type definition for the ASMDG agent

The ASMDG agent of the is represented by the ASMDG resource type in SF Oracle RAC.

```
type ASMDG (  
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/ASMDG"  
    static str ArgList[] = { Sid, Owner, Home, DBAUser,  
        DBAPword, DiskGroups, EnvFile, Encoding }  
    str Sid  
    str Owner  
    str Home  
    str DBAUser  
    str DBAPword  
    keylist DiskGroups  
    str EnvFile  
    str Encoding  
)
```

Attribute definition for the ASMDG agent

Review the description of the ASMDG agent attributes. The agent attributes are classified as required, optional, and internal.

[Table H-33](#) lists the required attributes. You must assign values to the required attributes.

Table H-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.
Owner	string-scalar	The Oracle user who has privileges to mount or unmount the ASM disk group. The agent also supports LDAP users as Oracle user.
Home	string-scalar	The \$ORACLE_HOME path to Oracle ASM binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path.

[Table H-34](#) lists the optional attributes for ASMDG agent. You can configure the optional attributes if necessary.

Table H-34 Optional attributes for ASMDG agent

Optional Attributes	Type and Dimension	Definition
DBAUser	string-scalar	The ASM user who has sysasm privileges to start or stop the ASM instance. You can create ASM users for Oracle 11g R1 and later.
DBAPword	string-scalar	Encrypted password for DBAUser. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.
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 H-35](#) lists the internal attribute for ASMDG agent. This attribute is for internal use only. recommends not to modify the value of this attribute.

Table H-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.

CRSResource agent

The `CRSResource` agent provides an alternative mechanism for monitoring the Oracle database in the absence of the VCS Oracle agent. It checks the status of the Oracle Clusterware resources, which include the Oracle database instance, the listener, and the virtual IP address (VIP). The agent supports multiple database configurations and ensures that the Oracle database is online and available to an application when it starts.

Note: For Oracle RAC 11g Release 2: The `CRSResource` agent is supported only in admin-managed database environments.

Functions of the CRSResource agent

[Table H-36](#) describes the functions of the CRSResource agent.

Table H-36 CRSResource agent entry points

Entry point	Description
Monitor	Checks the status of the following Oracle Clusterware resources: ■ Oracle database ■ Virtual IP address ■ Listener The monitor script checks the status of the resources every 60 seconds.
Online	A dummy entry point that uses a delay interval to ensure that the Oracle Clusterware resources come online before the monitor operation begins.

States of the CRSResource agent

Table H-37 lists the states of the CRSResource agent.

Table H-37 States of the CRSResource agent

State	Description
Online	Indicates that the Oracle database, VIP, and listener are running.
Faulted	Indicates that the Oracle database, VIP, and listener are not running or are faulted.
Unknown	Indicates the inability to determine the state of the resource due to incorrect attributes settings or other configuration issues.

Attributes of the CRSResource agent

Table H-38 Attributes of the CRSResource agent

Attribute	Dimension	Description
ResType	string-scalar	Indicates the type of the Oracle Clusterware resource. The types of resources are as follows: ■ DB ■ VIP ■ Listener
DBHome	string-scalar	The path to ORACLE_HOME containing the Oracle binaries and configuration files. For example, you could specify the path as /app/oracle/orahome. Note: Do not append a slash (/) at the end of the path.
CRSHome	string-scalar	The path to CRS_HOME containing the Oracle Clusterware binaries.
DBName	string-scalar	The name of the database on which the services are configured. This attribute is optional if you are configuring CRSResource to monitor VIP or listener.

Table H-38 Attributes of the CRSResource agent (*continued*)

Attribute	Dimension	Description
SID	string-scalar	The name of the Oracle instance in the variable \$ORACLE_SID. The SID is case-sensitive. This attribute is optional if you are configuring CRSResource to monitor VIP or listener.
Owner	string-scalar	The name of the user for the Oracle database. This attribute is optional if you are configuring CRSResource to monitor VIP or listener.

