

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

Linux

5.1 Service Pack 1



Veritas Storage Foundation™ Cluster File System for Oracle RAC Installation and Configuration Guide

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

Product version: 5.1 SP1

Document version: 5.1SP1.1

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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.

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

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

docs@symantec.com

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	17
Chapter 1 Introducing Veritas Storage Foundation Cluster File System for Oracle RAC	19
About Veritas Storage Foundation Cluster File System for Oracle RAC	19
Benefits of SFCFS RAC	20
About basic SFCFS RAC components	21
About SFCFS RAC optional components	21
About Symantec Product Authentication Service (AT)	21
Veritas Operations Manager	22
About Cluster Manager (Java Console)	23
About SFCFS RAC features	23
Symantec Operations Readiness Tools	23
About VCS notifications	23
About global clusters	24
About Veritas Volume Replicator	24
How SFCFS RAC works	24
Data flow	25
Communication requirements	26
Application and resource management in SFCFS RAC environments	27
SFCFS RAC cluster setup models	28
Typical configuration of four-node SFCFS RAC cluster	28
Typical configuration of SFCFS RAC clusters in secure mode	30
Typical configuration of VOM-managed SFCFS RAC clusters	31
Typical configuration of SFCFS RAC global clusters for disaster recovery	32

Chapter 2	System requirements	35
	Important preinstallation information	35
	Hardware requirements	35
	Supported Linux operating systems	36
	Required operating system patches	37
	Supported database software	37
	Supported replication technologies for global clusters	38
Chapter 3	Planning to install SFCFS RAC	39
	Planning your network configuration	39
	Planning the public network configuration for Oracle RAC	40
	Planning the private network configuration for Oracle RAC	40
	Planning the storage	41
	Planning the storage for SFCFS RAC	42
	Planning the storage for Oracle RAC	42
	Planning volume layout	44
	Planning file system design	45
	Planning for cluster management	45
	Planning for disaster recovery	46
	Planning a global cluster setup	46
	Data replication considerations	46
Chapter 4	Licensing SFCFS RAC	47
	About Veritas product licensing	47
	About SFCFS RAC licenses	48
	Setting or changing the product level for keyless licensing	49
	Installing Veritas product license keys	51
Section 2	Installation and configuration of SFCFS RAC	53
Chapter 5	Preparing to install and configure SFCFS RAC	55
	About preparing to install and configure SFCFS RAC	55
	Preparing to configure the clusters in secure mode	56
	Installing the root broker for the security infrastructure	60
	Creating authentication broker accounts on root broker system	61
	Creating encrypted files for the security infrastructure	62
	Preparing the installation system for the security infrastructure	64

Setting the umask before installation	65
Synchronizing time settings on cluster nodes	65
Setting up inter-system communication	66
Mounting the product disc	66
Setting up shared storage	67
Setting the environment variables	67
Configuring the I/O scheduler	68
Optimizing LLT media speed settings on private NICs	68
Guidelines for setting the media speed of the LLT interconnects	69
Verifying the systems before installation	69
Chapter 6	
Installing and configuring SFCFS RAC	71
About installing and configuring SFCFS RAC	71
About installation and configuration methods	72
Installing SFCFS RAC using the Veritas script-based installation program	73
Installing SFCFS RAC using Kickstart	76
Sample Kickstart configuration file	79
Configuring the SFCFS RAC components using the script-based installer	81
Configuring the SFCFS RAC cluster	82
Creation of SFCFS RAC configuration files	95
Stopping and starting SFCFS RAC processes	95
Installing SFCFS RAC using the Veritas Web-based installation program	96
Before using the Veritas Web-based installer	96
Starting the Veritas Web-based installer	97
Installing products with the Veritas Web-based installer	98
Configuring SFCFS RAC using the Web-based installer	99
Chapter 7	
Installing and configuring SFCFS RAC using a response file	105
About response files	105
About the installation simulator	107
Response file syntax	107
Installing and configuring SFCFS RAC	108
Sample response file for installing and configuring SFCFS RAC	109
Response file variables to install or uninstall SFCFS RAC	110
Response file variables to configure SFCFS RAC	113

Chapter 8	Performing post-installation and configuration tasks	123
	About enabling LDAP authentication for clusters that run in secure mode	123
	Enabling LDAP authentication for clusters that run in secure mode	125
	Verifying LLT, GAB, and cluster operation	131
	Verifying LLT	131
	Verifying GAB	134
	Verifying the cluster	135
	Verifying the cluster nodes	136
Section 3	Upgrade of SFCFS RAC	139
Chapter 9	Preparing to upgrade	141
	About upgrading SFCFS RAC	141
	About types of upgrade	141
	Supported upgrade paths	142
	Preparing to upgrade to SFCFS RAC 5.1 SP1	144
Chapter 10	Performing a full upgrade to SFCFS RAC 5.1 SP1	147
	About full upgrades	147
	Upgrading SFCFS RAC and the operating system (major OS upgrade)	148
	Upgrading SFCFS RAC and the operating system (minor OS upgrade)	149
	Upgrading SFCFS RAC using the Veritas script-based installation program	150
	Upgrading SFCFS RAC with the Veritas Web-based installer	151
	Upgrading SFCFS RAC using a response file	152
	Response file variables to upgrade SFCFS RAC	153
	Sample response file for upgrading SFCFS RAC	154
Chapter 11	Performing a phased upgrade to SFCFS RAC 5.1 SP1	157
	About phased upgrade	157
	Performing phased upgrade of SFCFS RAC from version 5.0	158

Chapter 12	Performing a rolling upgrade to SFCFS RAC 5.1 SP1	165
	About rolling upgrades	165
	Prerequisites for a rolling upgrade	165
	Preparing to perform a rolling upgrade to SFCFS RAC 5.1 SP1	166
	Performing a rolling upgrade using the installer	167
	Performing a rolling upgrade on kernel RPMs: phase 1	167
	Performing a rolling upgrade on non-kernel RPMs: phase 2	168
	Performing a rolling upgrade of SFCFS RAC using the Web-based installer	169
Chapter 13	Performing post-upgrade tasks	171
	Setting or changing the product license level	171
	Upgrading disk layout versions	171
	Upgrading VxVM disk group version	172
Section 4	Installation and upgrade of Oracle RAC	173
Chapter 14	Before installing Oracle RAC	175
	About preparing to install Oracle RAC	175
	Preparing to install Oracle RAC manually	175
	Identifying the public virtual IP addresses for use by Oracle	176
	Setting the kernel parameters	176
	Creating Oracle user and groups	177
	Creating storage for OCR and voting disk manually	177
	Creating Oracle Clusterware and Oracle database home directories manually	183
	Setting up user equivalence	189
	Editing the Oracle user profile	189
	Adding private IP addresses to the /etc/hosts file manually	190
Chapter 15	Installing Oracle RAC	191
	Installing Oracle Clusterware using the Oracle Universal Installer	191
	Installing the Oracle RAC database using the Oracle Universal Installer	193
	Verifying the Oracle Clusterware and database installation	195
	Completing the post-installation tasks	195

	Relinking with ODM	195
	Creating Oracle databases	196
	Increasing the peer inactivity timeout of LLT	201
	Setting the start and stop init sequence for VCS and Oracle Clusterware	201
	Configuring LLT to use bonded network interfaces (optional)	202
Chapter 16	Upgrading Oracle RAC and migrating the database	203
	Supported upgrade paths	203
	Upgrading the Oracle database	204
Section 5	Adding or removing nodes from an SFCFS RAC cluster	205
Chapter 17	Adding a node to SFCFS RAC clusters using Oracle RAC	207
	About adding a node to a cluster	207
	Before adding a node to a cluster	207
	Meeting hardware and software requirements	208
	Setting up the hardware	208
	Preparing to add a node to a cluster	209
	Adding a node to a cluster	210
	Adding a node to a cluster using the SFCFS RAC installer	210
	Adding a node using the Web-based installer	213
	Adding the node to a cluster manually	214
Chapter 18	Removing a node from SFCFS RAC clusters using Oracle RAC	223
	Removing a node from a cluster	223
Section 6	Configuration of disaster recovery environments	225
Chapter 19	Setting up a replicated global cluster	227
	Replication in the SFCFS RAC environment	227
	Requirements for SFCFS RAC global clusters	228

Supported software and hardware for SFCFS RAC	228
Supported replication technologies for SFCFS RAC	228
Enabling a keyless license for SFCFS RAC with VVR for a global cluster	230
About setting up a global cluster in an SFCFS RAC environment	231
Configuring a cluster at the primary site	231
Configuring a cluster at the secondary site	233
Setting up the cluster on the secondary site	233
Setting up the database for the secondary site	234
Configuring replication on clusters at both sites	236
Modifying the ClusterService group for global clusters	236
Modifying the global clustering configuration using the wizard	237
Defining the remote cluster and heartbeat objects	238
Configuring the VCS service groups for global clusters	241
Chapter 20 Configuring a global cluster using VVR	243
About configuring global clustering using VVR	243
Setting up replication using VVR on the primary site	244
Creating the SRL volume on the primary site	244
Setting up the Replicated Volume Group (RVG) on the primary site	245
Setting up replication using VVR on the secondary site	247
Creating the data and SRL volumes on the secondary site	247
Editing the /etc/vx/vras/.rdg files	248
Setting up IP addresses for RLINKs on each cluster	248
Setting up the disk group on secondary site for replication	249
Starting replication of Oracle RAC database volume	251
Starting replication using automatic synchronization	251
Starting replication using full synchronization with Checkpoint	252
Verifying replication status	253
Configuring VCS to replicate the database volume using VVR	253
About modifying the VCS configuration for replication	253
Modifying the VCS Configuration on the Primary Site	255
Modifying the VCS Configuration on the Secondary Site	259
Using VCS commands on SFCFS RAC global clusters	263
Using VVR commands on SFCFS RAC global clusters	263
About migration and takeover of the primary site role	264
Migrating the role of primary site to the secondary site	264
Migrating the role of new primary site back to the original primary site	265

	Taking over the primary role by the remote cluster	266
	VCS agents to manage wide-area failover	270
Section 7	Uninstallation of SFCFS RAC	273
Chapter 21	Preparing to uninstall SFCFS RAC from a cluster	275
	About uninstalling SFCFS RAC from a cluster	275
	Options for uninstalling SFCFS RAC	275
	Preparing to uninstall SFCFS RAC from a cluster	276
Chapter 22	Uninstalling SFCFS RAC from a cluster	279
	Uninstalling SFCFS RAC from a cluster using the script-based installer	279
	Uninstalling SFCFS RAC with the Veritas Web-based installer	280
	Uninstalling SFCFS RAC using a response file	281
	Response file variables to uninstall SFCFS RAC	281
	Sample response file for uninstalling SFCFS RAC	282
Section 8	Installation reference	285
Appendix A	SFCFS RAC installation RPMs	287
	SFCFS RAC installation RPMs	287
Appendix B	Installation scripts	291
	About installation scripts	291
	Starting and stopping processes for the Veritas products	292
	Restarting the installer after a failed connection	292
	Installation program has improved failure handling	292
	Installation script options	293
Appendix C	Configuration files	299
	About the LLT and GAB configuration files	299
	Sample main.cf file for configuring a volume and file system under VCS	301

Appendix D	Automatic Storage Management	307
	About ASM in SFCFS RAC environments	307
	ASM configuration with SFCFS RAC	308
	Configuring ASM in SFCFS RAC environments	309
	Setting up Automatic Storage Management	310
	Creating database storage on ASM	311
	Creating ASM disk groups and instances	312
	Verify the ASM installation	314
	Configuring VCS service groups for database instances on ASM	314
Appendix E	High availability agent information	317
	About agents	317
	VCS agents included within SFCFS RAC	318
	CVMCluster agent	318
	Entry points for CVMCluster agent	318
	Attribute definition for CVMCluster agent	319
	CVMCluster agent type definition	319
	CVMCluster agent sample configuration	320
	CVMVxconfigd agent	320
	Entry points for CVMVxconfigd agent	321
	Attribute definition for CVMVxconfigd agent	321
	CVMVxconfigd agent type definition	325
	CVMVxconfigd agent sample configuration	325
	CVMVolDg agent	325
	Entry points for CVMVolDg agent	326
	Attribute definition for CVMVolDg agent	327
	CVMVolDg agent type definition	329
	CVMVolDg agent sample configuration	329
	CFSMount agent	329
	Entry points for CFSMount agent	330
	Attribute definition for CFSMount agent	330
	CFSMount agent type definition	333
	CFSMount agent sample configuration	334

1

Section

Installation overview and planning

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

Introducing Veritas Storage Foundation Cluster File System for Oracle RAC

This chapter includes the following topics:

- [About Veritas Storage Foundation Cluster File System for Oracle RAC](#)
- [About basic SFCFS RAC components](#)
- [About SFCFS RAC optional components](#)
- [About SFCFS RAC features](#)
- [How SFCFS RAC works](#)
- [SFCFS RAC cluster setup models](#)

About Veritas Storage Foundation Cluster File System for Oracle RAC

Veritas Storage Foundation™ Cluster File System for Oracle® RAC (SFCFS RAC) leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of Oracle RAC on the Linux platform. 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 Disk Manager (VRTSodm), Veritas Cluster

File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Benefits of SFCFS RAC

SFCFS RAC provides the following benefits:

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

About basic SFCFS RAC components

The basic components of SFCFS RAC are as follows:

Veritas Cluster Server	Veritas Cluster Server (VCS) manages Oracle RAC databases and infrastructure components in a clustered environment.
Cluster Volume Manager	Cluster Volume Manager (CVM) enables simultaneous access to the shared volumes that are based on technology from Veritas Volume Manager (VxVM).
Cluster File System	Cluster File System (CFS) enables simultaneous access to the shared file systems that are based on technology from Veritas File System (VxFs).
Oracle Disk Manager	Oracle Disk Manager (ODM) is a disk and file management interface that is provided by Oracle to improve disk I/O performance. ODM enables Oracle to allocate and release disk space, manage tablespaces, and read/write disk blocks directly. SFCFS RAC uses a custom driver that enables applications to use ODM for enhanced file system performance and easy file administration.

About SFCFS RAC optional components

You can add the following optional components to SFCFS RAC:

Symantec Product Authentication Service	See “ About Symantec Product Authentication Service (AT) ” on page 21.
Veritas Operations Manager	See “ Veritas Operations Manager ” on page 22.
Cluster Manager (Java console)	See “ About Cluster Manager (Java Console) ” on page 23.

About Symantec Product Authentication Service (AT)

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

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

- Root broker

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

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

- Authentication brokers

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

See Symantec Product Authentication Service documentation for more information.

See “[Preparing to configure the clusters in secure mode](#)” on page 56.

Veritas Operations Manager

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

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

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

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

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

About Cluster Manager (Java Console)

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

You can download the console from http://go.symantec.com/vcsm_download.

About SFCFS RAC features

You can configure the following features in an SFCFS RAC cluster:

- Symantec Operations Readiness Tools
See “[Symantec Operations Readiness Tools](#)” on page 23.
- VCS notifications
See “[About VCS notifications](#)” on page 23.
- Global clusters
See “[About global clusters](#)” on page 24.
- Veritas Volume Replicator
See “[About Veritas Volume Replicator](#)” on page 24.

Symantec Operations Readiness Tools

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

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

To access SORT, go to:

<http://sort.symantec.com>

About VCS notifications

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

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

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

About 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.

How SFCFS RAC works

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

In traditional environments, only one instance accesses a database at a specific time. SFCFS RAC enables all nodes to concurrently run Oracle instances and execute transactions against the same database. This software coordinates access to the shared data for each node to provide consistency and integrity.

Each node adds its processing power to the cluster as a whole and can increase overall throughput or performance.

At a conceptual level, SFCFS RAC is a cluster that manages applications (instances), networking, and storage components using resources that are contained in service groups.

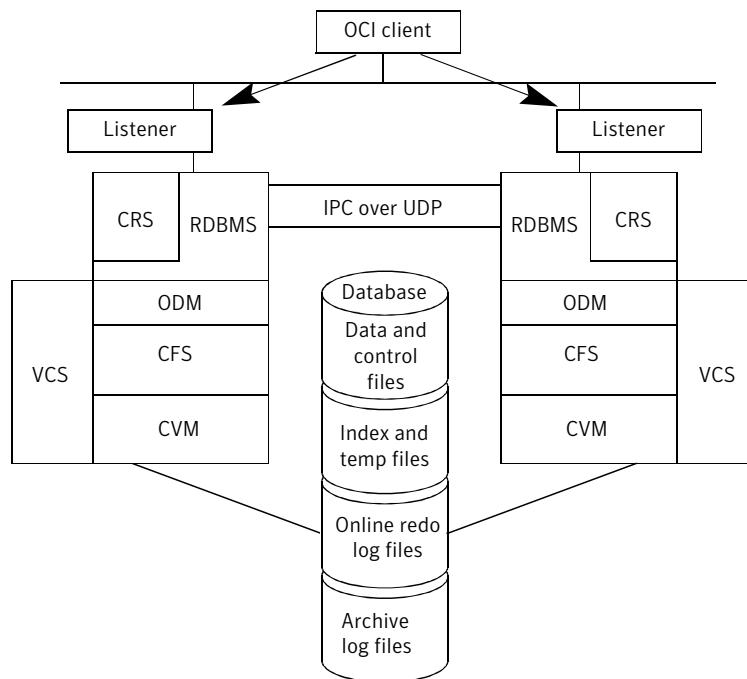
SFCFS RAC clusters have many of the same properties as VCS clusters, such as the following:

- Each node runs its own operating system.
- A cluster interconnect enables cluster communications.
- A public network connects each node to a LAN for client access.
- Shared storage is accessible by each node that needs to run the application.

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

[Figure 1-1](#) describes the SFCFS RAC architecture.

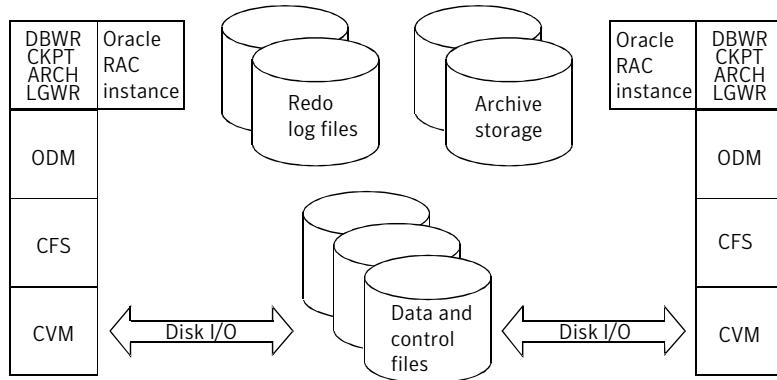
Figure 1-1 SFCFS RAC architecture



Data flow

The CVM, CFS, ODM, and Oracle RAC elements reflect the overall data flow, or data stack, from an instance running on a server to the shared storage. The various Oracle processes that compose an instance (such as DB Writers, Checkpoint, Archiver, Log Writer, and Server) read and write data to the storage through the I/O stack. Oracle communicates through the ODM interface to CFS, which in turn accesses the storage through CVM.

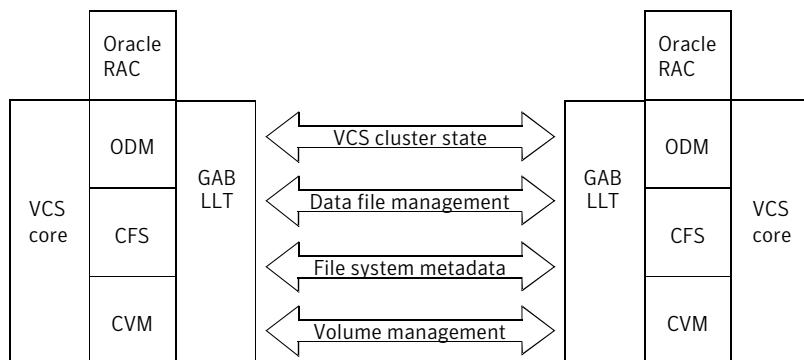
[Figure 1-2](#) describes the data flow in SFCFS RAC.