VCS service group dependencies with the CRSResource agent

In a service group configuration with the CRSResource agent, Oracle Clusterware controls the database. An online local firm dependency exists between the groups—Application group, Oracle Clusterware group, and the CVM group.

[Figure H-4](#) shows a schematic illustration of the service group dependencies.

Figure H-4 Service group dependencies with CRSResource agent



In the configuration:

- When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware, and starts Oracle Clusterware. Oracle Clusterware starts the database and the application is brought online. CRSResource comes online when the Oracle Clusterware resources (database/VIP/listener) are started by Oracle Clusterware.

Note: When the system starts, all volumes and mount points **MUST** be online for the dependent service groups to be online.

- The oradata_mnt and oradata_voldg resources are configured as non-critical resources (critical=0) for managing failure scenarios.
See [“How CRSResource agent handles failures”](#) on page 733.
- When CRSResource faults for any of the Oracle Clusterware resources, the application is brought offline.

The limitations of this configuration are as follows:

- The CFSSMount and CVMVolDg resources can not be set as critical resources in the group.
If the mount points and volume disk groups for all the databases are configured as critical in a single service group, then failure of any of them results in the whole group being FAULTED or brought offline. To ensure that a resource failure does not affect other resources in the group, the attribute Critical is set to zero for the CFSSMount and CVMVolDg resources.
However, if any of the database mounts fail to come online or a volume does not start, the whole service group fails to come online.
- CRSResource reports as FAULTED until Oracle Clusterware brings up the database instance, VIP, and listener. Even after Oracle Clusterware starts the database instance, VIP and listener, CRSResource remains in the FAULTED state for the OfflineMonitorInterval period. The status of CRSResource cannot be changed.


How CRSResource agent handles failures

The CRSResource agent ensures that faults in the resources of an application do not adversely impact other applications running on the system. To isolate failures, the oradata_mnt and oradata_voldg resources are configured as non-critical resources (critical=0). This ensures that storage issues in either of these resources do not affect the other databases and the dependent application continues to be online.

Note: The resources are considered non-critical only for the purpose of managing failure scenarios.

[Figure H-5](#) illustrates a failure scenario.

Figure H-5 How CRSResource agent handles failures



Fault configurations with CRSResource agent

This section discusses scenarios that cause CRSResource to report FAULTED.

Scenario 1: CRSResource fault at system startup

1. The CVM group brings online the volume and mount points for the databases when the system starts.
2. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware and starts Oracle Clusterware.


Oracle Clusterware is unable to bring the database or other resources online within the start duration configured in the agent.

3. The agent starts the monitoring operation after the set duration.

If the Oracle Clusterware resources are not brought online yet, the CRSResource appears faulted.

Figure H-6 illustrates the scenario.

Figure H-6 CRSResource fault at system startup



Scenario 2: CRSResource fault when resource is brought offline

Any or all of the Oracle Clusterware resources (database, listener, or VIP) are brought offline due to a fault or for administrative purposes.

CRSResource reports FAULTED until Oracle Clusterware brings the resources online.

Resource type definition for the CRSResource agent

The following extract from the `CRSResource.cf` file describes the type definition for the CRSResource agent.

```
type CRSResource (
    static int MonitorTimeout = 300
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { ResType, DBHome, CRSHome,
                           DBName, SID, Owner }
    static str Operations = None
    str ResType
    str DBHome
    str CRSHome
    str DBName
    str SID
    str Owner
)
```

Sample configuration for the CRSResource agent

The following sample configuration describes the DB, VIP, and Listener resource configurations for the CRSResource agent.

```
CRSResource crsr_oradb_db (
    ResType = DB
    DBHome = "/app/oracle/orahome"
    CRSHome = "/app/crshome"
    DBName = oradb
    SID @galaxy = oradb1
    SID @nebula = oradb2
    Owner = oracle
)

CRSResource crsr_oradb_vip (
    ResType = VIP
    DBHome = "/app/oracle/orahome"
    CRSHome = "/app/crshome"
```



```
)  
CRSResource crsr_oradb_lsnr (  
    ResType = Listener  
    DBHome = "/app/oracle/orahome"  
    CRSHome = "/app/crshome"  
)
```