Figure 1-2 Data flow in SFCFS RAC

Communication requirements

Users on a client system are unaware that they are accessing a database hosted by multiple instances. The key to performing I/O to a database accessed by multiple instances is communication between the processes. Each layer or component in the data stack must reliably communicate with its peer on other nodes to function properly. RAC instances communicate to coordinate the protection of data blocks in the database. ODM processes communicate to coordinate the protection of data files and access across the cluster. CFS coordinates metadata updates for file systems, and CVM coordinates the status of logical volumes and maps.

[Figure 1-3](#) describes the communication requirements in SFCFS RAC.

Figure 1-3 Communication requirements in SFCFS RAC

Note: Oracle cache fusion and lock management traffic goes over UDP since SFCFS RAC does not contain Inter-Process Communication library (VCSIPC) and Low Latency Transport Multiplexer (LMX) module provided by Veritas.

Application and resource management in SFCFS RAC environments

In SFCFS RAC environments, the following resources can be configured under VCS:

- Disk groups and mount points for Oracle Clusterware and Oracle database
- Oracle database binaries
- OCR and voting disk
- Oracle database data files

Oracle Clusterware manages the database and other Oracle-specific resources.

Both VCS and Oracle Clusterware may independently attempt to bring resources online leading to application startup failures. For example, if Oracle Clusterware attempts to start the Oracle database without the availability of underlying Oracle database volumes and mount points managed by VCS, the startup process fails. Application startups in such scenarios may require manual intervention. This section discusses a few application startup scenarios that require manual intervention.

Scenario 1: Oracle Clusterware starts before OCR and voting disk resources are brought online by VCS

In this scenario, the OCR and voting disk is configured under VCS. Oracle Clusterware starts before VCS brings the OCR and voting disk resources online.

Oracle Clusterware will continue to wait until VCS brings the OCR and voting disk resources online. When the resources come online, Oracle Clusterware will start. This is true even for private IP addresses that are required by Oracle Clusterware.

Scenario 2: Oracle Clusterware starts Oracle database before the database volumes and mount points are brought online by VCS

In this scenario, the Oracle database volumes and mount points are configured under VCS. Oracle Clusterware manages the Oracle database. Oracle Clusterware attempts to start the Oracle database before VCS brings the database volumes and mount points online. Oracle database fails to start as the required resources are not available.

Oracle Clusterware does not retry starting the database after the volumes and mount points come online. You must manually start the database after the underlying resources come online.

SFCFS RAC cluster setup models

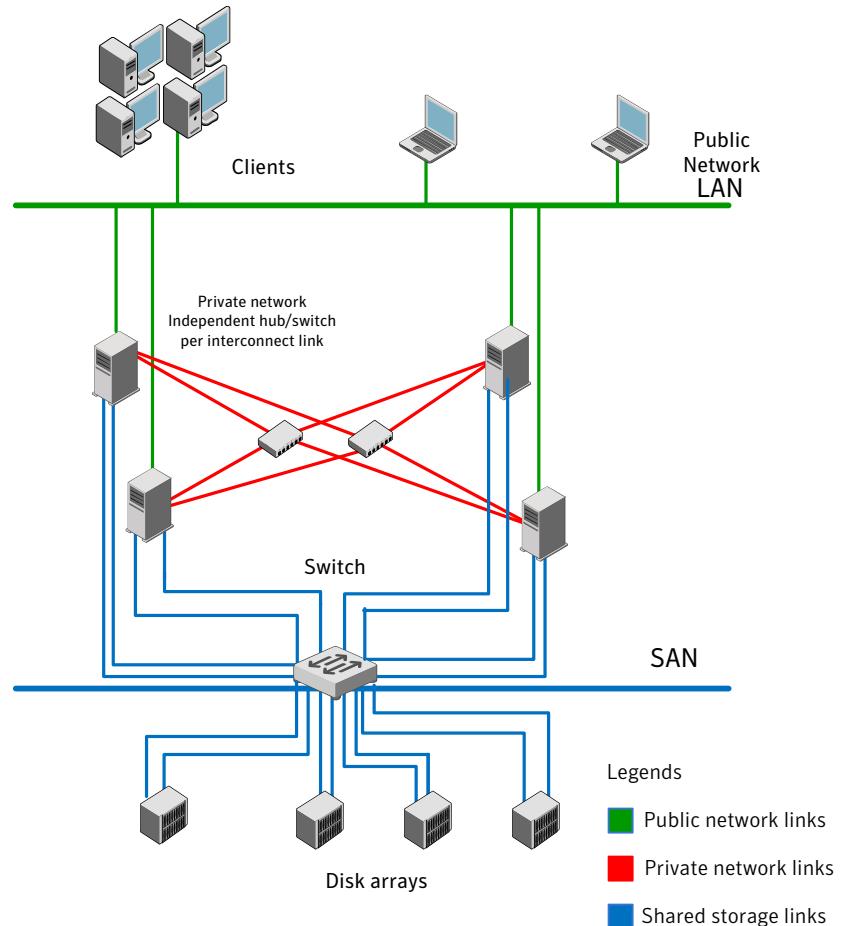
SFCFS RAC supports a variety of cluster configurations.

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

- Basic setup
See “[Typical configuration of four-node SFCFS RAC cluster](#)” on page 28.
- Secure setup
See “[Typical configuration of SFCFS RAC clusters in secure mode](#)” on page 30.
- Central management setup
See “[Typical configuration of VOM-managed SFCFS RAC clusters](#)” on page 31.
- Global cluster setup
See “[Typical configuration of SFCFS RAC global clusters for disaster recovery](#)” on page 32.

Typical configuration of four-node SFCFS RAC cluster

[Figure 1-4](#) depicts a high-level view of a basic SFCFS RAC configuration for a four-node cluster.

Figure 1-4 Sample four-node SFCFS 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.

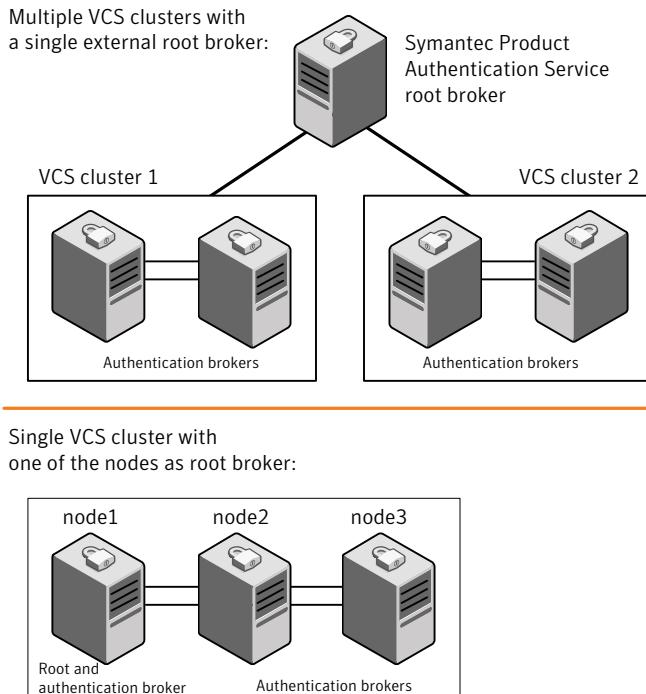
- Nodes that are connected with private network links using similar network devices and matching port numbers.
For example, if you use eth1 on one end of a link, it is recommended that the other end also use eth1.
- Every system must have a local disk.
- The Oracle Cluster Registry and vote disks configured 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.
- VCS manages the resources that are required by Oracle RAC. The resources must run in parallel on each node.

Typical configuration of SFCFS RAC clusters in secure mode

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

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 21.

Figure 1-5 illustrates typical configuration of VCS clusters in secure mode. You can use one of the cluster nodes as AT root broker or you can use a stable system outside the cluster as AT root broker.

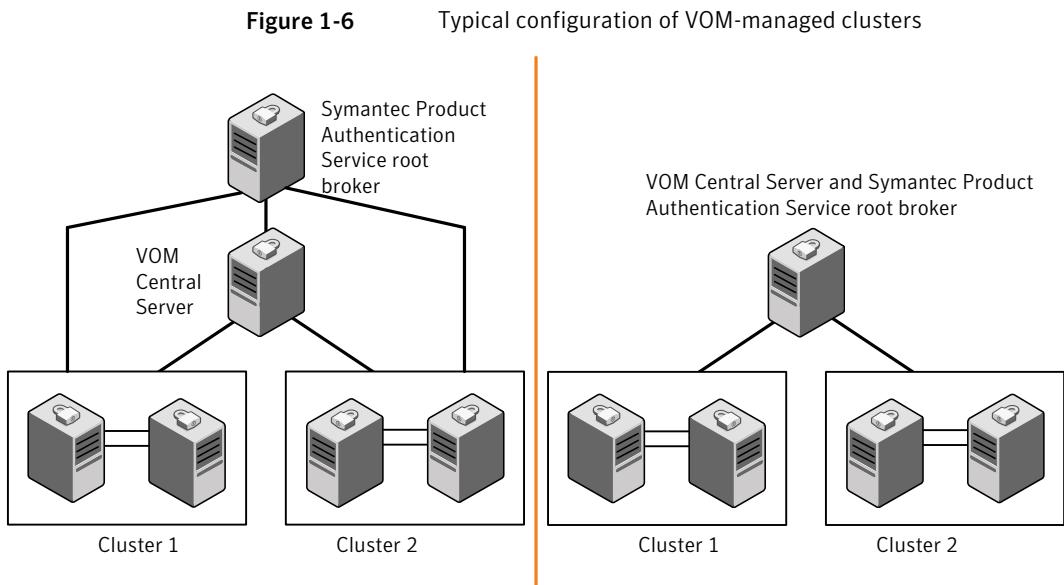
Figure 1-5 Typical configuration of VCS clusters in secure mode

Typical configuration of VOM-managed SFCFS RAC clusters

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

See “[Veritas Operations Manager](#)” on page 22.

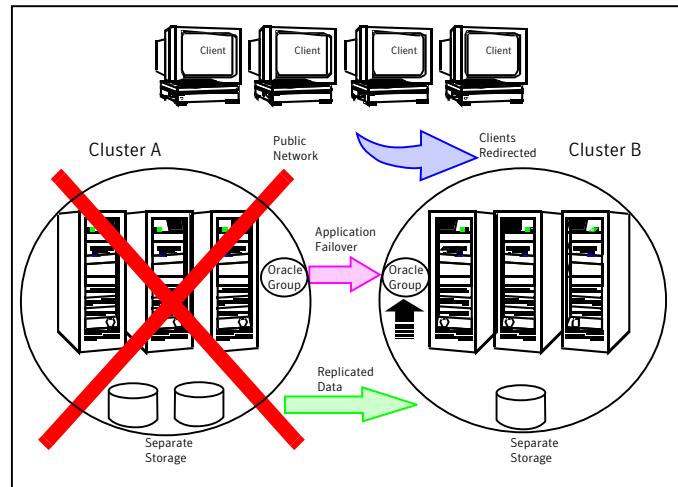
Figure 1-6 illustrates a typical setup of SFCFS RAC clusters that are centrally managed using Veritas Operations Manager. You can install Symantec Product Authentication Service root broker on the same system as that of VOM Central Server or on a different system.



Typical configuration of SFCFS RAC global clusters for disaster recovery

SFCFS 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 software-based replication technologies.

Figure 1-7 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 a VCS_GCO 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 licenses at a later time or you may add the licenses during the SFCFS RAC installation.

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

System requirements

This chapter includes the following topics:

- [Important preinstallation information](#)
- [Hardware requirements](#)
- [Supported Linux operating systems](#)
- [Required operating system patches](#)
- [Supported database software](#)
- [Supported replication technologies for global clusters](#)

Important preinstallation information

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

- Hardware compatibility list for information about supported hardware:
<http://entsupport.symantec.com/docs/330441>
- Disk storage array support information:
<http://entsupport.symantec.com/docs/283282>
- Latest information on support for Oracle database versions:
<http://www.symantec.com/docs/TECH44807>
- Oracle documentation for additional requirements pertaining to your version of Oracle.

Hardware requirements

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

Table 2-1 Hardware requirements for basic clusters

Item	Description
SFCFS RAC systems	Two to sixteen systems with two or more CPUs at 2GHz or higher.
DVD drive	A DVD drive on one of the nodes in the cluster.
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: <pre># ./installsfcrac -precheck node_name</pre> For details on the additional space that is required for Oracle, see the Oracle documentation.
RAM	Each SFCFS RAC system requires at least 2 GB. Symantec recommends additional amount of at least twice the Oracle SGA size.
Swap space	See the Oracle Metalink document: 169706.1
Network links	Two or more private links and one public link. Links must be 100BaseT or gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication. These links must be of the same type; you cannot mix 100BaseT and gigabit. Symantec recommends gigabit Ethernet using enterprise-class switches for the private links.
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

Supported Linux operating systems

This section lists the supported operating systems for this release of Veritas products.

The Veritas 5.1 Service Pack 1 release supports the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel) or SP3 on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5-default kernel) or SUSE Linux Enterprise Server 11 (SLES 11) with SP1 on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux 5 (OEL 5) with Update 3 or later (Red Hat compatible kernel mode only)

Note: 64-bit operating systems are only supported.

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Enterprise Linux, you must upgrade it before attempting to install the Veritas software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your system.

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat errata and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

<http://entsupport.symantec.com/docs/335001>

Required operating system patches

Oracle RAC requires the following patch on systems running SLES10 SP2:
kernel 2.6.16.60-0.39.3

Supported database software

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

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

Note: SFCFS RAC supports only 64-bit Oracle.

The following database versions are supported:

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

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

Supported replication technologies for global clusters

SFCFS RAC supports the software replication technology Veritas Volume Replicator (VVR) for global cluster configurations.

Planning to install SFCFS RAC

This chapter includes the following topics:

- [Planning your network configuration](#)
- [Planning the storage](#)
- [Planning volume layout](#)
- [Planning file system design](#)
- [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.

Planning the public network configuration for Oracle RAC

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

Planning the private network configuration for Oracle RAC

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

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:

- Oracle Clusterware interconnects need to be protected against NIC failures and link failures. Configure the Oracle Clusterware interconnects over bonded NIC interfaces. LLT recognizes aggregated links and treats the aggregated link as a single network interface. Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link.
- Configure Oracle Cache Fusion traffic to take place through the private network. Symantec also recommends that all UDP cache-fusion links be LLT links. 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	<p>Use a native NIC bonding solution to provide redundancy, in case of NIC failures.</p> <p>Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link.</p> <p>When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership:</p> <ul style="list-style-type: none"> ■ 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>

Planning the storage

SFCFS 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

See “[Planning for Oracle ASM over CVM](#)” on page 43.

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.

- 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 SFCFS RAC

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

Table 3-2 Type of storage required for SFCFS RAC

SFCFS RAC files	Type of storage
SFCFS RAC binaries	Local

Planning the storage for Oracle RAC

Oracle Cluster Registry (OCR) and voting disk are not supported on Cluster File System (CFS). Symantec recommends keeping OCR and voting disk on volumes created by the Cluster Volume Manager (CVM).

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

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 and Oracle recommend 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.
- 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.
- To support the execution of multiple versions of Oracle at the same time, store each version of the Oracle software in a directory path as follows:

`/mount_point/StdDirectoryName/OwnerOfDirectory/product/version/type`

For example:`/u01/app/oracle/product/10.2.0/dbhome`

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	<p>Local</p> <p>The ORACLE_BASE directory specified during the Oracle Clusterware installation is used to store the Oracle Inventory files. Oracle requires that the Oracle inventory files be placed on the local file system.</p>
Oracle Clusterware binaries	Local
Oracle database binaries	<p>Local or shared</p> <p>Placing the Oracle database binaries on local disks enables rolling upgrade of the cluster.</p>
Database datafiles	<p>Shared</p> <p>Store the Oracle database files on CFS rather than on raw device or CVM raw device for easier management. Create separate clustered file systems for each Oracle database. Keeping the Oracle database datafiles on separate mount points enables you to unmount the database for maintenance purposes without affecting other databases.</p> <p>Note: Set the CVM disk group fail policy to <code>leave</code> for disk groups containing the data files.</p> <p>If you plan to store the Oracle database on ASM, configure the ASM disk groups over CVM volumes to take advantage of dynamic multi-pathing.</p>
Database recovery data (archive, flash recovery)	<p>Shared</p> <p>Place archived logs on CFS rather than on local file systems.</p>

Planning for Oracle ASM over CVM

ASM provides storage for data files, control files, online and archive log files, and backup files.

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

The following practices offer high availability and better performance:

- Use CVM mirrored volumes with dynamic multi-pathing for creating ASM disk groups. Select external redundancy while creating ASM disk groups.
- The CVM raw volumes used for ASM must be used exclusively for ASM. Do not use these volumes for any other purpose, such as creation of file systems. Creating file systems on CVM raw volumes used with ASM may cause data corruption.
- Do not enable ODM when databases are installed on ASM. ODM is a disk management interface for data files that reside on the Veritas File System.
- Use a minimum of two Oracle ASM disk groups. Store the data files, one set of redo logs, and one set of control files on one disk group. Store the Flash Recovery Area, archive logs, and a second set of redo logs and control files on the second disk group.
For more information, see Oracle's ASM best practices document.
- Do not configure DMP meta nodes as ASM disks for creating ASM disk groups. Access to DMP meta nodes must be configured to take place through CVM.
- Do not combine DMP with other multipathing software in the cluster.
- Do not use coordinator disks, which are configured for I/O fencing, as ASM disks.
- Volumes presented to a particular ASM diskgroup should be of the same geometry, speed, and type.

Planning volume layout

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

- Mirror the volumes across two or more storage arrays, if using VxVM mirrors.
- Separate the Oracle recovery structures from the database files to ensure high availability when you design placement policies.
- Separate redo logs and place them on the fastest storage (for example, RAID 1+ 0) for better performance.
- Use "third-mirror break-off" snapshots for cloning the Oracle log volumes. Do not create Oracle log volumes on a Space-Optimized (SO) snapshot.
- Create as many Cache Objects (CO) as possible when you use Space-Optimized (SO) snapshots for Oracle data volumes.
- Distribute the I/O load uniformly on all Cache Objects when you create multiple Cache Objects.

- If using VxVM mirroring, keep the Fast Mirror Resync 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.
- 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.

Planning for cluster management

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

Table 3-4 List of agents

Agent	Description
VCS agent for CVM	Volume management An SFCFS RAC installation automatically configures the CVMCluster resource and the CVMVxconfig resource. You must configure the CVMVolDg agent for each shared disk group.
VCS agent for CFS	File system management If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.

Planning for disaster recovery

SFCFS RAC supports global cluster configurations in multi-site clusters for disaster recovery. 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. SFCFS RAC supports the software replication technology Veritas Volume Replicator for data replication.

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

Planning a global cluster setup

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

Global clustering involves two steps:

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

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

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

Data replication considerations

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

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

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

Licensing SFCFS RAC

This chapter includes the following topics:

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

About Veritas product licensing

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

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

www.symantec.com/techsupp/

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

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

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

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

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
See “[Setting or changing the product level for keyless licensing](#)” on page 49.
See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
See “[Installing Veritas product license keys](#)” on page 51.
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 SFCFS RAC licenses

[Table 4-1](#) lists the various SFCFS RAC license levels and the corresponding features.

Table 4-1 SFCFS RAC license levels

License	Description	Features enabled
SFCFSRACENT	SFCFS RAC Enterprise Edition	The license enables the following features: <ul style="list-style-type: none"> ■ 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
SFCFSRACENT_VVR	SFCFS RAC Enterprise Edition with VVR	The license enables the following features: <ul style="list-style-type: none"> ■ 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 <p>Note: If you plan to set up global clusters in an SFCFS RAC environment, you also need to have the VCS_GCO license.</p>
VCS_GCO	Cluster Server with GCO option	This license enables the global cluster option for SFCFS RAC. <p>Note: Keyless licensing is not supported for the global cluster option in SFCFS RAC.</p>

Setting or changing the product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 View the current setting for the product level.

```
# vxkeyless -v display
```

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

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

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

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas product license keys

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

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

Even though other products are included on the enclosed software discs, you can only 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
```


2

Section

Installation and configuration of SFCFS RAC

- [Chapter 5. Preparing to install and configure SFCFS RAC](#)
- [Chapter 6. Installing and configuring SFCFS RAC](#)
- [Chapter 7. Installing and configuring SFCFS RAC using a response file](#)
- [Chapter 8. Performing post-installation and configuration tasks](#)

Preparing to install and configure SFCFS RAC

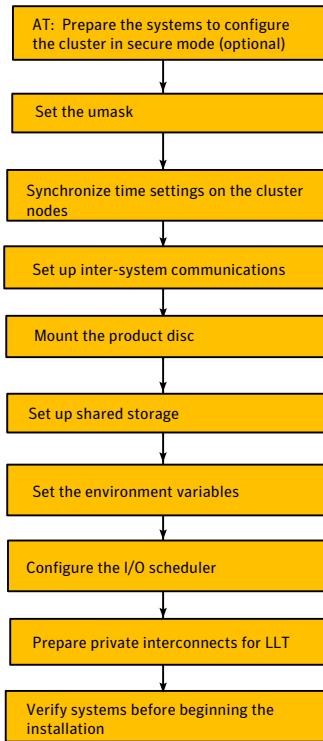
This chapter includes the following topics:

- [About preparing to install and configure SFCFS RAC](#)
- [Preparing to configure the clusters in secure mode](#)
- [Setting the umask before installation](#)
- [Synchronizing time settings on cluster nodes](#)
- [Setting up inter-system communication](#)
- [Mounting the product disc](#)
- [Setting up shared storage](#)
- [Setting the environment variables](#)
- [Configuring the I/O scheduler](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Verifying the systems before installation](#)

About preparing to install and configure SFCFS RAC

[Figure 5-1](#) illustrates an overview of the pre-installation steps for SFCFS RAC.

Figure 5-1 SFCFS RAC pre-installation tasks



Preparing to configure the clusters in secure mode

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

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

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

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

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

- To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system.

See “[Installing the root broker for the security infrastructure](#)” on page 60.

- To use one of the cluster nodes as root broker, the installer does not require you to do any preparatory tasks.

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

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

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

- For external root broker, an authentication broker (AB) account for each node in the cluster is set up on the root broker system.

See “[Creating authentication broker accounts on root broker system](#)” on page 61.

- The system clocks of the external root broker and authentication brokers must be in sync.

The script-based installer provides the following configuration modes:

Automatic mode The external root broker system must allow rsh or ssh passwordless login to use this mode.

Semi-automatic mode This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode.

The nodes in the cluster must allow rsh or ssh passwordless login.

Manual mode This mode requires root_hash file and the root broker information from the AT administrator to configure a cluster in secure mode.

The nodes in the cluster must allow rsh or ssh passwordless login.

[Figure 5-2](#) depicts the flow of configuring SFCFS RAC cluster in secure mode.

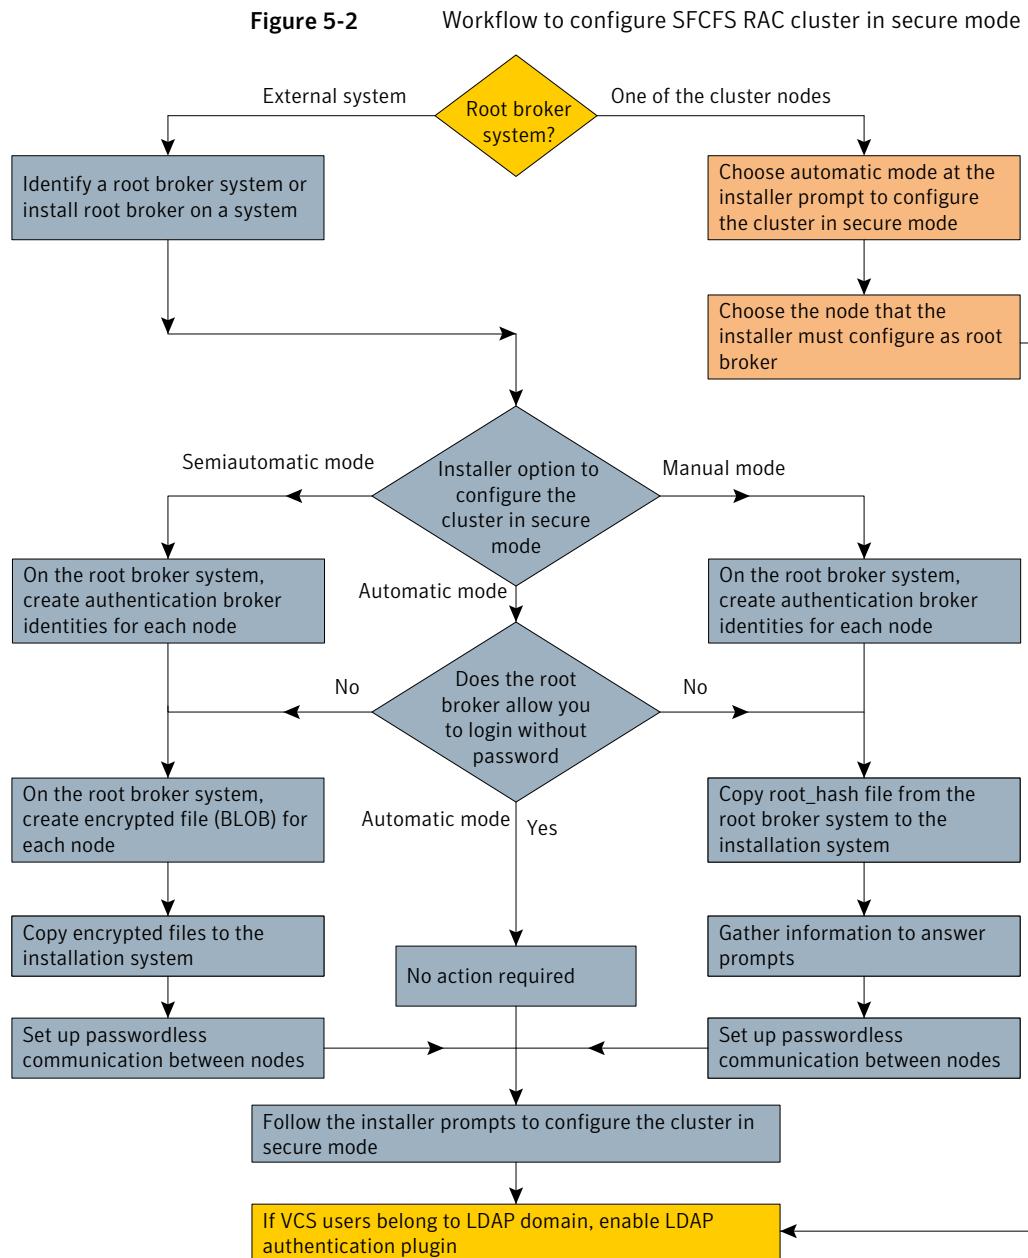


Table 5-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform. These preparatory tasks apply only when you use an external root broker system for the cluster.

Table 5-1 Preparatory tasks to configure a cluster in secure mode (with an external root broker)

Tasks	Who performs this task
<p>Decide one of the following configuration modes to set up a cluster in secure mode:</p> <ul style="list-style-type: none"> ■ Automatic mode ■ Semi-automatic mode ■ Manual mode 	VCS administrator
<p>Install the root broker on a stable system in the enterprise.</p> <p>See “Installing the root broker for the security infrastructure” on page 60.</p>	AT administrator
<p>To use the semi-automatic mode or the manual mode, on the root broker system, create authentication broker accounts for each node in the cluster.</p> <p>See “Creating authentication broker accounts on root broker system” on page 61.</p> <p>The AT administrator requires the following information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker 	AT administrator
<p>To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator.</p> <p>See “Creating encrypted files for the security infrastructure” on page 62.</p> <p>The AT administrator requires the following additional information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Administrator password for each authentication broker <p>Typically, the password is the same for all nodes.</p>	AT administrator
<p>To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.</p>	AT administrator

Table 5-1 Preparatory tasks to configure a cluster in secure mode (with an external root broker) (*continued*)

Tasks	Who performs this task
<p>Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure SFCFS RAC.</p> <p>See “Preparing the installation system for the security infrastructure” on page 64.</p>	VCS administrator

Installing the root broker for the security infrastructure

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

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

See Symantec Product Authentication Service documentation for more information.

See “[About Symantec Product Authentication Service \(AT\)](#)” on page 21.

To install the root broker

- 1 Mount the product disc and start the installer.

```
# ./installer
```

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

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

- 7 Review the output as the installer does the following:
 - Checks to make sure that AT supports the operating system
 - Checks if the RPMs are already on the system.The installer lists the RPMs that the program is about to install on the system.
Press Enter to continue.
- 8 Review the output as the installer installs the root broker on the system.
- 9 After the installation, configure the root broker.
- 10 Select a mode to configure the root broker from the three choices that the installer presents:

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

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

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

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

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

Creating authentication broker accounts on root broker system

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

To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

"Failed To Get Attributes For Principal"

Proceed to step 3.

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

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

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

Creating encrypted files for the security infrastructure

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

To create encrypted files

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

hash

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

```
venus> # vssat showbrokerhash
```

root_domain

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

```
venus> # vssat showalltrustedcreds
```

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

identity

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

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

See “[Creating authentication broker accounts on root broker system](#)” on page 61.

password

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

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

See “[Creating authentication broker accounts on root broker system](#)” on page 61.

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

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

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

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

Preparing to configure the clusters in secure mode

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

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

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

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

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

For example:

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

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

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

Preparing the installation system for the security infrastructure

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

To prepare the installation system for the security infrastructure

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

Automatic mode Do the following:

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

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

Setting the umask before installation

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

```
# umask 0022
```

Synchronizing time settings on cluster nodes

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, use the `yast` command:

```
# yast ntp-client
```

Setting up inter-system communication

By default, the installer uses SSH for inter-system communication. You must grant permissions to allow the root user to invoke passwordless SSH or RSH communication between cluster nodes and local node.

Check the following resources for information on setting up passwordless inter-system communication:

- See the manual page for `ssh(1M)`, and `rsh(1M)` for information on how to configure them for passwordless communication.
If you encounter issues during configuration, contact the operating system support provider.
- To access online manuals and other resources, visit the OpenSSH website:
<http://openssh.org>

Mounting the product disc

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

You can unmount the product disc after completing the SFCFS 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 SFCFS 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
```

- 4 Navigate to the location of the RPMs:

Depending on the OS distribution and architecture, type the command:

```
RHEL 5(x86_64)  # cd /dvd_mount/rhel5_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
```

```
SLES 10(x86_64) # cd /dvd_mount/sles10_x86_64/\
storage_foundation_cluster_file_system_for_oracle_rac
```

```
SLES 11 (x86_64) # cd /dvd_mount/sles11_x86_64/\\
storage_foundation_cluster_file_system_for_oracle_rac
```

Setting up shared storage

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

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

```
# fdisk -l
```

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

Setting the environment variables

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

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

For sh, ksh, or bash `# export MANPATH=$MANPATH:\\
/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.
- If you use directly connected Ethernet links (using crossover cables), Symantec recommends that you set the media speed to the highest value common to both cards, typically 1000_Full_Duplex.

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

Verifying the systems before installation

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

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

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

The program proceeds in a non-interactive mode, examining the systems for licenses, RPMs, disk space, and system-to-system communications. The

Verifying the systems before installation

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.

Installing and configuring SFCFS RAC

This chapter includes the following topics:

- [About installing and configuring SFCFS RAC](#)
- [About installation and configuration methods](#)
- [Installing SFCFS RAC using the Veritas script-based installation program](#)
- [Installing SFCFS RAC using Kickstart](#)
- [Configuring the SFCFS RAC components using the script-based installer](#)
- [Installing SFCFS RAC using the Veritas Web-based installation program](#)
- [Configuring SFCFS RAC using the Web-based installer](#)

About installing and configuring SFCFS RAC

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

The following packages are installed on each cluster node:

- Veritas Cluster Server (VCS)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Oracle Disk Manager (ODM)

You can configure the following components for SFCFS RAC:

- Veritas Cluster Server (VCS)

Note: You can not configure VCS to manage Oracle Clusterware.

- CVM (Veritas Volume Manager enabled for clusters)
- CFS (Veritas File System enabled for clusters)

About installation and configuration methods

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

Table 6-1 Installation and configuration methods

Method	Description
Interactive installation and configuration using the script-based installer Note: If you obtained SFCFS RAC from an electronic download site, you must use the <code>installsfcfsrac</code> program script instead of the <code>installer</code> script.	You can use one of the following script-based installers: <ul style="list-style-type: none">■ Common product installer script: <code>installer</code> The common product installer script provides a menu that simplifies the selection of installation and configuration options. Use this method to install other products, such as the Symantec Product Authentication Service (AT), along with SFCFS RAC.■ Product-specific installation script: <code>installsfcfsrac</code> program■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsfcfsrac</code> program script is identical to specifying SFCFS RAC from the <code>installer</code> script. Use this method to install or configure only SFCFS RAC.

Table 6-1 Installation and configuration methods (*continued*)

Method	Description
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 105.
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. . /webinstaller See “ Installing SFCFS RAC using the Veritas Web-based installation program ” on page 96.
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 SFCFS RAC using the Veritas script-based installation program

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, 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 SFCFS RAC installation.

- Installs the SFCFS RAC 5.1 SP1 RPMs.

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

To install SFCFS RAC

1 Log in as the superuser on one of the systems.

2 Start the installation program:

SFCFS RAC installer Navigate to the directory that contains the installation program. The program is located in the product directory.

Run the program:

```
# ./installsfccsrac galaxy 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.

Common product installer Navigate to the directory that contains the installation program.

Run the program:

```
# ./installer galaxy 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.

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

From the displayed list of products to install, choose **Storage Foundation Cluster File System for Oracle RAC (SFCFS RAC)**

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

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

The installer has discovered an existing installer process.

The process exited while performing configure of SFCFS RAC on galaxy.

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

- 5 Enter y to agree to the End User License Agreement (EULA).
- 6 Select the type of package installation—Minimal, Recommended, All.
 - 1) Install minimal required Storage Foundation Cluster File System for Oracle RAC rpms
 - 2) Install recommended Storage Foundation Cluster File System for Oracle RAC rpms
 - 3) Install all Storage Foundation Cluster File System for Oracle RAC 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 SFCFS RAC RPMs and patches.

- 7 Select the appropriate license option.
 - 1) Enter a valid license key
 - 2) Enable keyless licensing

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

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

Enter a SFCFS 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 SFCFS RAC without entering a key. However, you must still acquire a valid license to install and use SFCFS RAC. Keyless licensing requires that you manage the systems with a Management Server.

Enter y if you want to configure Veritas Volume Replicator (VVR).

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

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

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

Proceed to configure the SFCFS RAC components.

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

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

Installing SFCFS RAC using Kickstart

You can install SFCFS RAC using Kickstart. Kickstart is supported only for Red Hat Enterprise Linux (RHEL5).

To install SFCFS RAC using Kickstart

- 1 Create a directory for the Kickstart configuration files.

```
# mkdir /kickstart_files/
```

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

- To generate scripts and configuration files for all products, including Veritas Storage Foundation™ (SF), Storage Foundation Cluster File System (SFCFS), Storage Foundation for Oracle® RAC (SFRAC), and Storage Foundation Cluster File System for Oracle® RAC (SFCFSRAC), enter the following command:

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

The system lists the scripts and files as they are generated.

- To only generate scripts and configuration files for SFCFSRAC, enter the following command:

```
# ./installsfcfsrac -kickstart /kickstart_files/
```

The command output includes the following:

```
The kickstart script for SFCFSRAC51 is generated at
/kickstart_files/kickstart_sfccfsrac51.ks
The installer script to configure SFCFS Oracle RAC is generated at
/kickstart_files/installsfccfsrac
The installer script to uninstall SFCFS Oracle RAC is generated at
/kickstart_files/uninstallsfccfsrac
```

- 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 entire directory of RPMs and scripts from the CD/DVD to the NFS location.
- 5 Copy the required SFCFS RAC `installproduct` and `uninstallproduct` scripts you generated in step 2 to the following NFS directory.

```
# cp /kickstart_files/installsfccfsrac /nfs_mount_kickstart
# cp /kickstart_files/uninstallsfccfsrac /nfs_mount_kickstart
```

- 6 Verify the contents of the directory. They should look similar to the following:

```
# ls /nfs_mount_kickstart/
rpms      scripts      installsfcfsrac  uninstallsfcfsrac
```

- 7 In the SFCFS 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"
```

- 8 Append the entire modified contents of the Kickstart configuration file to the operating system `ks.cfg` file.
- 9 Launch the Kickstart installation for the RHEL5 operating system.
- 10 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.

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

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

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

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

For example:

```
# /opt/VRTS/install/installsfcrac -configure galaxy nebula
```

- 12** Verify that all the configured llt 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
```

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

```
# hastatus -summ
-- 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
```

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

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

```
# hastatus -summary
-- Group           System       Probed     AutoDisabled  State
B   VxSS           galaxy        Y          N             ONLINE
B   VxSS           nebula        Y          N             ONLINE
```

Sample Kickstart configuration file

The following is a sample Kickstart configuration file.

```
%post --nochroot
# Add necessary scripts or commands here to your need
# This generated kickstart file is only for the automated installation of products
# in the DVD

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

#
# Notice:
# * Modify the BUILDRC below according to your real environment
# * The location specified with BUILDRC should be NFS accessible
#   to the Kickstart Server
```

```
# * Copy the whole directories of rpms and scripts from
#   installation media to the BUILDSRC
# * Put generated install<product> scripts into the BUILDSRC
#
BUILDSRC="<hostname_or_ip>:/path/to/rpms_and_scripts"

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

# define path variables
ROOT=/mnt/sysimage
BUILDDIR="${ROOT}/build"
RPMDIR="${BUILDDIR}/rpms"
SCRIPTDIR="${BUILDDIR}/scripts"
CPIDIR="${ROOT}/opt/VRTS/install"
CPIMODDIR="${ROOT}/opt/VRTSperl/lib/site_perl/UXRT51"

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

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

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

# install rpms one by one
for RPM in VRTSvllic VRTSperl VRTSspt VRTSvxvm VRTSaslapm VRTSob
VRTSlvconv VRTSsfmh VRTSvxfs VRTSfssdk VRTSatClient VRTSatServer
VRTSllt VRTSgab VRTSvxfen VRTSamf VRTSvcs VRTScps VRTSvcsag
VRTSvcsdr VRTSvcsea VRTSglm VRTScavf VRTSgms VRTSodm

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

# copy CPI perl modules
mkdir -p ${CPIMODDIR}
for MODULE in EDR CPIC CPIP
do
    if [ -d "${SCRIPTDIR}/${MODULE}" ]
```

```

then
    cp -rp ${SCRIPTDIR}/${MODULE} ${CPIMODDIR}
    echo "Perl modules from ${SCRIPTDIR}/${MODULE} are copied to ${CPIMODDIR}" \
        >> ${KSLOG}
else
    echo "Required perl modules ${SCRIPTDIR}/${MODULE} not found" >> ${KSLOG}
fi
done

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

# copy CPI installer scripts
mkdir -p ${CPIDIR}
touch "${CPIDIR}/.cpi5"
for FILE in $(find ${BUILDDIR} -maxdepth 1 -name '*install*')
do
    cp -p ${FILE} ${CPIDIR}
    echo "Copy ${FILE} to ${CPIDIR}" >> ${KSLOG}
done

umount ${BUILDDIR}

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

exit 0

```

Configuring the SFCFS 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.

See “[About preparing to install and configure SFCFS RAC](#)” on page 55.

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 SFCFS RAC, before you start the installer you must stop all the resources that are under VCS control using the hastop command or the hagrps –offline command.

Note: SFCFS RAC is not licensed to support I/O fencing. It is configured in disabled mode during SFCFS RAC configuration.

To configure the SFCFS RAC components

- 1 Configure the Veritas Cluster Server component to set up the SFCFS RAC cluster.
See “[Configuring the SFCFS RAC cluster](#)” on page 82.
- 2 Add VCS users.
See “[Adding VCS users](#)” on page 89.
- 3 Configure SMTP email notification.
See “[Configuring SMTP email notification](#)” on page 90.
- 4 Configure SNMP trap notification.
See “[Configuring SNMP trap notification](#)” on page 92.
- 5 Configure global clusters, if you chose to enable GCO during the installation.