SF Oracle RAC deployment scenarios

This appendix includes the following topics:

- [SF Oracle RAC cluster with UDP IPC and PrivNIC agent](#)
- [SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent](#)
- [SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent](#)
- [SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent](#)
- [Configuration diagrams for setting up server-based I/O fencing](#)
- [Deploying Storage Foundation for Databases \(SFDB\) tools in a Storage Foundation for Oracle RAC environment](#)

SF Oracle RAC cluster with UDP IPC and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:

Deployment scenario	Oracle RAC configured with UDP IPC for database cache fusion.
Recommendation	Use the PrivNIC agent.
Sample main.cf configuration file	The following sample main.cf file describes the configuration: <code>/etc/VRTSvcs/conf/sample_rac/sfrac02_main.cf</code>


In the illustrated configuration:

- A common IP address is used for Oracle Clusterware communication and Oracle database cache fusion.

- Oracle Clusterware communication and Oracle database cache fusion traffic flows over one of the LLT private interconnect links.
- The CFS/CVM/VCS metadata travels through LLT over the two physical links that are configured as LLT links.
- In the event of a link failure, the PrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure I-1 illustrates the logical view of a two-node SF Oracle RAC cluster with UDP IPC and PrivNIC agent.

Figure I-1 SF Oracle RAC cluster with UDP IPC and PrivNIC agent



Legends

Public link (GigE)	Oracle inter-process communication	↔
Private Interconnect (GigE)	Active connection	↔
Oracle Clusterware Heartbeat	Failover connection for	↔
Oracle DB Cache Fusion	PrivNIC Agent	↔
CFS /CVM / VCS Metadata		

SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent

This section illustrates the recommended configuration for the following scenario:


Deployment scenario	Oracle RAC is configured with UDP IPC for database cache fusion.
Recommendation	Use the MultiPrivNIC agent.
Sample main.cf configuration file	The following sample main.cf file describes the configuration: <code>/etc/VRTSvcs/conf/sample_rac/sfrac03_main.cf</code>

In the illustrated configuration:

- One private IP address is used for each database for database cache fusion. One of the private IP addresses used for the database cache fusion is shared by Oracle Clusterware communication. The CFS/CVM/VCS metadata also travels through these links.
- In the event of a link failure, the MultiPrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure I-2 illustrates the logical view of a two-node SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent.

Figure I-2 SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent



SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent

This section illustrates the recommended configuration for the following scenario:


Deployment scenario	Oracle RAC database cache fusion traffic is isolated from the CFS/CVM/VCS metadata.
Recommendation	Use the MultiPrivNIC agent.
Sample main.cf configuration file	The following sample main.cf file describes the configuration: <code>/etc/VRTSvcs/conf/sample_rac/sfrac06_main.cf</code>

In the illustrated configuration:

- The private network IP address used for database cache fusion is configured over a dedicated link for each database. These links are configured as low-priority LLT links.
- The private network IP address used for Oracle Clusterware communication is configured over a high-priority LLT link. This link is also used for the CFS/CVM/VCS metadata transfer.
- In the event of a link failure, the MultiPrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure I-3 illustrates the logical view of a two-node SF Oracle RAC cluster with Oracle traffic isolated from VCS / CVM / CFS traffic.

Figure I-3 SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent



SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:


Deployment scenario	Oracle RAC with UDP IPC is configured to use a bonded NIC interface for distribution of Oracle database cache fusion traffic. A second link is configured as a standby link.
Recommendation	Use the PrivNIC agent.
Sample main.cf configuration file	The following sample main.cf file describes the configuration: <pre>/etc/VRTSvcs/conf/sample_rac/sfrac02_main.cf</pre> Note: You must replace the eth1 interface in the sample file with the bonded NIC interface you use in the scenario and eth2 in the sample file with the NIC3 interface in the scenario.