Note: Perform this step only if you provided the global cluster license information during the installation.

See “[Configuring global clusters](#)” on page 94.

- 6 Stop the SFCFS RAC resources.
See “[Stopping and starting SFCFS RAC processes](#)” on page 95.

Configuring the SFCFS RAC cluster

You must configure the Veritas Cluster Server component to set up the SFCFS RAC cluster. You can configure a basic cluster with only the required components, or an advanced cluster with any or all the optional features that meet your configuration requirements.

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

Configuring the cluster name and ID

Enter the cluster information when the installer prompts you.

To configure the cluster

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

```
Enter the unique cluster name: [q,?] rac_cluster101
Enter a unique Cluster ID number between 0-65535: [b,q,?] 7
```

Configuring private heartbeat links

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

The following procedure helps you configure LLT over Ethernet.

To configure private heartbeat links

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

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 can use either the standard interfaces or the aggregated interfaces (bonded NICs).

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)

Do you want to configure an additional low priority heartbeat link? [y,n,q,b,?] (n)

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

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

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

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

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

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

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

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

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

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

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

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 To configure virtual IP, enter **y** at the prompt.
- 3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

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

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (eth0)
```

- 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 eth0 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

If you want to configure the cluster in secure mode, make sure that you meet the prerequisites for secure cluster configuration.

The installsfcfsrac program provides different configuration modes to configure a secure cluster. Make sure that you completed the pre-configuration tasks for the configuration mode that you want to choose.

To configure the cluster in secure mode

1 Choose whether to configure SFCFS RAC to use Symantec Product Authentication Service.

```
Would you like to configure VCS to use Symantec Security  
Services? [y,n,q] (n) y
```

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**. You must add VCS users when the configuration program prompts. See “[Adding VCS users](#)” on page 89.

2 Select one of the options to enable security.

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

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

Option 1. Automatic configuration Enter the name of the Root Broker system when prompted.

Requires a remote access to the Root Broker.

Review the output as the installer verifies communication with the Root Broker system, checks vxatd process and version, and checks security domain.

Option 2. Semiautomatic configuration Enter the path of the encrypted file (BLOB file) for each node when prompted.

Option 3. Manual configuration

Enter the following Root Broker information as the installer prompts you:

```
Enter root Broker name:  
east.symantecexample.com  
Enter root broker FQDN: [b]  
(symantecexample.com)  
symantecexample.com  
Enter root broker domain: [b]  
(root@east.symantecexample.com)  
root@east.symantecexample.com  
Enter root broker port: [b] (2821) 2821  
Enter path to the locally accessible  
root hash [b] (/var/tmp/  
installvcs-1Lcljr/root_hash)  
/root/root_hash
```

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

```
Enter authentication broker principal name on  
galaxy [b]  
(galaxy.symantecexample.com)  
galaxy.symantecexample.com  
Enter authentication broker password on galaxy:  
Enter authentication broker principal name on  
nebula [b]  
(nebula.symantecexample.com)  
nebula.symantecexample.com  
Enter authentication broker password on nebula:
```

- 3 After you provide the required information to configure the cluster in secure mode, the program prompts you to configure SMTP email notification.

Note that the installer does not prompt you to add VCS users if you configured the cluster in secure mode. However, you must add VCS users later.

See *Veritas Cluster Server User's Guide* for more information.

Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

To add VCS users

- 1 Review the required information to add VCS users.
- 2 Reset the password for the Admin user, if necessary.

```
Do you want to set the username and/or password for the Admin user  
(default username = 'admin', password='password')? [y,n,q] (n) y  
Enter the user name: [b,q,?] admin  
Enter the password:  
Enter again:
```

- 3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

- 4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [b,q,?] smith  
Enter New Password:*****
```

```
Enter Again:*****  
Enter the privilege for user smith (A=Administrator, O=Operator,  
G=Guest): [b,q,?] a
```

- 5 Enter **n** at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```

- 6 Review the summary of the newly added users and confirm the information.

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.
- 2 Specify whether you want to configure the SMTP notification.

```
Do you want to configure SMTP notification? [y,n,q,?] (n) y
```

If you do not want to configure the SMTP notification, you can skip to the next configuration option.

See “[Configuring SNMP trap notification](#)” on page 92.

- 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 SFCFS RAC based on the configuration details you provided.

See “[Configuring global clusters](#)” on page 94.

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

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 SFCFS 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 SFCFS RAC configuration files

The program consolidates all the information gathered in the preceding configuration tasks and creates configuration files.

If you chose to configure the cluster in secure mode, the installer also configures the Symantec Product Authentication Service. Depending on the mode you chose to set up Authentication Service, the installer creates security principal or executes the encrypted file to create security principal on each node in the cluster. The installer creates the VxSS service group, creates Authentication Server credentials on each node in the cluster, and Web credentials for VCS users, and sets up trust with the root broker. Then, the installer proceeds to start SFCFS RAC in secure mode.

Review the output as the configuration program creates security principal, starts VCS, creates VCS configuration files, and copies the files to each node.

Stopping and starting SFCFS RAC processes

The installer stops and starts SFCFS RAC processes and configures the SFCFS RAC agents.

Note: Do not opt to start SFCFS 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 SFCFS RAC processes

- 1 Enter **y** to stop SFCFS RAC processes.

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

- 2 Review the output as the installer stops and starts the SFCFS RAC processes.

Installing SFCFS RAC using the Veritas Web-based installation program

The installation of SFCFS 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 96.
- Start the Veritas Web installer.
See “[Starting the Veritas Web-based installer](#)” on page 97.
- Install SFCFS RAC.
See “[Installing products with the Veritas Web-based installer](#)” on page 98.

During the installation, the installer performs the following tasks:

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

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

- Installs the SFCFS RAC 5.1 SP1 RPMs.

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 SFCFS RAC 5.1 SP1.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and must be at one of the supported operating system update levels.
Administrative system	The system where you run the Web browser to perform the installation.	Must have a Web browser. Supported browsers: <ul style="list-style-type: none">■ Internet Explorer 6, 7, and 8■ Firefox 3.x

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.

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

Secure Connection Failed

Obtain a security exception for your browser.

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

Installing products with the Veritas Web-based installer

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

To install SFCFS RAC

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

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

For more information, go to the following website:

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

Complete the following information:

Choose whether you want to enable Global Cluster option.

Click Register.

- Enter license key

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

- 9 The installer prompts you to configure the cluster. Select Yes to continue with configuring the product.
If you select No, you can exit the installer. You must configure the product before you can use SFCFS RAC.
After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.
- 10 Select the checkbox to specify whether you want to send your installation information to Symantec.
- Would you like to send the information about this installation to Symantec to help improve installation in the future?
- Click **Finish**. The installer prompts you for another task.

Configuring SFCFS RAC using the Web-based installer

Before you begin to configure SFCFS RAC using the Web-based installer, review the configuration requirements.

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

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 SFCFS RAC on a cluster

- 1 Start the Web-based installer.
See “[Starting the Veritas Web-based installer](#)” on page 97.
- 2 On the Select a task and a product page, select the task and the product as follows:

Task	Configure a Product
Product	Storage Foundation for Cluster File System for Oracle RAC

Click **Next**.

- 3 On the Select Systems page, enter the system names where you want to configure SFCFS RAC, and click **Validate**.

Example: **galaxy nebula**

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

Click **Next** after the installer completes the system verification successfully.

- 4 On the Set Cluster Name/ID page, specify the following information for the cluster.

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID.
LLT Type	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 6 . 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.
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**.

- 5 On the Set Cluster Heartbeat page, select the heartbeat link details for the LLT type you chose on the Set Cluster Name/ID page.

For **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**.

- 6** In the Confirmation dialog box that appears, choose whether or not to configure the cluster in secure mode using Symantec Product Authentication Service (AT).

To configure the cluster in secure mode, click **Yes**.

If you want to perform this task later, click **No**. You can use the `installsfcfsrac -security` command. Go to step **8**.

- 7** On the Security Options page, choose an option to enable security and specify the required information.

Do not configure security services

Choose this option if you do not want to enable security.

The installer takes you to the next page to configure optional features of SFCFS RAC.

Configure security automatically

Choose this option to use an external root broker.

Enter the name of the root broker that is already configured for your enterprise environment, and click **Validate**. The installer configures the cluster in secure mode.

Configure one node as RAB and the others as AB

Select the system that you want to configure as RAB node.

The installer configures the cluster in secure mode.

Click **Next**.

- 8** On the Optional Configuration page, decide the optional VCS features that you want to configure. Click the corresponding tab to specify the details for each option:

- | | |
|-------------------|--|
| Virtual IP | <ul style="list-style-type: none">■ 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 | <ul style="list-style-type: none">■ Reset the password for the Admin user, if necessary.■ Click Add to add a new user.
Specify the user name, password, and user privileges for this user. |
| SMTP | <ul style="list-style-type: none">■ 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 | <ul style="list-style-type: none">■ 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 Cluster File System for Oracle RAC Installation and Configuration Guide* for instructions to set up SFCFS RAC global clusters.

- Select the **Configure GCO** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select a NIC.
- Enter a virtual IP address and value for the netmask.

Click **Next**.

- 9 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.
- 10 On the Start Processes page, click **Next** after the installer performs the configuration based on the details you provided and starts all the processes successfully.
- 11 Click **Next** to complete the process of configuring SFCFS 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.

Installing and configuring SFCFS RAC using a response file

This chapter includes the following topics:

- [About response files](#)
- [Installing and configuring SFCFS RAC](#)
- [Sample response file for installing and configuring SFCFS RAC](#)
- [Response file variables to install or uninstall SFCFS RAC](#)
- [Response file variables to configure SFCFS RAC](#)

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 SFCFS RAC
- Uninstalling SFCFS RAC

[Table 7-1](#) lists the various options available for creating or obtaining a response file.

Table 7-1 Options for obtaining a response file

Option	Description
Create a response file	<p>Create a response file based on the response file template provided with SFCFS RAC.</p> <p>The file is located at <code>/opt/VRTSvcs/rac/install</code>.</p>
Reuse or customize the response files generated by an installation	<p>The Veritas installation programs generate a response file during the installation, configuration, or uninstallation of SFCFS RAC.</p> <p>The response file generated by the installer is located in the following directory:</p> <pre>/opt/VRTS/install/logs/installsfcrac-installernumber\ /installsfcrac-installernumber.response file</pre> <p>You can reuse or customize the response files as needed.</p> <p>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 SFCFS RAC configuration option.</p>
Use the installation simulator	<p>Create a response file by specifying the <code>-makeresponsefile</code> option with the SFCFS RAC installer.</p> <p>Mount the product disc and navigate to the folder that contains the installation program. Start the installation program.</p> <pre># ./installsfcrac -makeresponsefile</pre> <p>The option creates a response file by simulating an installation. The response file is created in the directory <code>/opt/VRTS/install/logs/</code>.</p> <p>Note: You can use the <code>-makeresponsefile</code> option to create response files only for installing, configuring, and uninstalling SFCFS RAC.</p> <p>For more information:</p> <p>See “About the installation simulator” on page 107.</p>

At the end of the SFCFS 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 SFCFS 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*.

About the installation simulator

The SFCFS RAC installer includes the option (`-makeresponsefile`) to simulate the installation, configuration, or uninstallation of the product. The simulation option steps through the installation script, including all of the preinstallation checks on the systems. However, the simulation does not actually install the RPMs, uninstall previously installed RPMs, or start or stop any processes.

Use the installation simulator in the following situations:

- To understand the information that is required when you install, configure, or uninstall SFCFS RAC
 - As the simulator steps through the same code that is used by the installer, you can use the simulation process to preview the installation or configuration steps without disrupting your existing installation. You can skip the preinstallation checks, if not required.
- To create a response file
 - The simulation process enables you to create a response file that can be used as a template for installing, configuring, or uninstalling SFCFS RAC. You can customize the response file, as required.

To simulate an installation or configuration, specify the `-makeresponsefile` option with the installer or product installation script at the command line.

To simulate an uninstallation, specify the `-makeresponsefile` option with the installer or the product uninstall script at the command line.

Response file syntax

The Perl statement syntax that is included in the response file varies, depending on whether “Scalar” or “List” values are required by the variables.

For example,

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=[“value”, “value”, “value”];
```

Installing and configuring SFCFS RAC

You can create a single response file or separate response files for installing and configuring SFCFS RAC.

The installer performs the following tasks:

- Installs SFCFS RAC.
- Configures SFCFS RAC.

The following sample procedure uses a single response file for installing and configuring SFCFS RAC.

To install and configure SFCFS 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 and configure SFCFS 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 105.

Note: You must replace the host names in the response file with that of the new systems in the cluster.

For a sample response file:

See "[Sample response file for installing and configuring SFCFS RAC](#)" on page 109.

See "[Sample response file for installing and configuring SFCFS RAC](#)" on page 109.

- 4 Mount the product disc and navigate to the product directory that contains the installation program.

5 Start the installation and configuration:

```
# ./installsfcfsrac -responsefile /tmp/response_file
```

Where /tmp/response_file is the full path name of the response file.

6 Complete the SFCFS 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 SFCFS RAC

The following sample response file installs and configures SFCFS RAC on two nodes, galaxy and nebula.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{rsh}=1;
$CFG{prod}="SFCFSRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=1234;
$CFG{vcs_clustername}="my_clus";
$CFG{vcs_lltlink1}{galaxy}="eth2";
$CFG{vcs_lltlink1}{nebula}="eth2";
$CFG{vcs_lltlink2}{galaxy}="eth3";
$CFG{vcs_lltlink2}{nebula}="eth3";
$CFG{vcs_userenpw}=[ qw(aPQiPKpMQlQQoYQkPN) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];

1;
```

Response file variables to install or uninstall SFCFS RAC

Table 7-2 Response file variables specific to installing or uninstalling SFCFS RAC

Variable	Description
CFG{opt}{install}	Installs SFCFS RAC RPMs. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	Instructs the installer to install SFCFS RAC RPMs based on the variable that has the value set to 1: <ul style="list-style-type: none"> ■ <code>installallpkgs</code>: Installs all RPMs ■ <code>installrecpkgs</code>: Installs recommended RPMs ■ <code>installminpkgs</code>: Installs minimum RPMs <p>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.</p> List or scalar: scalar Optional or required: required
CFG{accepteula}	Specifies whether you agree with the <code>EULA.pdf</code> file on the media. List or scalar: scalar Optional or required: required
\$CFG{opt}{vxkeyless}	Installs the product with keyless license. List of scalar: scalar Optional or required: optional
CFG{keys}{hostname}	List of keys to be registered on the system if the variable <code>\$CFG{opt}{vxkeyless}</code> is set to 0. List or scalar: scalar Optional or required: optional

Table 7-2 Response file variables specific to installing or uninstalling SFCFS RAC (*continued*)

Variable	Description
CFG{systems}	<p>List of systems on which the product is to be installed or uninstalled.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>
CFG{systemscfs}	<p>List of systems for configuration if secure environment prevents the installer to install SFCFS RAC on all systems at once.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>
CFG{systemscfg}	<p>List of systems for configuration if secure environment prevents the installer to install SFCFS RAC on all systems at once.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{prod}	<p>Defines the product to be installed or uninstalled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>
CFG{opt}{keyfile}	<p>Defines the location of an ssh keyfile that is used to communicate with all remote systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{at_rootdomain}	<p>Defines the name of the system where the root broker is installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{pkxpath}	<p>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.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table 7-2 Response file variables specific to installing or uninstalling SFCFS RAC (*continued*)

Variable	Description
CFG{opt}{tmppath}	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. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{donotinstall} {RPM}	Instructs the installation to not install the optional RPMs in the list. List or scalar: list Optional or required: optional
CFG{donotremove} {RPM}	Instructs the uninstallation to not remove the optional RPMs in the list. List or scalar: list Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
\$CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFCFS RAC RPMs. List or scalar: scalar Optional or required: optional

Response file variables to configure SFCFS RAC

[Table 7-3](#) lists the response file variables that you can define to configure SFCFS RAC.

Table 7-3 Response file variables specific to configuring SFCFS RAC

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (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 configured. (Required)
CFG{prod}	Scalar	Defines the product to be configured. (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 ssh as the communication method between systems. (Optional)

Table 7-3 Response file variables specific to configuring SFCFS RAC (continued)

Variable	List or Scalar	Description
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
\$CFG{uploadlogs}	Scalar	Defines 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 (the csgnic, csgvip, and csgnetmask variables) must be defined if any are defined. The same is true for the SMTP notification (the smtpserver, smtpprecp, and smtprsev variables), the SNMP trap notification (the snmpport, snmpcons, and snmpcsev variables), and the Global Cluster Option (the gconic, gcovip, and gconetmask variables).

[Table 7-4](#) lists the response file variables that specify the required information to configure a basic SFCFS RAC cluster.

Table 7-4 Response file variables specific to configuring a basic SFCFS RAC cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)

Table 7-4 Response file variables specific to configuring a basic SFCFS RAC cluster (*continued*)

Variable	List or Scalar	Description
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 7-5](#) lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 7-5 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 7-5 Response file variables specific to configuring private LLT over Ethernet (*continued*)

Variable	List or Scalar	Description
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	<p>Defines a low-priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication.</p> <p>If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on.</p> <p>You must enclose the system name within double quotes.</p> <p>(Optional)</p>

Table 7-6 lists the response file variables that specify the required information to configure LLT over UDP.

Table 7-6 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	<p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p>
CFG{vcs_udplink<n>_address} {<system1>}	Scalar	<p>Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1.</p> <p>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.</p> <p>(Required)</p>

Table 7-6 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)
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)

Response file variables to configure SFCFS RAC

Table 7-7 lists the response file variables that specify the required information to configure virtual IP for SFCFS RAC cluster.

Table 7-7 Response file variables specific to configuring virtual IP for SFCFS 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 7-8 lists the response file variables that specify the required information to configure the SFCFS RAC cluster in secure mode.

Table 7-8 Response file variables specific to configuring SFCFS RAC cluster in secure mode

Variable	List or Scalar	Description
CFG{at_rootdomain}	Scalar	Defines the name of the system where the root broker is installed. (Optional)
CFG{at_rootbroker}	Scalar	Defines the root broker's name.
CFG{vcs_securitymenuopt}	Scalar	Specifies the menu option to choose to configure the cluster in secure mode. <ul style="list-style-type: none"> ■ 1—Automatic ■ 2—Semi-automatic ■ 3—Manual (Optional)

Table 7-8 Response file variables specific to configuring SFCFS RAC cluster in secure mode (*continued*)

Variable	List or Scalar	Description
CFG{vcs_vssdefport}	Scalar	Specifies the default port address of the root broker. (Optional)
CFG{vcs_roothashpath}	Scalar	Specifies the path of the root hash file. (Optional)
CFG{vcs_ab_prplname} {system}	Scalar	Specifies the authentication broker's principal name on system. (Optional)
CFG{vcs_ab_password} {system}	Scalar	Specifies the authentication broker's password on system. (Optional)
CFG{vcs_blobpath} {system}	Scalar	Specifies the path of the encrypted BLOB file for system. (Optional)

Table 7-9 lists the response file variables that specify the required information to configure VCS users.

Table 7-9 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	<p>List of encoded passwords for VCS users.</p> <p>The value in the list can be "Administrators Operators Guests."</p> <p>Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list.</p> <p>(Optional)</p>
CFG{vcs_username}	List	<p>List of names of VCS users.</p> <p>(Optional)</p>

Table 7-9 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 7-10 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 7-10 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_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, and SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table 7-11 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 7-11 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, and SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

[Table 7-12](#) lists the response file variables that specify the required information to configure SFCFS RAC global clusters.

Table 7-12 Response file variables specific to configuring SFCFS 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 SFCFS RAC

Performing post-installation and configuration tasks

This chapter includes the following topics:

- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Verifying LLT, GAB, and cluster operation](#)

About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

See “[Enabling LDAP authentication for clusters that run in secure mode](#)” on page 125.

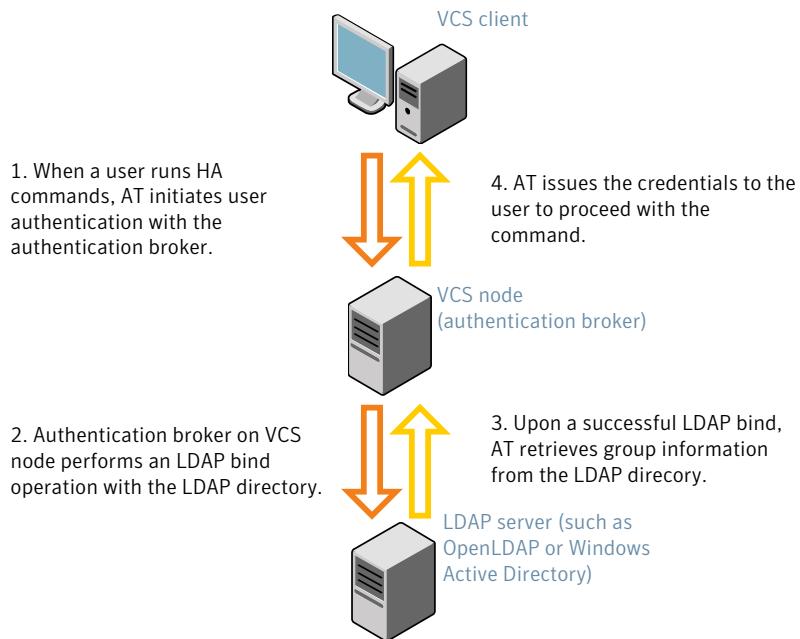
If you have not already added VCS users during installation, you can add the users later.

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

[Figure 8-1](#) depicts the SFCFS RAC cluster communication with the LDAP servers when clusters run in secure mode.

Figure 8-1

Client communication with LDAP servers



See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is posixAccount)
 - UserObject Attribute (the default is uid)
 - User Group Attribute (the default is gidNumber)
 - Group Object Class (the default is posixGroup)
 - GroupObject Attribute (the default is cn)
 - Group GID Attribute (the default is gidNumber)
 - Group Membership Attribute (the default is memberUid)
- URL to the LDAP Directory

- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 5.0.32.0 or later.

```
# /opt/VRTSat/bin/vssat showversion
vssat version: 5.0.32.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSat/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com" \
--server_url "ldap://my_opendap_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 SFCFS RAC nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, vcsadmin1.

```
galaxy# /opt/VRTSat/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:2821
```

```
Enter password for vcsadmin1: #####
```

```
authenticate
-----
-----
```

```
Authenticated User vcsadmin1
-----
```

3 Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

4 Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster rac_cluster101 (
    SecureClus = 1
    Administrators = {
        "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/
        DC=com@myenterprise.symantecdomain.com" }
    AdministratorGroups = {
        "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/
        DC=com@myenterprise.symantecdomain.com" }
)
...
...
```

5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=myenterprise.symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute      Value
galaxy       Attribute      RUNNING
nebula        Attribute      RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFS RAC node using the VCS Cluster Manager (Java Console).

7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool atldapconf using the -d option. The -d option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSat/bin/atldapconf -d  
-s domain_controller_name_or_ipaddress  
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSat/bin/atldapconf -d -s 192.168.20.32 \  
-u Administrator -g "Domain Admins"  
Search User provided is invalid or Authentication is required to  
proceed further.  
Please provide authentication information for LDAP server.
```

Username/Common Name: **symantecdomain\administrator**
Password:

Attribute file created.

- 2 Run the LDAP configuration tool atldapconf using the -c option. The -c option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSat/bin/atldapconf -c -d windows_domain_name
```

For example:

```
# /opt/VRTSat/bin/atldapconf -c -d symantecdomain.com  
Attribute list file not provided, using default AttributeList.txt.  
CLI file name not provided, using default CLI.txt.
```

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool atldapconf using the -x option. The -x option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSat/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSat/bin/vssat listldapdomains
```

```
Domain Name : symantecdomain.com
Server URL : ldap://192.168.20.32:389
SSL Enabled : No
User Base DN : CN=people,DC=symantecdomain,DC=com
User Object Class : account
User Attribute : cn
User GID Attribute : gidNumber
Group Base DN : CN=group,DC=symantecdomain,DC=com
Group Object Class : group
Group Attribute : cn
Group GID Attribute : cn
Auth Type : FLAT
Admin User :
Admin User Password :
Search Scope : SUB
```

- 5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute     Value
galaxy       Attribute    RUNNING
nebula        Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFS 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.

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 131.
- 4 Verify GAB operation.
See “[Verifying GAB](#)” on page 134.
- 5 Verify the cluster operation.
See “[Verifying the cluster](#)” on page 135.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

- 1** Log in as superuser on the node galaxy.
- 2** Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
  Node          State      Links
*0 galaxy      OPEN       2
  1 nebula      OPEN       2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
  Node          State      Links
* 0 galaxy     OPEN       2
  1 nebula      OPEN       2
  2 saturn      OPEN       1
```

- 3** Log in as superuser on the node nebula.
- 4** Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
  Node          State      Links
  0 galaxy      OPEN       2
*1 nebula      OPEN       2
```

- 5** To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv active
```

The output on galaxy resembles:

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
- A MAC address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port Usage      Cookie
    0   gab        0x0
          opens:   0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
          connects: 0 1
    7   gab        0x7
          opens:   0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
          connects: 0 1
   63   gab        0x1F
          opens:   0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
          connects: 0 1
```

Verifying GAB

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information. The output displays the nodes that have membership with the modules you installed and configured. You can use GAB port membership as a method of determining if a specific component of the SFCFS RAC stack communicates with its peers.

Table 8-1 lists the different ports that the software configures for different functions.

Table 8-1 GAB port description

Port	Function
a	GAB
d	Oracle Disk Manager (ODM)
f	Cluster File System (CFS)
h	Veritas Cluster Server (VCS: High Availability Daemon)
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)

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 d gen  ada409 membership 01
Port f gen  ada41c membership 01
Port h gen  ada40f membership 01
Port u gen  ada41a membership 01
Port v gen  ada416 membership 01
Port w gen  ada418 membership 01
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for a description of system states and the transitions between them.

To verify the cluster

- 1 To verify the status of the cluster, type the following command:

```
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
B   VxSS           galaxy       Y      N      ONLINE
B   VxSS           nebula       Y      N      ONLINE
```

Note that the VxSS service group is displayed only if you have configured the cluster in secure mode.

- 2 Review the command output for the following information:

- The system state

If the value of the system state is RUNNING, the cluster is successfully started.

- The cvm group state

In the sample output, the group state lists the cvm group, which is ONLINE on both the nodes galaxy and nebula.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for information about the system attributes for VCS.

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	Wed 14 Oct 2009 17:22:48
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.1.10.0
galaxy	FencingWeight	0
galaxy	Frozen	0

galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 0 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO
galaxy	Limits	
galaxy	LinkHbStatus	
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	
galaxy	SourceFile	./main.cf
galaxy	SwapThresholdLevel	Critical 90 Warning 80 Note 70 Info 60
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	DR
galaxy	VCSMode	

3

Section

Upgrade of SFCFS RAC

- [Chapter 9. Preparing to upgrade](#)
- [Chapter 10. Performing a full upgrade to SFCFS RAC 5.1 SP1](#)
- [Chapter 11. Performing a phased upgrade to SFCFS RAC 5.1 SP1](#)
- [Chapter 12. Performing a rolling upgrade to SFCFS RAC 5.1 SP1](#)
- [Chapter 13. Performing post-upgrade tasks](#)

Preparing to upgrade

This chapter includes the following topics:

- [About upgrading SFCFS RAC](#)
- [About types of upgrade](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade to SFCFS RAC 5.1 SP1](#)

About upgrading SFCFS RAC

SFCFS RAC supports various ways of upgrading your cluster to the latest version. Choose a method that best suits your environment and planned upgrade path.

For the supported types of upgrade:

See “[About types of upgrade](#)” on page 141.

For the supported upgrade paths:

See “[Supported upgrade paths](#)” on page 142.

Note: SFCFS RAC software must be at the same version across all nodes in an SFCFS RAC cluster after the upgrade, that is 5.1 Service Pack 1.

About types of upgrade

[Table 9-1](#) lists the supported types of upgrade.

Table 9-1 Types of upgrade

Type of upgrade	Method of upgrade	Procedures
Full upgrade	Veritas script-based installation programs <ul style="list-style-type: none">■ Interactive mode■ Non-interactive mode using response files	Complete the following steps: <ul style="list-style-type: none">■ Preparing to upgrade■ Upgrading to SFCFS RAC 5.1 SP1 See the chapter <i>Performing a full upgrade to SFCFS RAC 5.1 SP1</i>.■ Completing post-upgrade tasks See the chapter "Performing post-upgrade tasks".
Phased upgrade	Combination of manual steps and the Veritas script-based installation programs	Complete the steps in the chapter <i>Performing a phased upgrade to SFCFS RAC 5.1 SP1</i> .
Rolling upgrade	Veritas script-based installation programs	Complete the steps in the chapter <i>Performing a rolling upgrade to SFCFS RAC 5.1 SP1</i> .

Supported upgrade paths

Table 9-2 lists the supported upgrade paths if you plan to upgrade SFCFS RAC on RHEL or OEL.

Table 9-2 Supported upgrade paths on RHEL and OEL

From product version	On OS version	To SFCFS RAC version	To OS version	Supported upgrade type
SFCFS RAC 5.0 MP3/5.0 MP4	RHEL 4 Update 3 to Update 7 RHEL 5 Update 1 to Update 3 OEL 4 Update 4 to Update 8 OEL 5 Update 1 to Update 3	SFCFS RAC 5.1 SP1	RHEL 5 Update 4 or later OEL 5 Update 4 or later	Full or phased upgrade Note: If you are planning a major operating system upgrade, for example from RHEL 4 to RHEL 5, you must uninstall the existing version of SFCFS RAC and install SFCFS RAC 5.1 SP1.

Table 9-2 Supported upgrade paths on RHEL and OEL (*continued*)

From product version	On OS version	To SFCFS RAC version	To OS version	Supported upgrade type
SFCFS RAC 5.1	RHEL 5 Update 3 or OEL 5 Update 3	SFCFS RAC 5.1 SP1	RHEL 5 Update 4 or later	Full or rolling upgrade
	RHEL 5 Update 4 or OEL 5 Update 4		OEL 5 Update 4 or later	
SFCFS RAC 5.1 Patch 1	RHEL 5 Update 3 or OEL 5 Update 3	SFCFS RAC 5.1 SP1	RHEL 5 Update 4 or later	Full or rolling upgrade
	RHEL 5 Update 4 or OEL 5 Update 4		OEL 5 Update 4 or later	
SFCFS RAC 5.1 Rolling Patch 1 (RP1) or Rolling Patch 2 (RP2)	RHEL 5 Update 3 or OEL 5 Update 3	SFCFS RAC 5.1 SP1	RHEL 5 Update 4 or later	Full or rolling upgrade
	RHEL 5 Update 4 or OEL 5 Update 4		OEL 5 Update 4 or later	

[Table 9-3](#) lists the supported upgrade paths if you plan to upgrade SFCFS RAC on SLES.

Note: If you are planning a major operating system upgrade, for example from SLES 10 to SLES 11, then the full upgrade involves uninstallation of the existing version of SFCFS RAC and installation of SFCFS RAC 5.1 SP1.

Table 9-3 Supported upgrade paths on SLES

From product version	On OS version	To SFCFS RAC version	To OS version	Supported upgrade type
SFCFS RAC 5.0 MP3/5.0 MP4	SLES 9 SP4	SFCFS RAC 5.1 SP1	SLES 10 SP3	Full or phased upgrade
	SLES10 SP2		SLES 11	
	SLES10 SP3			
	SLES11			
SFCFS RAC 5.1	SLES 10 SP2	SFCFS RAC 5.1 SP1	SLES 10 SP3 SLES 11	Full or rolling upgrade

Table 9-3 Supported upgrade paths on SLES (continued)

From product version	On OS version	To SFCFS RAC version	To OS version	Supported upgrade type
SFCFS RAC 5.1 Patch 1	SLES 10 SP2 SLES 10 SP3	SFCFS RAC 5.1 SP1	SLES 10 SP3 SLES 11	Full or rolling upgrade
SFCFS RAC 5.1 Rolling Patch 1 (5.1 RP1) or Rolling Patch 2 (RP2)	SLES 10 SP2 SLES 10 SP3	SFCFS RAC 5.1 SP1	SLES 10 SP3 SLES 11	Full or rolling upgrade

Preparing to upgrade to SFCFS RAC 5.1 SP1

Perform the preparatory steps in this section if you are performing a full upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

To prepare to upgrade SFCFS RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleASMTypes.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.backup
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.backup
```

- 3 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

Use native application commands to stop the application.

- 4 Stop the Oracle RAC resources.

```
$ srvctl stop database -d db_name
```

5 Stop Oracle Clusterware:

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

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
```


Performing a full upgrade to SFCFS RAC 5.1 SP1

This chapter includes the following topics:

- [About full upgrades](#)
- [Upgrading SFCFS RAC and the operating system \(major OS upgrade\)](#)
- [Upgrading SFCFS RAC and the operating system \(minor OS upgrade\)](#)
- [Upgrading SFCFS RAC using the Veritas script-based installation program](#)
- [Upgrading SFCFS RAC with the Veritas Web-based installer](#)
- [Upgrading SFCFS RAC using a response file](#)

About full upgrades

A full upgrade involves upgrading all the nodes in the cluster at the same time. The cluster remains unavailable for the duration of the upgrade.

You can perform the upgrade using one of the following Veritas script-based installation programs:

- Common product installer (`installer`)
The common product installer provides menu options for installing and configuring multiple Veritas products.
- SFCFS RAC installation script (`installsfcfsrac`)
The SFCFS RAC installation script provides menu options for installing and configuring SFCFS RAC.

Note: If you obtained SFCFS RAC from an electronic download site, you must use the product installer (`installsfcfsrac`) instead of the common product installer (`installer`).

You can also perform a full upgrade using a response file. You can create a response file using the response file template or customize an installer-generated response file.

Note: After you upgrade to version 5.1 SP1 from version 5.1 or related 5.1 rolling patches, Symantec recommends that you do not roll back to version 5.1 or corresponding rolling patches.

For more information about response files:

See “[About response files](#)” on page 105.

Upgrading SFCFS RAC and the operating system (major OS 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 SFCFS RAC.

To upgrade SFCFS RAC and the operating system

- 1 Uninstall SFCFS RAC.

For instructions, see the chapter *Uninstalling SFCFS RAC from a cluster*.

- 2 Upgrade the operating system.

For instructions, see the operating system documentation.

- 3 Install SFCFS RAC 5.1 Service Pack 1.

For instructions, see the chapter *Installing and configuring SFCFS RAC*.

- 4 Complete the post-installation and configuration tasks.

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

- 5 Install Oracle RAC.

For instructions, see the chapter *Installing and configuring Oracle RAC* in this document.

Upgrading SFCFS RAC and the operating system (minor OS 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 SP3, along with SFCFS RAC.

To upgrade SFCFS RAC and the operating system

1 Upgrade to SFCFS RAC 5.1 SP1.

For upgrades from version 5.0x Upgrade to SFCFS RAC 5.1 SP1. You can perform a full or phased upgrade.

For full upgrade, using the script-based installation program:

See “[Upgrading SFCFS RAC using the Veritas script-based installation program](#)” on page 150.

For full upgrade, using the response file:

See “[Upgrading SFCFS RAC using a response file](#)” on page 152.

For phased upgrade:

See “[About phased upgrade](#)” on page 157.

For upgrades from version 5.1x Upgrade to SFCFS RAC 5.1 SP1. You can perform a full or rolling upgrade.

For full upgrade:

See “[Upgrading SFCFS RAC using the Veritas script-based installation program](#)” on page 150.

For rolling upgrade:

See the chapter *Performing a rolling upgrade to SFCFS RAC 5.1 SP1*.

2 Relink the ODM library:

See “[Relinking with ODM](#)” on page 195.

3 Upgrade Oracle RAC, if required.

For instructions, see the chapter *Upgrading Oracle RAC and migrating the database* in this document.

4 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 SFCFS RAC:

```
# ./installsfcfsrac -stop
```

- Upgrade the operating system. For instructions, see the operating system documentation.

- Reboot the nodes:

```
# shutdown -r now
```

Note: The Oracle resources might start automatically after the upgrade.

Upgrading SFCFS RAC using the Veritas script-based installation program

Perform the steps in the following procedure if you plan to upgrade only SFCFS RAC.

Note: After you upgrade to version 5.1 SP1 from version 5.1 or related 5.1 rolling patches, Symantec recommends that you do not roll back to version 5.1 or corresponding rolling patches.

To perform a full upgrade of SFCFS RAC

- 1 Mount the product disc containing the SFCFS RAC 5.1 SP1 software.
- 2 Change to the directory containing the installation program. The program is located in the product directory.
- 3 Start the installation program:

```
# ./installsfcfsrac -upgrade galaxy nebula
```

- 4 Select a suitable option to install the SFCFS RAC packages.
- 1) Install minimal required Storage Foundation Cluster File System for Oracle RAC rpms
 - 2) Install recommended Storage Foundation Cluster File System for Oracle RAC rpms
 - 3) Install all Storage Foundation Cluster File System for Oracle RAC 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 message that the current versions will be upgraded to 5.1 Service Pack 1.

Review the messages and make sure that you meet the requirements before proceeding with the upgrade.

- 5 Press **Enter** to continue with the upgrade.

Depending on the package installation option selected, the installer displays the list of packages that will be installed and prompts for the stoppage of SFCFS RAC processes before the upgrade.

- 6 Enter **y** to stop the SFCFS RAC processes.

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

The installer stops the processes and uninstalls SFCFS RAC. After the uninstallation, the installer installs SFCFS RAC 5.1 SP1 and starts SFCFS RAC 5.1 SP1 on all the nodes.

Upgrading SFCFS RAC with the Veritas Web-based installer

This section describes how to upgrade SFCFS RAC with the Veritas Web-based installer. 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.

To upgrade SFCFS 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 97.

- 3 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.

- 4 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.

- 5 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.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

Upgrading SFCFS RAC using a response file

Perform the steps in the following procedure to upgrade to SFCFS RAC 5.1 SP1 using a response file.

To upgrade SFCFS 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 105.

For response file variable definitions:

See "[Response file variables to upgrade SFCFS RAC](#)" on page 153.

For a sample response file:

See "[Sample response file for upgrading SFCFS RAC](#)" on page 154.

- 3 Navigate to the directory containing the SFCFS RAC installation program.

- 4 Start the installation:

```
# ./installsfcfsrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

- 5 Complete the post-upgrade steps.

Response file variables to upgrade SFCFS RAC

[Table 10-1](#) lists the response file variables that you can define to upgrade SFCFS RAC.

Table 10-1 Response file variables specific to upgrading SFCFS RAC

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades SFCFS RAC RPMs. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{opt}{systems}	List	List of systems on which the product is to be upgraded. (Required)
CFG{prod}	Scalar	Defines the product to be upgraded. (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)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{patchpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems. (Optional)

Table 10-1 Response file variables specific to upgrading SFCFS 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 SFCFS RAC

The following sample response file upgrades SFCFS RAC to version 5.1 Service Pack 1 on nodes, galaxy and nebula.

```
our  %CFG;

$CFG{accepteula}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{patchupgrade}=1;
$CFG{prod}="SFCFSRAC51";
```

```
$CFG{systems}=[ qw(galaxy nebula) ];  
$CFG{vcs_allowcomms}=1;  
1;
```


Performing a phased upgrade to SFCFS RAC 5.1 SP1

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing phased upgrade of SFCFS RAC from version 5.0](#)

About phased upgrade

The phased upgrade methodology involves upgrading half of the nodes in the cluster at a time.

SFCFS RAC supports phased upgrade of the cluster for version 5.0 and related maintenance packs.

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 SFCFS 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 SFCFS RAC from version 5.0

Before starting the upgrade on the first half of the cluster, back up the configuration files.