In the illustrated configuration:

- A common IP address is used for Oracle Clusterware communication and Oracle database cache fusion that is distributed over two underlying physical links of the bonded NIC interface. The bonded NIC interface is configured as a single LLT link.
- The CFS/CVM/VCS metadata also travels through the bonded link.
- In the event of a bonded link failure, the PrivNIC agent fails over the private network IP address from the failed link to the available standby LLT link. When the link is restored, the IP address is failed back to the original link.

Figure I-4 illustrates the logical view of a two-node SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent.

Figure I-4 SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent



Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:


- Two unique client clusters that are served by 3 CP servers:
See [Figure I-5](#) on page 748.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
[Figure I-6](#)
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
[Figure I-7](#)
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:
[Figure I-8](#)

Two unique client clusters served by 3 CP servers

[Figure I-5](#) displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

Figure I-5 Two unique client clusters served by 3 CP servers




Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure I-6 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfermode` file on the client nodes, `vxfermode` is set to `vxfer` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group vxsfencoorddg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-6 Client cluster served by highly available CP server and 2 SCSI-3 disks




Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure I-7 displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group `vxfencoorddg`. The third coordination point is a CP server on a single node VCS cluster.

Figure I-7 Two node campus cluster served by remote CP server and 2 SCSI-3




Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

Figure I-8 displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-8 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks




Deploying Storage Foundation for Databases (SFDB) tools in a Storage Foundation for Oracle RAC environment

If you are deploying the SFDB tools with Storage Foundation for Oracle RAC (multiple instance Oracle) your setup configuration will reflect the following conditions:

- A highly available parallel cluster with a multiple instances of Oracle is set up on galaxy and nebula with SF for Oracle RAC.
- The database is online on galaxy, nebula, and mercury.
- The datafiles are mounted and shared on galaxy, nebula, and mercury.
- The SFDB tools is mounted and shared on galaxy, nebula, and mercury.
- You can run the SFDB tools commands on galaxy, nebula, and mercury.
- Clustered ODM is supported for this configuration.

In the figure below the repository directory resides in the Oracle mount points.

Figure I-9 Deploying the database repository with Storage Foundation



For an SF Oracle RAC configuration, the systems are online in parallel and do not use failover mechanisms within the cluster.

Compatibility issues when installing Veritas Storage Foundation for Oracle RAC with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present](#)

Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present

Installing Storage Foundation when other Veritas products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host RPMs as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present

If you plan to install or upgrade Storage Foundation on systems where ApplicationHA has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not consider VCS as an installed product even though it uses the bundled VRTSvcS RPM.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not allow the installation or upgrade for products that use VCS. The following products cannot be installed or upgraded: VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SFCFSRAC, or SFSYBASECE.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer allows the installation or upgrade of VM, FS, SF, or DMP.
- When you uninstall Storage Foundation products where ApplicationHA is present, the installer does not uninstall VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx, VRTSicsco, and VRTSat.

Index

Symbols

\$CRS_HOME
 setting up 312
\$ORACLE_HOME
 setting up 312

A

adding
 users 127
agent operations
 Netlsnr agent 720
 Oracle agent 707
agents
 about 676
 CFSfsckd 690
 CFSMount 686, 690
 CSSD 703
 CVMCluster 678
 CVMVolDg 683
 CVMVxconfigd 680
 log files 365
 MultiPrivNIC 699
 of VCS 677
 of VCS for Oracle 677
 PrivNIC 693
ASM
 about 664
 configuring with CVM
 about 666
 configuring disk group 667
 creating shared raw volumes 668
 main.cf 672
 SF Oracle RAC configuration 665
 unlinking ODM 667
 verifying setup 669
ASMDG agent
 attribute definitions 727
 resource type 726
ASMDG agent attributes
 AgentDirectory 727
 DBAPword 727

ASMDG agent attributes *(continued)*
 DBAUser 727
 DiskGroups 727
 Encoding 727
 EnvFile 727
 Home 727
 Owner 727
 Sid 727
attribute definitions
 ASMDG agent 727
 Netlsnr agent 722
 Oracle agent 714
attributes
 about agent attributes 677
 CFSMount agent 687, 691
 CVMCluster agent 679
 CVMVolDg agent 679, 684
 CVMVxconfigd agent 681
 MultiPrivNIC agent 700
 PrivNIC agent
 optional 696
 required 694
 UseFence 151
Automatic Storage Management. *See* ASM

B

backup boot disk group 255
 rejoining 255
basic monitoring 709
 health check 709
 process 709

C

cables
 cross-over Ethernet 438
CFS
 creating databases 675
 stopping applications 569
 unmounting file systems 569
CFSfsckd agent 690
 attributes 691

- CFSMount agent 686, 690
 - attributes 687
 - entry points 686
 - sample configuration 689–690
 - type definition 689
- CFSTypes.cf 689
- cluster
 - removing a node from 497
 - verifying operation 175
- cluster configuration 117
- Cluster File System. *See* CFS
- Cluster Manager 31
- Cluster Volume Manager. *See* CVM
- clusters
 - basic setup 33
 - four-node illustration 34
- commands
 - gcoconfig 525
 - hastatus 175
 - hasys 176
 - lltconfig 632
 - lltstat 172
 - vradmin 540
 - vxassist 532, 535, 675
 - vxdisksetup (initializing disks) 140
 - vxedit 674
 - vxvol 532, 675
- configuration
 - required information
 - for global clusters 612
 - for I/O fencing 610
 - for VCS 609
- configuration files
 - removing 579
 - See also* main.cf samples
- configuring
 - rsh 661
 - ssh 661–662
- Configuring global clusters
 - for SF Oracle RAC 515
- configuring VCS
 - adding users 127
 - event notification 128, 130
 - global clusters 132
 - secure mode 122
- coordinator disks
 - setting up 149
- CSSD agent 703
 - attributes 705

- CVM
 - configuration worksheet 625
 - CVMTypes.cf file 680
 - installation worksheet 625–626
 - upgrading protocol version 257
- CVMCluster agent 678
 - attributes 679
 - entry points 678
 - sample configuration 680
 - type definition 680
- CVMTypes.cf
 - definition, CVMCluster agent 680
 - definition, CVMVolDg agent 685
 - definition, CVMVxconfigd agent 683
- CVMVolDg agent 683
 - attributes 684
 - entry points 683
 - sample configuration 686
 - type definition 685
- CVMVxconfigd agent 680
 - attributes 681
 - CVMTypes.cf 683
 - entry points 681
 - sample configuration 683
 - type definition 683

D

- databases
 - creating 673
 - creating tablespaces
 - on CFS 675
 - on raw volumes 674
- detail monitoring 709
- disks
 - adding and initializing 140
 - coordinator 149
 - testing with vxfcntlshdw 141
 - verifying node access 143

E

- environment variables
 - MANPATH 76
- Ethernet controllers 438

F

- files. *See* configuration files

G

- GAB
 - port memberships 174
- GAB ports 174
- gabtab file
 - verifying after installation 632
- gcoconfig command 525
- global cluster
 - return of primary role to original primary site 558
 - takeover of primary site role by secondary 558
 - troubleshooting VVR component 561
- Global Cluster Option (GCO)
 - overview 39
- global clusters
 - about 30
 - configuration 132
 - using VVR 524
 - configuration examples 543
 - configuring VCS service groups 524
 - illustration of dependencies 541
 - migration of primary to secondary site 557
 - overview 515
 - replicating database volumes 541
 - required information
 - for configuration 612
 - requirements 515
 - return migration to original primary site 557
- groups
 - log owner 542
 - RVG 542

H

- hastatus -summary command 175
- hasys -display command 176
- health check APIs 709
- health check monitoring 709
- hubs
 - independent 438

I

- I/O fencing
 - checking disks 141
 - required information
 - for configuration 610
 - setting up 149
 - shared storage 141
- installation
 - of Oracle Clusterware 326, 329

- installation (*continued*)
 - of Oracle Grid Infrastructure 326, 329
 - of Oracle RAC 265
 - preparation 605
- installation worksheets 605
- installing
 - yum 93
- installsrac
 - installing Clusterware 326, 329
 - installing Oracle Grid Infrastructure 326, 329
 - installing Oracle RAC 269
 - installing SF Oracle RAC 82
 - upgrading SF Oracle RAC 199