To perform phased upgrade of SFCFS RAC from version 5.0

- 1 Switch failover groups from the first half of the cluster, for example, from galaxy, to the second half of the cluster, for example to jupiter and mercury.

```
# hagrpg -switch failover_group -to jupiter  
  
# hagrpg -switch failover_group -to mercury
```

- 2 On the first half of the cluster, log in as the Oracle user and shut down the instances:

```
$ svrctl stop instance -d database_name \  
-i instance_name
```

- 3 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 4 On the first half of the cluster, 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 or shared volumes. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs  
# fuser -cu /mount_point
```

■ Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

- 5 On first half of the cluster, stop all VxVM and CVM volumes (for each disk group) that are not managed by VCS:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 6 On the first half of the cluster, take all the VCS service groups offline:

```
# hagrp -offline group_name -sys galaxy
```

```
# hagrp -offline group_name -sys nebula
```

- 7 Verify that all the VCS service groups are offline on the first half of the nodes in the cluster:

```
# hagrp -state group_name
```

- 8 Freeze the nodes in the first half of the cluster:

```
# haconf -makerw
# hasys -freeze -persistent galaxy
# hasys -freeze -persistent nebula
# haconf -dump -makero
```

- 9 Verify that only ports a, b, d, and h are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

- 10 On first half of the cluster, upgrade the operating system, if applicable.

For supported operating systems:

See “[Supported Linux operating systems](#)” on page 36.

For instructions, see the operating system documentation.

- 11 Make sure that you can run SSH or RSH from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 12 On the first half of the cluster, upgrade SFCFS RAC. Navigate to the directory that contains the installation program and run the program.

```
# ./installsfcfsrac 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.

- 13 On the first half of the cluster, relink the ODM library with Oracle.
See “[Relinking with ODM](#)” on page 195.
- 14 On the second half of the cluster, log in as the Oracle user on one of the nodes and shut down the instances:

```
$ svrctl stop instance -d database_name \
-i instance_name
```

Note: The downtime starts now.

- 15 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 16 On the second half of the cluster, 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 or shared volumes. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
# fuser -cu /mount_point
```

- Unmount the VxFS file system:

```
# umount /mount_point
```

- 17 On the second half of the cluster, stop all VxVM and CVM volumes (for each disk group) that are not managed by VCS:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 18 On the second half of the cluster, unfreeze all the VCS service groups:

```
# haconf -makerw
# hagrp -unfreeze group_name -persistent
# haconf -dump -makero
```

- 19 On the second half of the cluster, take all the VCS service groups offline:

```
# hagrp -offline group_name -sys jupiter
# hagrp -offline group_name -sys mercury
```

- 20 Verify that all the VCS service groups are offline on the second half of the cluster:

```
# hagrp -state group_name
```

- 21 Stop VCS on the second half of the cluster:

```
# hastop -local
```

- 22 On the second half of the cluster, stop the following SFCFS RAC modules: Fencing, ODM, VCS, GAB, and LLT

```
# /etc/init.d/vxqlm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 23 Restart the nodes in the first half of the cluster. When the nodes in the first half of the cluster come up, no GAB ports will be shown in the output of the gabconfig -a command.

```
# shutdown -r now
```

- 24 On first half of the cluster, force GAB to form a cluster after the upgraded nodes reboot.

```
# gabconfig -x
```

After GAB seeds the cluster membership, the GAB ports a, b, d and h will appear in the gabconfig -a command output.

- 25 On first half of the cluster, unfreeze the nodes:

```
# haconf -makervw
# hasys -unfreeze -persistent galaxy
# hasys -unfreeze -persistent nebula
# haconf -dump -makero
```

- 26 On the first half of the cluster, bring the VCS service groups online:

For parallel service groups:

```
# hagrp -online group_name -sys galaxy
# hagrp -online group_name -sys nebula
```

For failover service groups:

```
# hagrp -online group_name -any
```

Note: The downtime ends here.

Once the cvm service group comes online, the GAB ports v, w, and f come online; all the service groups pertaining to the CFS mounts also come online automatically.

The failover service groups must be brought online manually using the above command.

- 27 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 28 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

- 29 On the second half of the cluster, upgrade the operating system, if applicable.

Before you upgrade the operating system on the second half of the cluster, perform the following steps to ensure that LLT, GAB, fencing, and VCS do not start automatically after rebooting the systems:

- For RHEL:

```
# chkconfig llt off
# chkconfig gab off
# chkconfig vxfen off
# chkconfig vcs off
```

- For SLES:

```
# insserv -r llt
# insserv -r gab
# insserv -r vxfen
# insserv -r vcs
```

For supported operating systems:

See “[Supported Linux operating systems](#)” on page 36.

For instructions, see the operating system documentation.

- 30 Make sure that you can run SSH or RSH from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 31 On the second half of the cluster, upgrade SFCFS RAC. Navigate to the directory that contains the installation program and run the program.

```
# ./installsfcfsrac jupiter mercury
```

- 32 On the second half of the cluster, relink the ODM library with Oracle.

See “[Relinking with ODM](#)” on page 195.

- 33** Verify that the cluster UUID on the nodes in the second half of the cluster is the same as the cluster UUID on the nodes in the first half of the cluster.

Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus \
-display node_name
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first half of the cluster to the nodes in the second half of the cluster.

For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -copy \
-from_sys galaxy -to_sys jupiter mercury
```

- 34** Restart the nodes in the second half of the cluster.

When the nodes in the second half of the cluster come up, all the GAB ports a, b, d, h, v, w and f will be online. All the CFS mount service groups also come online automatically.

```
# shutdown -r now
```

- 35** On the second half of the cluster, manually mount the VxFS and CFS file systems that are not managed by VCS.

- 36** On the second half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Performing a rolling upgrade to SFCFS RAC 5.1 SP1

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Prerequisites for a rolling upgrade](#)
- [Preparing to perform a rolling upgrade to SFCFS RAC 5.1 SP1](#)
- [Performing a rolling upgrade using the installer](#)
- [Performing a rolling upgrade of SFCFS RAC using the Web-based installer](#)

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

Rolling upgrades take two discrete phases. In the first, you upgrade the product kernel RPMs. In the second, you upgrade the non-kernel RPMs such as VCS RPMs and agent RPMs.

You can perform a rolling upgrade from 5.1, 5.1 P1, 5.1 RP1, or 5.1 RP2 to 5.1 SP1.

Prerequisites for a rolling upgrade

Meet the following prerequisites before performing a rolling upgrade:

- Make sure the product that you want to upgrade supports rolling upgrades.
- Split your clusters into sub-clusters for the upgrade to keep the service groups available during upgrade.
- Make sure you logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.

Preparing to perform a rolling upgrade to SFCFS RAC 5.1 SP1

Perform the preparatory steps in this section if you are performing a rolling upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

Note: Perform the steps on the first subcluster.

To prepare to upgrade SFCFS RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system:`main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleASMTypes.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`

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
```

- 3 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrp -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.

```
$ srvctl stop database -d db_name
```

- 5 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
```

Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

Performing a rolling upgrade on kernel RPMs: phase 1

Note that in the following instructions that a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

To perform the rolling upgrade on kernel RPMs: phase 1

- 1 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
./installer -upgrade_kernelpkgs nodeA
```

- 2 Note that if the boot disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.

- 3 The installer checks system communications, RPM versions, product versions, and completes prechecks.

It then upgrades applicable product kernel RPMs.

- 4 The installer loads new kernel modules.

- 5 The installer starts all the relevant processes and brings all the service groups online.

- 6 Manually mount the VxFS and CFS file systems that are not managed by VCS.

- 7 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 8 Start Oracle Clusterware:

```
# $CRS_HOME/bin/crsctl start crs
```
- 9 Relink the ODM library:
See “[Relinking with ODM](#)” on page 195.
- 10 If the boot disk is encapsulated, reboot the first sub-cluster's system.
- 11 Before you proceed to phase 2, complete step 1 to 10 on the second subcluster.

Performing a rolling upgrade on non-kernel RPMs: phase 2

In this phase installer installs all non-kernel RPMs on all the nodes in cluster and restarts VCS cluster.

To perform the rolling upgrade on non-kernel RPMs: phase 2

- 1 Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installer -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

- 2 The installer checks system communications, RPM versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 The installer upgrades non-kernel RPMs.
- 4 The installer checks system communications, RPM versions, product versions, and completes prechecks. It verifies completion of phase 1. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.
- 5 Verify the cluster's status:

```
# hastatus -sum
```

- 6 If you want to upgrade 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 SFCFS RAC:

```
# ./installsfcfsrac -stop
```

- Upgrade the operating system. For instructions, see the operating system documentation.
- Reboot the nodes:

```
# shutdown -r now
```

Performing a rolling upgrade of SFCFS 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.

The rolling upgrade is divided into two phases. In the first phase, the installer upgrade kernel RPMs. In the second phase, it upgrades non-kernel RPMs. The second phase is required for upgrades that have high-availability components. When you perform a rolling upgrade, you need to divide the number of systems that you plan to upgrade roughly in half. Half of the systems' available capacity is needed to take over processes during the rolling upgrade.

To start the rolling upgrade—phase 1

- 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 97.

- 3 In the Task pull-down menu, select **Rolling Upgrade**.

In the Product pull-down menu, select the product that you want to upgrade using a rolling upgrade.

Note that the Upgrade Kernel packages for Rolling Upgrade Phase-1 radio button is selected.

Click the **Next** button to proceed.

- 4 In the Systems Names field, enter the sub-cluster's system names. Separate system names with a space.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

- 5 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.
- 6 When the upgrade completes, perform step 3 through step 6 on the second subcluster.

To upgrade the non-kernel components—phase 2

- 1 In the Task pull-down menu, make sure that **Rolling Upgrade** and the product are selected.

Note that the Upgrade Non-Kernel packages for Rolling Upgrade Phase-2 radio button is selected.

Click the **Next** button to proceed.

- 2 In the Systems Names field, enter the names of all the systems that you want to upgrade. Separate system names with a space.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

- 3 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.

Performing post-upgrade tasks

This chapter includes the following topics:

- [Setting or changing the product license level](#)
- [Upgrading disk layout versions](#)
- [Upgrading VxVM disk group version](#)

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 SFCFS RAC*.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 6, Version 7, and Version 8. No prior versions can be created or mounted.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 8.

See the `vxfsconvert` or `vxupgrade` man pages.

For more information about disk layouts, see the *Veritas File System Administrator's Guide*.

Upgrading VxVM disk group version

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 160.

Run the following command on the master node to upgrade the disk group version:

```
# vxdg -T 160 upgrade diskgroup_name
```

The default Cluster Volume Manager protocol version is 100 and does not need to be upgraded.

Run the following command to verify the CVM protocol version:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the protocol version is not 100, run the following command to upgrade the version:

```
# /opt/VRTS/bin/vxdctl upgrade
```

4

Section

Installation and upgrade of Oracle RAC

- [Chapter 14. Before installing Oracle RAC](#)
- [Chapter 15. Installing Oracle RAC](#)
- [Chapter 16. Upgrading Oracle RAC and migrating the database](#)

Before installing Oracle RAC

This chapter includes the following topics:

- [About preparing to install Oracle RAC](#)
- [Preparing to install Oracle RAC manually](#)

About preparing to install Oracle RAC

The examples in this chapter assume a two-node cluster comprising the nodes galaxy and nebula.

Before installing Oracle RAC, review the Oracle installation manuals and the appropriate Oracle support Web sites.

Note: Some of the pre-installation tasks, wherever indicated in the document, must be done in accordance with the instructions in the Oracle installation manuals. The instructions for these tasks are not provided in this document.

Preparing to install Oracle RAC manually

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 176.

- 2 Set the kernel parameters.

See “[Setting the kernel parameters](#)” on page 176.

- 3 Create Oracle user and groups.

See “[Creating Oracle user and groups](#)” on page 177.

- 4 Create the storage for OCR and voting disk.
See “[Creating storage for OCR and voting disk manually](#)” on page 177.
- 5 Set up Oracle user equivalence on all nodes.
See “[Setting up user equivalence](#)” on page 189.
- 6 Edit the Oracle user profile.
See “[Editing the Oracle user profile](#)” on page 189.
- 7 Add the private IP addresses to the `/etc/hosts` file.
See “[Adding private IP addresses to the /etc/hosts file manually](#)” on page 190.

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 manages the virtual IP addresses.

The IP address and the corresponding host name should be registered in the domain name service (DNS) for each public network interface. 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).

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

Creating Oracle user and groups

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.

Before creating Oracle users and groups, see the Oracle documentation for information about creating the oinstall (Oracle inventory), dba group, and Oracle user.

Creating the Oracle user and groups

Note: When you create the user and group, make sure that you specify a user and group ID that is not in use.

To create the operating system Oracle user and group on each system

- 1 Create the 'oinstall' group on each system.

```
# groupadd -g 1000 oinstall  
# groupadd -g 1001 dba
```

- 2 Create the Oracle user and the user home directory on each system:

```
# useradd -g oinstall -u 1000 \  
-G dba -m -d /home/oracle oracle
```

Creating storage for OCR and voting disk manually

Symantec recommends that you place the OCR and voting disk on CVM raw volumes. Oracle Cluster Registry (OCR) is supported on Cluster File System (CFS) from Oracle RAC version 11.1.0.7 and later. Voting disk is not supported on Cluster File System (CFS).

To create OCR and voting disk volumes on raw volumes

- 1 Log in as the root user.
- 2 Determine the CVM master:

```
# vxrctl -c mode
```

- 3 On the master node, create a shared disk group:

```
# vxrdg -s init ocrvotedg Disk_2 Disk_3
```

- 4 Create mirrored volumes in the shared group for OCR and voting disk:

```
# vxassist -g ocrvotedg make ocrvol 300M nmirrors=2
```

```
# vxassist -g ocrvotedg make votevol 300M nmirrors=2
```

- 5 Set the ownership for the volumes:

```
# vxedit -g ocrvotedg set group=oinstall user=oracle mode=660  
ocrvol
```

```
# vxedit -g ocrvotedg set group=oinstall user=oracle mode=660  
votevol
```

- 6 Start the volume:

```
# vxvol -g ocrvotedg startall
```

- 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 178.

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 CFSMount resources to the VCS configuration.

Use one of the following ways to add the storage resources to the VCS configuration:

Use the command line interface	See “ To add the storage resources created on raw volumes to the VCS configuration using CLI ” on page 179.
Edit main.cf	See “ To add the storage resources to the VCS configuration using the main.cf file ” on page 180. This procedure requires all the nodes to be restarted.

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 raw volumes 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 ocrvotedg_crvotedg CVMVolDg cvm
# hares -modify ocrvotedg_crvotedg CVMDiskGroup ocrvotedg
# hares -modify ocrvotedg_crvotedg CVMVolume -add ocrvol
# hares -modify ocrvotedg_crvotedg CVMVolume -add votevol
# hares -modify ocrvotedg_crvotedg CVMActivation sw
```

- 3 Link the parent and child resources:

```
# hares -link ocrvotedg_crvotedg cvm_clus
```

- 4 Enable the resources:

```
# hares -modify ocrvotedg_crvotedg Enabled 1
# haconf -dump -makero
```

5 Verify the configuration of the ocrvote_voldg_ocrvotedg resource in main.cf.

```
CVMVolDg ocrvote_voldg_ocrvotedg (
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvol, votevol }
    CVMActivation = sw
)
ocrvote_voldg_ocrvotedg requires cvm_clus
```

6 Verify that the ocrvote_voldg_ocrvotedg resource is online on all systems in the cluster.

```
# hares -state ocrvote_voldg_ocrvotedg
```

Note: The ocrvote_voldg_ocrvotedg resource is reported offline though the underlying volume is online. You need to manually bring the resource online on each node.

To bring the resource online manually:

```
# hares -online ocrvote_voldg_ocrvotedg -sys galaxy
# hares -online ocrvote_voldg_ocrvotedg -sys nebula
```

To add the storage resources to the VCS configuration using the main.cf file

1 Log in to one of the systems as the root user.

2 Save your existing configuration:

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

If your configuration is not writable, a warning appears: "Cluster not writable".

You may ignore the warning.

3 Stop the VCS engine on all systems and leave the resources available.

```
# hastop -all -force
```

4 Back up the /etc/VRTSvcs/conf/config/main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp -p main.cf main.cf.`date +%m_%d_%Y-%H_%M_%S`
```

5 Modify the CVM service group.

```
CVMVolDg ocrvvote_voldg_ocrvotedg (
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvol, votevol }
    CVMActivation = sw
)
ocrvvote_voldg_ocrvotedg requires cvm_clus
```

6 Save and close the file.

7 Verify the syntax of the `main.cf` file:

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

8 Start the VCS engine on one of the nodes:

```
# hastart
```

9 Verify the status of VCS:

```
# hastatus
```

10 When "LOCAL_BUILD" is listed in the message column, start VCS on the other system with the following command:

```
# hastart
```

11 Verify that the CVM group is online:

```
# hagrp -state
```

12 Verify that the `ocrvvote_voldg_ocrvotedg` resource is online:

```
# hares -state
```

Setting CVM disk policies for OCR, voting disk, and data disk groups

Table 14-1 describes the CVM disk policies you need to set for Oracle Cluster Registry (OCR), voting disk, and data disk groups.

Table 14-1 Disk group fail policy for OCR, voting disk, and data disk groups

Disk group	Policy setting
OCR and voting disk	<p>leave</p> <p>If the disk groups containing the OCR and voting disk uses the default policy setting (<code>dgdisable</code>), and the CVM master node loses connectivity to the storage, then a cluster-wide panic results. The CVM master node disables the corresponding disk group after a storage loss and prevents access to volumes from any cluster node. This in turn prevents the slave nodes from performing I/O operations to shared volumes even though issues with storage may be confined to the master node. If Oracle's clusterware is unable to perform I/O operations to voting disks, the corresponding node is panicked and evicted from the cluster.</p> <p>Setting the policy to <code>leave</code> for the disk group ensures that only the node which loses connectivity to the storage is panicked.</p>
Data disk groups	<p>leave</p> <p>If the disk groups containing the Oracle data uses the default policy setting (<code>dgdisable</code>), and if one of the nodes in the cluster loses connectivity to the storage, then the disk group is disabled across the cluster. This causes the database to freeze on all nodes in the cluster.</p> <p>Setting the policy to <code>leave</code> for the disk group ensures that only the node which loses connectivity to the storage is panicked.</p>

Note: Symantec strongly recommends retaining the default setting (`global`) for the disk detach policy. For other disk detach policy options, see the *Veritas Volume Manager Administrator's Guide*.

To set the disk policies for OCR, voting disk, and data disk groups

- 1 Log into the CVM master node as the root user.
- 2 Set the disk group fail policy as follows:

```
# vxldg -g dg_name set dgfailpolicy=leave
```

Creating Oracle Clusterware and Oracle database home directories manually

You can create the Oracle Clusterware and Oracle database home directories either on local storage or on shared storage. This step is mandatory if you plan to place the directories on shared VxVM disks. If you plan to place the directories on storage local to each node, you may or may not perform this step. When the installer prompts for the home directories at the time of installing Oracle Clusterware and Oracle database, it creates the directories locally on each node, if they do not exist.

Use one of the following options to create the directories:

Local storage	See “ To create the file system and directories on local storage for Oracle Clusterware and Oracle RAC database ” on page 183.
Shared storage	See “ To create the file system and directories on shared storage for Oracle Clusterware and Oracle database ” on page 184.

To create the file system and directories on local storage for Oracle Clusterware and Oracle RAC database

The sample commands in the procedure are for node galaxy. Repeat the steps on each node of the cluster.

- 1 As the root user, create a VxVM local diskgroup `bindg_hostname` on each node:

```
# vxdg init bindg_galaxy Disk_1
```

- 2 Create a volume `binvol_hostname` on each node:

```
# vxassist -g bindg_galaxy make binvol_galaxy 12G
```

- 3 Create a filesystem with the volume, `binvol_hostname`, on each node.

```
# mkfs -t vxfs /dev/vx/dsk/bindg_galaxy/binvol_galaxy
```

- 4 Mount the filesystem (`/app`) on each node:

```
# mount -t vxfs /dev/vx/dsk/bindg_galaxy/binvol_galaxy
/app
```

- 5 Create the following directories for Oracle RAC (ORACLE_BASE, CRS_HOME, ORACLE_HOME) on each node:

```
# mkdir -p /app/oracle  
# mkdir -p /app/crshome  
# mkdir -p /app/oracle/orahome
```

- 6 Change the ownership and permissions on each node:

```
# chown -R oracle:oinstall /app  
# chmod -R 744 /app
```

- 7 In /etc/fstab add:

```
/dev/vx/dete/bindg_hostname/binvol_hostname \  
/app vxfs defaults 01
```

- 8 Repeat all the steps on each node of the cluster.