J

- Java Console 31

L

- links
 - private network 632
- LLT
 - interconnects 78
 - verifying 172
- lltconfig command 632
- llthosts file
 - verifying after installation 632
- lltstat command 172
- llttab file
 - verifying after installation 632

M

- main.cf
 - after SF Oracle RAC installation 631
- main.cf files 654
- main.cf samples
 - for CSSD resource 706
- MANPATH environment variable 76
- media speed 78
 - optimizing 78
- migration, of Oracle 367
- monitoring
 - basic 709
 - detail 709
- MultiPrivNIC
 - for UDP 307
- MultiPrivNIC agent 699
 - attributes 700
 - entry point 700

MultiPrivNIC agent *(continued)*
 type definition 702
 MultiPrivNIC.cf 702
 sample configuration 703

N

Netlsnr agent
 attribute definitions 722
 operations 720
 resource type 721
 Netlsnr agent attributes
 AgentDebug 722
 AgentDirectory 722
 EnvFile 722
 Home 722
 Listener 722
 LsnrPwd 722
 MonScript 722
 Owner 722
 TnsAdmin 722
 nodes
 adding new nodes
 configuring GAB 447
 configuring LLT 447
 configuring VCSMM 447
 configuring VXFEN 447
 adding Oracle RAC nodes
 configuring GAB 445, 456
 configuring LLT 445, 456
 configuring VCSMM 445, 456
 configuring VXFEN 445, 456
 starting Volume Manager 444, 456
 preparing Oracle RAC nodes
 configuring CVM 452, 465
 configuring private IP addresses 484
 preparing \$CRS_HOME 486
 preparing \$ORACLE_HOME 486
 removing nodes
 GAB configuration 499
 LLT configuration 499
 removing nodes from a cluster 497
 removing nodes from VCS configuration files 498
 removing Oracle RAC 11g nodes
 modifying VCS configuration 500
 tasks for removing nodes 496

O

OCR
 creating
 on raw volumes 287
 creating file system
 using CLI 287
 operations
 Netlsnr agent 720
 Oracle agent 707
 optimizing
 media speed 78
 options
 SF Oracle RAC configuration 114
 Oracle
 agent log file 365
 creating databases 674–675
 stopping instances 566
 supported versions 48
 Oracle 11g
 migration from Oracle RAC 10g
 installing Oracle RAC 10g 369
 post-upgrade tasks 370
 Oracle 11g Release 2
 migration from Oracle RAC 10g
 of existing database 370
 pre-upgrade tasks 368
 Oracle agent
 attribute definitions 714
 operations 707
 resource type 713
 Oracle agent attributes
 AgentDebug 714
 AgentDirectory 714
 AutoEndBkup 714
 DBAPword 714
 DBAUser 714
 EnvFile 714
 Home 714
 IMF 714
 MonitorOption 714
 MonScript 714
 Owner 714
 Pfile 714
 Pword 714
 ShutDownOpt 714
 Sid 714
 StartupOpt 714
 Table 714
 User 714

- Oracle Grid Infrastructure
 - installing 326, 329
- Oracle RAC
 - database replication 539
 - installation
 - installing database 335, 337
 - workflow 324
 - manual installation
 - installing database binaries 340
 - installing Oracle Clusterware 332
 - installing Oracle Grid Infrastructure 332
 - linking ODM library 354
 - linking Veritas Membership library 354
 - pre-installation
 - creating OCR 286
 - creating Oracle disk groups 275
 - creating Oracle users and groups 271
 - creating voting disk volumes 286
 - identifying public virtual IP addresses 268
 - setting system parameters 268–269
 - setting up Oracle user equivalence 320
 - setting up storage 54
 - starting installsrac 269
 - workflow 265
- Oracle RAC 11g
 - agent log file 365
 - configuring service groups
 - preventing automatic starting of database 365
 - installation
 - configuring service groups 356
 - setting environment variables 327, 329
 - verifying 341
 - post-installation
 - adding patches 343
 - creating database 356
 - relinking SF Oracle RAC libraries 350
 - pre-installation
 - configuring MultiPrivNIC 307
 - preparing \$CRS_HOME 312
 - preparing \$ORACLE_HOME 312
 - setting environment variables 327, 329
- Oracle RAC 11g Clusterware
 - verifying installation 341
- Oracle RAC 11g Release 1 Clusterware
 - creating file system 488

P

- patches
 - adding 343
- patchsets
 - adding 343
- PATH variable
 - VCS commands 171
- ports
 - GAB 174
- primary site
 - creating SRL volume 532
 - setting up 517
 - setting up replication objects 533
 - VCS configuration 546, 550
- PrivNIC agent 693
 - attributes 694
 - entry point 693
 - sample configuration 698
 - type definition 698
- PrivNIC.cf 698
- process monitoring 709

R

- raw volumes
 - creating
 - Oracle DBCA procedure 668
- rejoining
 - backup boot disk group 255
- removing a node from a cluster 497
- replication
 - automatic synchronization 539
 - configuring on both sites 523
 - full synchronization with Storage Checkpoint 540
 - modifying VCS configuration 542
 - options 517
 - setting up primary site 517
 - setting up secondary site 519
 - supported hardware 515
 - supported software 515
 - using VVR 531
 - verifying status 540
- resource type
 - ASMDG 726
 - Netlsnr 721
 - Oracle 713
- resources
 - CVMVolDg 542
 - RVGSharedPri 542

- response file
 - about 373
 - syntax 376
- rolling upgrade 224
 - versions 219
- rsh 662
 - configuration 661

S

- SCSI-3 persistent reservations
 - verifying 149
- secondary site
 - configuring replication 534
 - creating SRL volume 535
 - setting up 519
 - setting up disk groups 537
 - setting up RLINKs 536
- secure clusters
 - adding users 611
- service groups
 - configuring 356
 - VCS, for global clusters 524
- setting umask, before installing 75
- SF Oracle RAC
 - about 26
 - coordinator disks 149
 - high-level view 33
 - illustration of global cluster dependencies 541
- SF Oracle RAC configuration
 - about 114
 - configuring clusters 117
 - creating configuration files 133
 - of components 115
 - options 114
 - preparation
 - worksheets 605
 - starting processes 133
- SF Oracle RAC installation
 - pre-installation tasks
 - mounting product disc 76
 - setting MANPATH 76
 - setting umask 75
 - synchronizing time settings 75
 - verifying systems 78
 - preinstallation information 41
 - preparation
 - worksheets 605
 - requirements
 - hardware 42
 - SF Oracle RAC installation *(continued)*
 - requirements *(continued)*
 - software 48
 - using installsfrac 82
 - verifying
 - cluster operations 171
 - GAB operations 171
 - LLT operations 171
 - SF Oracle RAC uninstallation
 - preparation
 - stopping applications, CFS 569
 - stopping applications, VxFS 570
 - stopping Oracle instances 566
 - stopping VCS 569
 - uninstalling 567
 - unmounting CFS file systems 569
 - unmounting VxFS file systems 570
 - removing configuration files 579
 - removing RPMs 577
 - using uninstallfrac program 577
 - workflow 563
- SF Oracle RAC upgrade
 - post-upgrade tasks
 - relinking Oracle RAC libraries 254
 - upgrading CVM protocol version 257
 - upgrading disk group version 257
 - preparation 194, 222
 - restoring configuration files 200
 - stopping cluster resources 194, 222
 - stopping Oracle RAC 194, 222
 - upgrading licenses 232
 - using installsfrac 199
- SFDB authentication 186
 - adding nodes 490
 - configuring vxdbd 186
- SMTP email notification 128
- SNMP trap notification 130
- ssh 662
 - configuration 661
 - configuring 662
- Storage Foundation for Oracle RAC configuration
 - verifying 174
- Symantec Product Authentication Service 122
- synchronizing time settings, before installing 75
- system communication using rsh
 - ssh 662
- system state attribute value 175