To create the file system and directories on shared storage for Oracle Clusterware and Oracle database

Perform the following steps on one of the nodes in the cluster.

- 1 As the root user, create a VxVM shared disk group bindg:

```
# vxrdg -s init bindg Disk_1
```

- 2 Create separate volumes for Oracle Clusterware (crsbinvol) and Oracle database (orabinvol):

```
# vxassist -g bindg make crsbinvol 5G  
# vxassist -g bindg make orabinvol 7G
```

- 3 Create the following directories for Oracle, ORACLE_BASE, CRS_HOME, ORACLE_HOME.

The file system and directories created on shared storage in this procedure are based on the following layout:

\$ORACLE_BASE /app/oracle
Both /app and /app/oracle are on local storage.

\$CRS_HOME /app/crshome
/app is on local storage.
/app/crshome is on shared storage.

\$ORACLE_HOME /app/oracle/orahome
/app/oracle is on local storage.
/app/oracle/orahome is on shared storage.

```
# mkdir -p /app/oracle  
# mkdir -p /app/crshome  
# mkdir -p /app/oracle/orahome
```

- 4 Create the file system with the volume orabinvol:

```
# mkfs -t vxfs /dev/vx/rdsk/bindg/crsbinvol  
# mkfs -t vxfs /dev/vx/rdsk/bindg/orabinvol
```

- 5 Mount the file systems. Perform this step on each node.

```
# mount -t vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \  
/app/crshome  
# mount -t vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \  
/app/oracle/orahome
```

6 Change the ownership and permissions on all nodes of the cluster.

Note: The ownership and permissions must be changed on all nodes of the cluster because `/app/oracle` must be owned by `oracle:oinstall`, otherwise `/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

```
# chown -R oracle:oinstall /app
# chmod -R 744 /app
```

7 Add the CVMVolDg and CFSMount resources to the VCS configuration.

See “[To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI](#)” on page 186.

To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI

1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

2 Configure the CVM volumes under VCS:

```
# hares -add crsorabin_voldg CVMVolDg cvm
# hares -modify crsorabin_voldg Critical 0
# hares -modify crsorabin_voldg CVMDiskGroup bindg
# hares -modify crsorabin_voldg CVMVolume -add crsbinvol
# hares -modify crsorabin_voldg CVMVolume -add orabinvol
# hares -modify crsorabin_voldg CVMActivation sw
```

3 Set up the file system under VCS:

```
# hares -add crsbin_mnt CFSMount cvm
# hares -modify crsbin_mnt Critical 0
# hares -modify crsbin_mnt MountPoint "/app/crshome"
# hares -modify crsbin_mnt BlockDevice \
"/dev/vx/dsk/bindg/crsbinvol"
# hares -add orabin_mnt CFSMount cvm
# hares -modify orabin_mnt Critical 0
# hares -modify orabin_mnt MountPoint "/app/oracle/orahome"
# hares -modify orabin_mnt BlockDevice \
"/dev/vx/dsk/bindg/orabinvol"
```

4 Link the parent and child resources:

```
# hares -link crsorabin_voldg cvm_clus
# hares -link crsbin_mnt crsorabin_voldg
# hares -link crsbin_mnt vxfsckd
# hares -link orabin_mnt crsorabin_voldg
# hares -link orabin_mnt vxfsckd
```

5 Enable the resources:

```
# hares -modify crsorabin_voldg Enabled 1
# hares -modify crsbin_mnt Enabled 1
# hares -modify orabin_mnt Enabled 1
# haconf -dump -makero
```

6 Verify the resource configuration in the main.cf file.

```
CFSMount crsbin_mnt (
    Critical = 0
    MountPoint = "/app/crshome"
    BlockDevice = "/dev/vx/dsk/bindg/crsbinvol"
)

CFSMount orabin_mnt (
    Critical = 0
    MountPoint = "/app/oracle/orahome"
    BlockDevice = "/dev/vx/dsk/bindg/orabinvol"
)

CVMVolDg crsorabin_voldg (
    Critical = 0
    CVMDiskGroup = bindg
    CVMVolume = { crsbinvol, orabinvol }
    CVMActivation = sw
)
crsbin_mnt requires crsorabin_voldg
crsbin_mnt requires vxfsckd
orabin_mnt requires crsorabin_voldg
orabin_mnt requires vxfsckd
crsorabin_voldg requires cvm_clus
```

7 Verify that the resources are online on all systems in the cluster.

```
# hares -state crsorabin_voldg
# hares -state crsbin_mnt
# hares -state orabin_mnt
```

Note: At this point, the crsorabin_voldg resource is reported offline, and the underlying volumes are online. Therefore, you need to manually bring the resource online on each node.

To bring the resource online manually:

```
# hares -online crsorabin_voldg -sys galaxy
# hares -online crsorabin_voldg -sys nebula
```

Setting up user equivalence

You must establish 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.

Editing the Oracle user profile

Edit the Oracle user `.profile` file to set the paths to ORACLE_BASE, CRS_HOME, and ORACLE_HOME on each node.

In the following sample procedure, `ksh` is the shell environment and the Oracle user home directory is `/home/oracle`.

To edit the Oracle user profile

- As the Oracle user, set the proper environment variables on each node.

```
$ export ORACLE_BASE=/app/oracle
$ export ORACLE_HOME=/app/oracle/orahome
$ export CRS_HOME=/app/crshome
$ export PATH=$PATH:$CRS_HOME/bin:$ORACLE_HOME/bin
$ export CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME\
jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
```

- Check whether the profile file contains commands, such as "stty erase" or other commands that modify the screen display.

Note: The presence of output-producing commands in the profile file garble the screen display when the installer runs.

If the file contains commands that display information on screen, wrap them in a conditional test loop as follows:

```
if test -t 1; then
    uname -a
    stty erase
    echo -uname -n
fi
```

- Verify the profile changes:

```
$ . /home/oracle/.profile
```

Adding private IP addresses to the /etc/hosts file manually

Edit the `/etc/hosts` file to add the private IP address and corresponding node name for each node in the cluster.

To edit the /etc/hosts file

- Log in to each system as the root user.
- Add the following entries to the `/etc/hosts` file:

```
192.168.12.1  galaxy-priv
```

```
192.168.12.2  nebula-priv
```

Installing Oracle RAC

This chapter includes the following topics:

- [Installing Oracle Clusterware using the Oracle Universal Installer](#)
- [Installing the Oracle RAC database using the Oracle Universal Installer](#)
- [Verifying the Oracle Clusterware and database installation](#)
- [Completing the post-installation tasks](#)

Installing Oracle Clusterware using the Oracle Universal Installer

This section provides instructions for installing Oracle Clusterware using the Oracle Universal Installer. The software is installed on each node in the Oracle Clusterware home directory.

To install Oracle Clusterware using the Oracle Universal Installer

1 Log in as the Oracle user. On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

2 Start the Oracle Universal Installer on the first node.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3 On the **Specify Inventory Details and Credentials** screen of the Oracle Universal Installer, enter the following information.

Full path of the inventory directory /app/oracle/oraInventory

Operating system group name oinstall

- 4 On the **Specify Home Details** screen, enter the following information.

Name CRS_HOME

Path /app/crshome

The Oracle Universal Installer performs product-specific prerequisite checks and verifies that your environment meets all of the minimum requirements for installing Oracle RAC. You must manually verify and confirm the items that are flagged with warnings and items that require manual checks.

The OUI displays the full path of the oraInventory logs. Make a note of the log file path to verify the installation at a later time.

- 5 The Oracle Universal Installer displays the cluster and the nodes to be managed by Oracle Clusterware. Verify the displayed information.

Note: The `galaxy-priv` and `nebula-priv` private node names are used for Oracle Clusterware heartbeat.

- 6 The Oracle Universal Installer displays the node's network interfaces. Identify the planned use for each interface: Public, Private, or Do Not use.

The interfaces that are **Private** are stored in OCR as a 'cluster_interconnect' for database cache fusion traffic. Oracle recommends the use of a common private interface for both Oracle Clusterware and Oracle RAC database.

- 7 Enter the full path of the location where you want to store the OCR information.

For example, enter: `/dev/vx/rdsk/ocrvotedg/ocrvol`

Note: Select the option **External Redundancy**. OCR mirroring is performed by CVM.

- 8 Enter the full path of the location where you want to store the voting disk information.

For example, enter: /dev/vx/rdsk/ocrvotedg/votevol.

Note: Select the option **External Redundancy**. Voting disk redundancy is provided by CVM.

- 9 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware installation.
- 10 Run the orainstRoot.sh script as prompted by the Oracle Universal Installer.
- 11 Run the root.sh script on each node as prompted by the Oracle Universal Installer:

```
# cd $CRS_HOME
# ./root.sh
```

The CRS daemons are started on the node where you enter the command.

- 12 Click **OK** to exit the installer.

Note: For Oracle RAC 10g Release 2: If vipca fails to run silently, run the script manually on one of the nodes as the root user.

```
# export DISPLAY=10.20.12.150:0.0
# cd $CRS_HOME/bin
# ./vipca
```

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 the Oracle RAC database using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the DISPLAY variable.
 - For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

2 Start the Oracle Universal Installer.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

3 Enter the following information when prompted by the Oracle Universal Installer:

Select installation type

Select **Enterprise Edition**.

Specify home details

Review or enter the ORACLE_HOME and ORACLE_BASE directory paths.

Specify Hardware Cluster

Select **Cluster Installation**.

Installation Mode

Select the nodes on which you want to install the Oracle RAC database software.

The Oracle Universal Installer runs a product-specific prerequisite check. Any items that are flagged must be manually checked and configured.

4 On the Select Configuration Option screen, select the option **Install database software only.**

Note: Do not select the option **Create a database**. Symantec recommends that you create the database later

5 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle database installation.

6 Run the root.sh script as prompted by the Oracle Universal Installer.

Verifying the Oracle Clusterware and database installation

The following procedure verifies the Oracle Clusterware 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:

```
# $CRS_HOME/bin/crs_stat
```

Verify in the command output that the Oracle Clusterware processes are online on the nodes:

Name	Type	Target	State	Host
ora.galaxy.vip	application	ONLINE	ONLINE	galaxy
ora.galaxy.gsd	application	ONLINE	ONLINE	galaxy
ora.galaxy.ons	application	ONLINE	ONLINE	galaxy
ora.nebula.vip	application	ONLINE	ONLINE	nebula
ora.nebula.gsd	application	ONLINE	ONLINE	nebula
ora.nebula.ons	application	ONLINE	ONLINE	nebula

To verify the Oracle RAC database installation, check the oraInventory logs at `/app/oracle/oraInventory/logs/` (`/app/oracle` is the ORACLE_BASE directory).

Completing the post-installation tasks

Perform the following tasks after installing Oracle RAC:

- [Relinking with ODM](#)
- [Creating Oracle databases](#)
- [Increasing the peer inactivity timeout of LLT](#)
- [Setting the start and stop init sequence for VCS and Oracle Clusterware](#)
- [Configuring LLT to use bonded network interfaces \(optional\)](#)

Relinking with ODM

After installing Oracle database, you must relink Oracle database with Veritas Extension for Oracle Disk Manager (ODM).

If ORACLE_HOME is on a shared file system, run the following commands from any node, otherwise run them on each node.

ORACLE_HOME is the location where Oracle database binaries have been installed.

To configure Veritas Extension for Oracle Disk Manager

- 1 Log in as `oracle` user.
- 2 If the Oracle database is running, then shut down the Oracle database.
- 3 Verify that the file `/opt/VRTSodm/lib64/libodm.so` exists.
- 4 Change to the `$ORACLE_HOME/lib` directory:

```
# cd $ORACLE_HOME/lib
```

- 5 Back up the Oracle ODM file.

For Oracle RAC 10g:

```
# mv libodm10.so libodm10.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
```

For Oracle RAC 11g:

```
# mv libodm11.so libodm11.so.oracle-`date '+%m_%d_%y-%H_%M_%S'`
```

- 6 Copy the Veritas ODM library to the `$ORACLE_HOME/lib` directory:

For Oracle RAC 10g:

```
# cp /usr/lib64/libodm.so libodm10.so
```

For Oracle RAC 11g:

```
# cp /usr/lib64/libodm.so libodm11.so
```

- 7 Start the Oracle database.

- 8 To confirm that the Oracle database starts with Veritas Extension for ODM, check the alert log for the following text:

Veritas <version> ODM Library

where <version> is the ODM library version shipped with the product.

The alert log location depends on the Oracle version used.

For more information on the exact location of the alert log, see the Oracle documentation.

Creating Oracle databases

This section provides instructions for creating Oracle RAC 10g and Oracle RAC 11g database tablespaces. You can create the database on shared raw VxVM volumes, on CFS, or on ASM.

Before you create the database, make sure that CRS daemons are running.

To verify the status of Oracle Clusterware, enter:

```
# $CRS_HOME/bin/crsctl check crs
```

The following text displays a sample output that verifies the status of CRS daemons:

```
Cluster Synchronization Services appears healthy
Cluster Ready Services appears healthy
Event Manager appears healthy
```

Creating database tablespaces on shared raw VxVM volumes

This section provides instructions for creating database tablespaces on shared raw VxVM volumes.

To create database tablespace on shared raw VxVM volumes

- 1 Log in as the root user.
- 2 On any node in the cluster, enter the following command to locate the CVM master:

```
# vxdctl -c mode

mode: enabled: cluster active - MASTER
master: galaxy
```

The above sample output indicates that `galaxy` is the CVM master.

- 3 On the CVM master, identify the spare disks that can be used for creating shared disk group for Oracle database tablespaces:

```
# vxdisk -o alldgs list

DEVICE TYPE DISK GROUP STATUS
sda auto:none -- online invalid
sdb auto:none -- online invalid
sdc auto:cdsdisk - tempdg online shared
sdd auto:none - ocrvotedg online shared
sde auto:cdsdisk -- online shared
sdf auto:cdsdisk -- online shared
```

The above sample output indicates that the shared disks `sde` and `sdf` are free and may be used for Oracle database tablespaces.

Check if the disks are of sufficient size. If the size is not sufficient for the available disks, then you may need to add additional disks to the system.

For more information on size requirements, see the Oracle documentation.

- 4 On the CVM master node, create a shared disk group:

```
# vxdg -s init oradatadg sde sdf
```

- 5 Create a volume in the shared disk group for each of the required tablespaces.

For example, enter:

```
# vxassist -g oradatadg make VRT_volume01 1000M
# vxassist -g oradatadg make VRT_volume02 10M
.
.
.
```

For more information, see the Oracle documentation specific to the Oracle database release to determine the tablespace requirements.

- 6 Define the access mode and permissions for the volumes that store Oracle data. For each volume required for Oracle database tablespaces, run the `vxedit` command as follows:

```
# vxedit -g disk_group set group=group user=user mode=660 \
<volume_name>
```

For example:

```
# vxedit -g oradatadg set group=oinstall user=oracle mode=660 \
VRT_volume01
```

In this example, `VRT_volume01` is the name of one of the volumes.

Repeat the command to define access mode and permissions for each volume in the `oradatadg`.

For more information about the command, see the `vxedit` (1M) manual page.

To automatically start the shared disk group by VCS, you need to configure the shared disk group under VCS.

For more information on configuring a volume and file system under VCS:

See “[Sample main.cf file for configuring a volume and file system under VCS](#)” on page 301.

- 7 Create the database using the Oracle documentation.

Creating database tablespaces on CFS

If you plan to use CFS to store the Oracle database, use the following procedure to create the file system.

To create database tablespaces on CFS

- 1 Log in as the root user.
- 2 On any node in the cluster, enter the following command to locate the CVM master:

```
# vxdctl -c mode

mode: enabled: cluster active - MASTER
master: galaxy
```

The above sample output indicates that `galaxy` is the CVM master.

- 3 On the CVM master, identify the spare disks that can be used for creating shared disk group for Oracle database tablespaces:

```
# vxdisk -o alldgs list
```

```
DEVICE TYPE DISK GROUP STATUS
sda auto:none -- online invalid
sdb auto:none -- online invalid
sdc auto:cdsdisk - tempdg online shared
sdd auto:none - ocrvotedg online shared
sde auto:cdsdisk -- online shared
sdf auto:cdsdisk -- online shared
```

The above sample output indicates that shared disks `sde` and `sdf` are free and can be used for Oracle database tablespaces.

Check if the disks are of sufficient size. If the size is not sufficient for the available disks, then you may need to add additional disks to the system.

For more information on size requirements, see the Oracle documentation.

- 4 Create a shared disk group. For example, enter:

```
# vxdg -s init oradatadg sde sdf
```

- 5 Create a single shared volume that is large enough to contain a file system for all tablespaces.

The following command assumes 6.8 GB of space for the tablespaces:

```
# vxassist -g oradatadg make oradatavol 6800M
```

For more information about tablespace sizes, see the Oracle documentation specific to the Oracle database release.

- 6 Create a VxFS file system in this volume:

```
# mkfs -t vxfs /dev/vx/rdsk/oradatadg/oradatavol
```

- 7 Create a mount point for the shared file system:

```
# mkdir /oradata
```

- 8 From the same node, mount the file system:

```
# mount -t vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol \
/oradata
```

To automatically start the file system by VCS, you need to configure the file system under VCS.

For more information on configuring a volume and file system under VCS:

See “[Sample main.cf file for configuring a volume and file system under VCS](#)” on page 301.

- 9 Set `oracle` as the owner of the file system, and set `775` as the permission:

```
# chown oracle:oinstall /oradata
# chmod 775 /oradata
```

- 10 On the other nodes, complete steps 7 through 8.

- 11 Create the Oracle database using the Oracle documentation.

Increasing the peer inactivity timeout of LLT

SFCFS RAC does not support Symantec’s implementation of I/O fencing. Oracle Clusterware must handle any split-brain situations. In the presence of two clusterwares (VCS and Oracle Clusterware), there is a high possibility of data corruption due to the lack of co-ordination between the clusterwares.

Note: To prevent data corruption, you must modify the LLT peer inactivity timeout settings.

For instructions, refer to the following technote:

<http://entsupport.symantec.com/docs/306411>

Setting the start and stop init sequence for VCS and Oracle Clusterware

VCS and Oracle Clusterware are interdependent services in SFCFS RAC. The mounts and volumes on which Oracle Clusterware resides may be controlled by VCS. Moreover, the OCR and Vote disks used by Oracle Clusterware are configured under VCS. This implies that VCS must start before Oracle Clusterware. Likewise, the volumes and mount points must not be in use by Oracle Clusterware when VCS attempts to take them offline. Therefore, Oracle Clusterware must stop before VCS. Since there is no inherent coordination between VCS and Oracle Clusterware

in SFCFS RAC, you need to ensure that VCS and Oracle Clusterware are started and stopped in the correct order by modifying the numbering of the start and stop scripts for these services in the appropriate run levels.

To start VCS before Oracle Clusterware, modify the numbering such that the VCS start script (S*vcs) ranks lower in number to the Oracle Clusterware start script (S*init.crs) in the appropriate run levels.

Note: If the sequence for the start script is not set correctly, Oracle Clusterware fails to start as the binaries may not be available when the Oracle Clusterware start script is invoked.

To stop Oracle Clusterware before VCS, modify the numbering such that the Oracle Clusterware stop script (K*init.crs) ranks lower in number to the VCS stop script (K*vcs) in the appropriate run levels.

Note: If the sequence for the stop script is not set correctly, the 'shutdown -r now' command hangs indefinitely in VCS as a result of shared volumes and mount points in use by Oracle Clusterware.

Configuring LLT to use bonded network interfaces (optional)

This is an optional task that may be performed after installation.

If you have configured LLT to use a single bonded network interface, GAB reports jeopardy membership even if there is more than one interface beneath the bonded interface.