T

- tunables file
 - about setting parameters 593
 - parameter definitions 598
 - preparing 597
 - setting for configuration 594
 - setting for installation 594
 - setting for upgrade 594
 - setting parameters 597
 - setting with no other operations 595
 - setting with un-integrated response file 596

U

- uninstallation
 - of SF Oracle RAC 572
- uninstallsfrac program
 - removing RPMs 577
- unsuccessful upgrade 255
- upgrading
 - rolling 224
- users
 - creating 271
 - for &ProductNameShort;, adding 611

V

- VCS
 - agent log file 365
 - command directory path variable 171
 - configuration, for database volume replication 541
 - configuring service groups 524
 - notifications 30
 - required information
 - for configuration 609
 - stopping 569
- VCS configuration
 - for replication 542
- VCS Global cluster option. *See* GCO
- VCS notifications
 - SMTP notification 30
 - SNMP notification 30
- Veritas File System
 - stopping applications 570
 - unmounting 570
- Veritas Operations Manager 31
- Veritas Volume Replicator. *See* VVR
- virtual fire drill 712
- voting disk
 - creating
 - on raw volumes 287
- voting disks
 - creating file system
 - using CLI 287
- vradm command 540
- VVR
 - about 31
 - configuration examples 543
 - configuration worksheet 625
 - configuring global clusters 524
 - configuring on both sites 523
 - database volume replication
 - configuring VCS 541
 - defining heartbeat cluster objects 526
 - defining remote clusters 526
 - global cluster overview 530
 - installation worksheet 625–626
 - Oracle RAC database replication 539
 - primary site
 - creating SRL volume 532
 - setting up replication objects 533
 - replication
 - using automatic synchronization 539
 - using full synchronization with Storage Checkpoint 540
 - replication agents 517
 - secondary site
 - configuring replication 534
 - creating SRL volume 535
 - setting up disk groups 537
 - setting up RLINKs 536
 - setting up primary site 517
 - setting up replication 531
 - setting up secondary site 519
 - types of replication 517
 - VCS configuration
 - CVMoldG resource 542
 - log owner group 542
 - Oracle RAC database service group 542
 - primary site 546
 - RVG group 542
 - RVGSharedPri resource 542
 - secondary site 550
 - verifying replication 540
- vxassist command 532, 535, 675
- vxdisksetup command 140
- VxFS. *See* Veritas File System

vxvol command 532, 675

W

worksheets

- for &ProductNameShort; 606

- for CVM 625

- for replication

 - SRDF 626

 - VVR 625

- for VVR 625

Y

yum

- installing 93

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 Multi-Pathing)	A feature designed to provide greater reliability and performance by using path failover and load balancing for multiported disk arrays connected to host systems through multiple paths. DMP detects the various paths to a disk using a mechanism that is specific to each supported array type. DMP can also differentiate between different enclosures of a supported array type that are connected to the same host system.
SmartTier	A feature with which administrators of multi-volume VxFS file systems can manage the placement of files on individual volumes in a volume set by defining placement policies that control both initial file location and the circumstances under which existing files are relocated. These placement policies cause the files to which they apply to be created and extended on specific subsets of a file system's volume set, known as placement classes. The files are relocated to volumes in other placement classes when they meet specified naming, timing, access rate, and storage capacity-related conditions.

Failover	A failover occurs when a service group faults and is migrated to another system.
GAB (Group Atomic Broadcast)	A communication mechanism of the VCS engine that manages cluster membership, monitors heartbeat communication, and distributes information throughout the cluster.
HA (high availability)	The concept of configuring the product to be highly available against system failure on a clustered network using Symantec 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 <code>vxdisk</code> 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.
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 (Symantec 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 Symantec Storage Foundation product suite that provides high performance and online management capabilities to facilitate the creation and maintenance of file systems. A file system is a collection of directories organized into a structure that enables you to locate and store files.
VxVM (Veritas Volume Manager)	A Symantec product installed on storage clients that enables management of physical disks as logical devices. It enhances data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems.
VVR (Veritas Volume Replicator)	A data replication tool designed to contribute to an effective disaster recovery plan.