To prevent GAB from reporting jeopardy membership, it is recommended that you add the following line in the /etc/llttab file:

```
set-dbg-minlinks 2
```

After you update the /etc/llttab file, when LLT is restarted, GAB does not report jeopardy membership even if only one bonded interface is specified in the /etc/llttab file.

For more information, see the following technote:

<http://entsupport.symantec.com/docs/308107>

For more information, see the following documents:

Veritas Volume Manager Administrator's Guide

Veritas Storage Foundation for Cluster File System Administrator's Guide

Veritas Storage Foundation Installation Guide

Upgrading Oracle RAC and migrating the database

This chapter includes the following topics:

- [Supported upgrade paths](#)
- [Upgrading the Oracle database](#)

Supported upgrade paths

[Table 16-1](#) lists the upgrade paths for Oracle RAC.

Table 16-1 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 1
Oracle RAC 10g Release 1	Oracle RAC 10g Release 2
	Oracle RAC 11g Release 1
Oracle RAC 10g Release 2	Oracle RAC 11g Release 1

Note: When you upgrade to a different version of Oracle RAC, make sure that the full path of the Oracle Clusterware 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:

- SFCFS RAC 5.1 SP1
- A supported version of the operating system

Upgrading the Oracle database

For instructions on upgrading the Oracle database, see the appropriate Oracle Metalink documents on the Oracle Web site.

To upgrade Oracle database See Oracle Metalink Doc ID 316889.1
to 10.2.0.4

To upgrade Oracle database See Oracle Metalink Doc ID 429825.1
to 11.1.0.6

Note: After upgrading the database, relink the Oracle library with the Veritas ODM library.

See “[Relinking with ODM](#)” on page 195.

5

Section

Adding or removing nodes from an SFCFS RAC cluster

- [Chapter 17. Adding a node to SFCFS RAC clusters using Oracle RAC](#)
- [Chapter 18. Removing a node from SFCFS RAC clusters using Oracle RAC](#)

Adding a node to SFCFS RAC clusters using Oracle RAC

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Preparing to add a node to a cluster](#)
- [Adding a node to a cluster](#)

About adding a node to a cluster

After you install SFCFS RAC and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 16 nodes.

You can add a node:

- Using the product installer
- Manually

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing SFCFS RAC cluster, verify the following:

- Hardware and software requirements are met.
See “[Meeting hardware and software requirements](#)” on page 208.
- Hardware is set up for the new node.
See “[Setting up the hardware](#)” on page 208.
- The existing cluster is an SFCFS RAC cluster and that SFCFS RAC is running on the cluster.
- The new system has the same identical operating system versions and patch levels as that of the existing cluster.

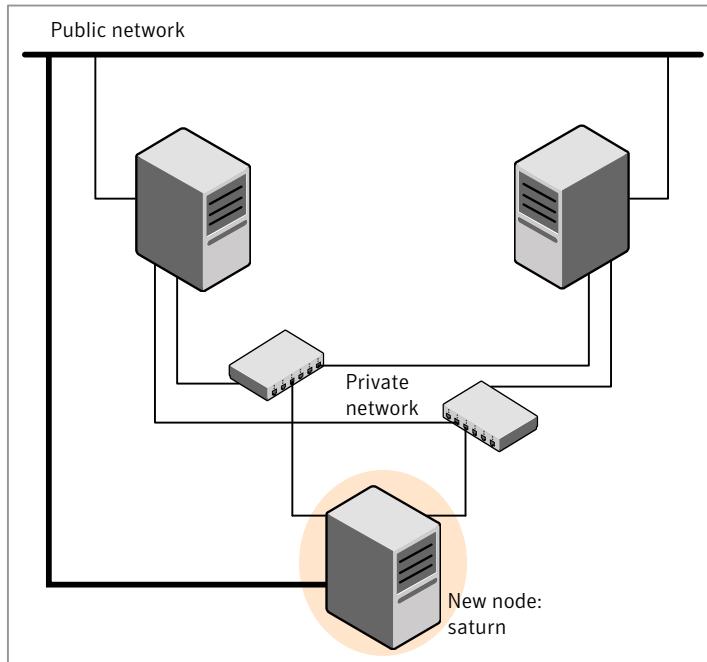
Meeting hardware and software requirements

The system you add to the cluster must meet the hardware and software requirements.

Setting up the hardware

[Figure 17-1](#) shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 17-1 Adding a node to a two-node cluster using two switches



To set up the hardware

1 Connect the SFCFS RAC private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

[Figure 17-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.
- The node must have private network connections to two independent switches for the cluster.
For more information, see the *Veritas Cluster Server Installation Guide*.
- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Preparing to add a node to a cluster

Complete the following preparatory steps on the new node before you add the node to an existing SFCFS RAC cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
# ./installsfcfsrac -precheck saturn
```

- 2 Install SFCFS RAC on the new system.

Note: Use the `-install` option to install SFCFS RAC. Do not configure SFCFS RAC after the installation.

```
Would you like to configure SFCFS RAC on saturn [y, n, q] (n)
```

You can configure the new node later using the configuration from the existing cluster nodes.

See “[About installing and configuring SFCFS RAC](#)” on page 71.

Adding a node to a cluster

You can use one of the following methods to add a node to an existing SFCFS RAC cluster:

SFCFS RAC installer	See “ Adding a node to a cluster using the SFCFS RAC installer ” on page 210.
	See “ Adding a node using the Web-based installer ” on page 213.
Manual	See “ Adding the node to a cluster manually ” on page 214.

Note: Before you add the node, make sure that SFCFS RAC is not configured on the node.

Adding a node to a cluster using the SFCFS RAC installer

You can add a node using the `-addnode` option with the SFCFS RAC installer.

The SFCFS RAC installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.

- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
`/etc/lltab`
`/etc/VRTSvcs/conf/sysname`
- Updates the following configuration files and copies them on the new node:
`/etc/llthosts`
`/etc/gabtab`
`/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node:
`/etc/vx/.uids/clusuuid`
- Configures security on the new node if the existing cluster is a secure cluster.

Warning: If the root broker system has failed, then you must recover or reconfigure the root broker system before you add a new node to the cluster.

- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

- Starts SFCFS RAC processes and configures CVM and CFS on the new node.
- At the end of the process, the new node joins the SFCFS RAC cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

Caution: If you plan to use the SFCFS RAC installer for completing the Oracle pre-installation tasks on the new node, do not quit the installer after adding the node to the cluster. If you quit the installer, you must perform the Oracle pre-installation tasks manually.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SFCFS RAC installer with the -addnode option.

```
# cd /opt/vrts/install  
# ./installsfcfsrac -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 SFCFS RAC cluster. The installer uses the node information to identify the existing cluster.

```
Enter a node of SFCFS RAC cluster to which  
you want to add a node: galaxy
```

- 4 Review and confirm the cluster information.
- 5 Enter the name of the systems that you want to add as new nodes to the cluster.

```
Enter the system names separated by spaces  
to add to the cluster: saturn
```

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

If there are IP addresses already configured on the interface, confirm whether you want to use the interface 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
```

- 7 Enter **y** to configure a second private heartbeat link.

Note: At least two private heartbeat links must be configured for high availability of the cluster.

Would you like to configure a second private heartbeat link? [y,n,q,b,?] (y)

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

If there are IP addresses already configured on the interface, confirm whether you want to use the interface as the second private heartbeat link.

Enter the NIC for the second private heartbeat link on saturn: [b,q,?] **eth2**

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

Enter the NIC for VCS to use on saturn: **eth3**

The installer starts the SFCFS RAC processes and configures CVM and CFS on the new node. The new node is now part of the cluster.

To add the new node into the Oracle cluster and database, see the appropriate Oracle documentation.

- 12 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 and indicates the location of the log file with details of the actions performed.

- 13 Confirm using `lltstat -n` and `gabconfig -a`.

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 In the System Names field enter a name of a node in the cluster where you plan to add the node.
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.
- 3 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 **Validate** button to check if the system can work in the cluster.
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.
- 4 From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.
- 5 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 this procedure after you install SFCFS RAC only if you plan to add the node to the cluster manually.

To add the node manually to the cluster

- 1 Start the Volume Manager.
See "[Starting Volume Manager on the new node](#)" on page 215.
- 2 If the existing cluster is a secure cluster, set up the new node to run in secure mode.
See "[Setting up the node to run in secure mode](#)" on page 217.
- 3 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 220.

4 Start VCS.

See “[To start VCS on the new node](#)” on page 221.

5 Configure CVM and CFS.

See “[Configuring CVM and CFS on the new node](#)” on page 221.

6 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 222.

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 `installsfcfsrac` 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.

Perform the steps in the following procedure to configure LLT and GAB 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. 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 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \
-from_sys galaxy -to_sys saturn
```

- 9 Start the LLT, GAB, and ODM drivers on the new node:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/odm restart
```

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 17-1](#) uses the following information for the following command examples.

Table 17-1 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.
RB1	RB1.brokers.example.com	The root broker for the cluster
RB2	RB2.brokers.example.com	Another root broker, not the cluster's RB

To verify the existing security setup on the node

- 1 If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.
- 2 Find out the root broker to which the node saturn belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.
- 4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.
 - Kill /opt/VRTSat/bin/vxatd process.
 - Remove the credential that RB2 has given to AB on node saturn.

```
# vssat deletecred --domain type:domainname \
--prplname prplname
```

For example:

```
# vssat deletecred --domain vx:root@RB2.brokers.example.com \
--prplname saturn.nodes.example.com
```

Configuring the authentication broker on node saturn

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

To configure the authentication broker on node saturn

- 1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \
--prplname prplname --password password \
--prpltype service
```

For example:

```
# vssat addprpl --pdrtype root \
--domain root@RB1.brokers.example.com \
--prplname saturn.nodes.example.com \
--password flurbdicate --prpltype service
```

- 2 Ensure that there is no clock skew between the times on node saturn and RB1.
- 3 Copy the /opt/VRTSat/bin/root_hash file from RB1 to node saturn.
- 4 Configure AB on node saturn to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -z 2821 -h roothash_file_path
```

For example:

```
# vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \
-x vx -y root@RB1.brokers.example.com -q RB1 \
-z 2821 -h roothash_file_path
```

- 5 Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

Setting up SFCFS RAC related security configuration

Perform the following steps to configure SFCFS RAC related security settings.

Setting up SFCFS RAC related security configuration

1 Start /opt/VRTSat/bin/vxatd process.

2 Create HA_SERVICES domain for SFCFS RAC.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

3 Add SFCFS RAC and webserver principal to AB on node saturn.

```
# vssat addprpl --pdrtype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

4 Create /etc/VRTSvcs/conf/config/.secure file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

To add a node in the VxSS group using the CLI

1 Make a backup copy of the main.cf file. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

3 Add the new node to the VCS configuration:

```
# hasys -add saturn
```

- 4 Add the node saturn to the existing VxSS group.

```
# hagrp -modify VxSS SystemList -add saturn 2
# hagrp -modify VxSS AutoStartList -add saturn
```

- 5 Save the configuration by running the following command from any node in the cluster:

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

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 If you are using disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

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

- 3 On the new node, verify that the GAB port memberships are a, b, and d:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
```

After adding the new node

Start VCS on the new node.

To start VCS on the new node

- 1 Start VCS on the new node:

```
# hastart
```

VCS brings the CVM and CFS groups online.

- 2 Verify that the CVM and CFS groups are online:

```
# hagrp -state
```

Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

To configure CVM and CFS on the new node

- 1 Make a backup copy of the main.cf file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagrp -modify cvm SystemList -add saturn 2  
  
# hagrp -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:

```
# scp /etc/VRTSvcs/conf/config/main.cf \
saturn:/etc/VRTSvcs/conf/config/main.cf
# scp /etc/VRTSvcs/conf/config/CFSTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf
# scp /etc/VRTSvcs/conf/config/CVMTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

- 1** On an existing node, for example galaxy, write-enable the configuration:

```
# haconf -makerw
```

- 2** Add the node saturn to the existing ClusterService group.

```
# hagrp -modify ClusterService SystemList -add saturn 2
# hagrp -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
```

Removing a node from SFCFS RAC clusters using Oracle RAC

This chapter includes the following topics:

- [Removing a node from a cluster](#)

Removing a node from a cluster

Perform the following steps to remove a node from a cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

To remove a node from a cluster

- 1 Log in as superuser on a node other than *saturn*.
- 2 Use the `cfsunmount` command to unmount the file system `/mnt` on all the nodes:

```
# cfsunmount /mnt saturn
```

- 3 Stop all the cluster components:

```
# cfscluster stop -f saturn
```

- 4 Open the VCS configuration for writing:

```
# haconf -makerw
```

- 5 Remove saturn from the CVM and SFCFS RAC service groups:

```
# hagrp -modify service_group SystemList -delete saturn  
# hagrp -modify cvm AutoStartList -delete saturn  
# hagrp -modify cvm SystemList -delete saturn
```

where *service_group* is the name of the service group displayed by the `hagrp -dep cvm` command.

- 6 Write the new VCS configuration to disk:

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

- 7 Edit the `/etc/llthosts` file on the remaining nodes of the cluster, and remove the entry corresponding to the node being removed.

- 8 Edit the `/etc/gabtab` file on the remaining nodes of the cluster and edit the `gabconfig` command to reflect the correct and new number of nodes in the cluster.

- 9 Change to the `install` directory:

```
# cd /opt/VRTS/install
```

- 10 Run the `uninstallsfcfsrac` script and remove SFCFS RAC on saturn:

```
# ./uninstallsfcfsrac
```

If you do not want to remove the Veritas Cluster Server software, enter **n** when prompted to uninstall VCS.

See the *Veritas Cluster Server Installation Guide*.

- 11 Log in to *saturn* and remove the following files:

```
# rm /etc/vxfenmode  
# rm /etc/llthosts  
# rm /etc/llttab  
# rm /etc/gabtab
```

6

Section

Configuration of disaster recovery environments

- [Chapter 19. Setting up a replicated global cluster](#)
- [Chapter 20. Configuring a global cluster using VVR](#)

Setting up a replicated global cluster

This chapter includes the following topics:

- [Replication in the SFCFS RAC environment](#)
- [Requirements for SFCFS RAC global clusters](#)
- [Enabling a keyless license for SFCFS RAC with VVR for a global cluster](#)
- [About setting up a global cluster in an SFCFS RAC environment](#)
- [Configuring a cluster at the primary site](#)
- [Configuring a cluster at the secondary site](#)
- [Configuring replication on clusters at both sites](#)
- [Modifying the ClusterService group for global clusters](#)
- [Defining the remote cluster and heartbeat objects](#)
- [Configuring the VCS service groups for global clusters](#)

Replication in the SFCFS RAC environment

You can set up a primary SFCFS RAC cluster for replication to a secondary SFCFS 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 SFCFS RAC.
- Supported hardware-based replication technologies. Using hardware-based replication you can replicate data from a primary array to a secondary array.

Note: You will need an SFCFSRAC_VVR license and a VCS_GCO license to enable this option.

- Using SFCFS RAC with VVR you can run a fire drill to verify the disaster recovery capability of your configuration.

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

Requirements for SFCFS RAC global clusters

Review the requirements information to make sure your configuration is supported for SFCFS RAC.

For product licensing information:

See “[About Veritas product licensing](#)” on page 47.

Supported software and hardware for SFCFS RAC

For supported hardware and software:

-
- See the current compatibility list in the Veritas Technical Support website to confirm the compatibility of your hardware:
<http://entsupport.symantec.com/docs/283161>

Supported replication technologies for SFCFS RAC

SFCFS RAC supports the following replication technologies through the use of Veritas replication agents:

Table 19-1 Supported replication options for SFCFS RAC global clusters

Replication technology	Supported modes	Supported software
Veritas Volume Replicator (VVR) Supporting agents <ul style="list-style-type: none">■ RVGShared■ RVGSharedPri■ RVGLogOwner	<ul style="list-style-type: none">■ Asynchronous replication■ Synchronous replication	Host-based replication
EMC SRDF Supporting agent: SRDF	<ul style="list-style-type: none">■ Asynchronous replication■ Synchronous replication	All versions of Solutions Enabler
Hitachi True Copy Supporting agent: HTC	<ul style="list-style-type: none">■ 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	<ul style="list-style-type: none">■ Asynchronous replication■ Synchronous replication	SSH access to the SVC
EMC Mirror View Supporting agent: MirrorView	<ul style="list-style-type: none">■ 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.

You can use the Veritas replication agents listed in the table above for global clusters that run SFCFS 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 SFCFS RAC nodes. The agents make sure that the system which has the resource online also has safe and exclusive access to the configured devices.

This information is current at the time this document is released. For more current information on the replicated agents, see:

- *Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM Metro Mirror Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM SVC Installation and Configuration Guide*
- *Veritas Cluster Server Agent for EMC MirrowView Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Oracle Data Guard Installation and Configuration Guide*
- Technical Support TechNote for the latest updates or software issues for replication agents:
<http://entsupport.symantec.com/docs/282004htm>

Enabling a keyless license for SFCFS RAC with VVR for a global cluster

If you install using the installer and select the keyless option, the SFCFS RAC with VVR selection does not include the GCO option required for global clusters.

To enable the GCO option for SFCFS RAC with VVR with a keyless license

- ◆ Do one of the following before configuring a global cluster using VVR:
 - Add a license key for SFCFS for Oracle RAC with VVR and for VCS GCO using the installer.
 - Manually add a keyless license SFCFS for Oracle RAC with VVR and for VCS GCO:

```
vxkeyless set SFCFSRACENT_VVR, VCS_GCO
```

About setting up a global cluster in an SFCFS RAC environment

Configuring a global cluster for Oracle RAC requires the coordination of many component setup tasks. The procedures provided in this document are guidelines.

The tasks required to set up a global cluster:

- Configure an SFCFS RAC cluster at the primary site
- Configure an SFCFS RAC cluster at the secondary site
- Configure replication on clusters at both sites
- Configure VCS service groups for replication
- Test the HA/DR configuration
- Upon successful testing, bring the environment into production

Some SFCFS RAC HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration. Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

For requirements:

Configuring a cluster at the primary site

You can use an existing SFCFS RAC cluster or you can install a new SFCFS RAC cluster for your primary site.

For planning information:

See “[Typical configuration of SFCFS RAC global clusters for disaster recovery](#)” on page 32.

If you are using an existing cluster as the primary and you want to set up a global cluster, skip the steps below and proceed to configure your secondary cluster.

See “[Configuring a cluster at the secondary site](#)” on page 233.

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 SFCFS RAC global cluster, follow the steps below.

To set up the cluster and database at the primary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.
- 4 Install and configure SFCFS RAC. Prepare for your installation according to your configuration needs.

For preparation:

See "[About preparing to install and configure SFCFS RAC](#)" on page 55.

For installation:

See "[About installing and configuring SFCFS RAC](#)" on page 71.

- 5 Verify the CVM group is online on all nodes in the primary cluster:

```
# hagrp -state cvm
```

- 6 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and RDBMS software.

You will need to set up:

- Local storage for Oracle RAC and CRS binaries
- Shared storage for OCR and Vote disk which is not replicated
- Replicated storage for database files

- 7 Install and configure the Oracle RAC binaries:

See "[About preparing to install Oracle RAC](#)" on page 175.

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.
See “[Creating storage for OCR and voting disk manually](#)” on page 177.
- 10 Configure the VCS service groups for the database.
- 11 Verify that all VCS service groups are online.

Configuring a cluster at the secondary site

To set up a multi-node or single-node cluster on the secondary site:

- Set up the cluster
- Set up the database

The setup requirements for the secondary site parallel the requirements for the primary site with a few additions or exceptions as noted below.

Important requirements for global clustering:

- Cluster names on the primary and secondary sites must be unique.
- Make sure that you use the same OS user and group IDs for Oracle for installation and configuration on both the primary and secondary clusters.

Setting up the cluster on the secondary site

To set up the cluster on secondary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.

- 4 Install and configure SFCFS RAC. Prepare for your installation according to your configuration needs.

For preparation:

See "[About preparing to install and configure SFCFS RAC](#)" on page 55.

For installation:

See "[About installing and configuring SFCFS RAC](#)" on page 71.

- 5 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and RDBMS software.

You will need to set up:

- Local storage for Oracle RAC and CRS binaries
- Shared storage for OCR and Vote disk which is not replicated
- Replicated storage for database files

- 6 Install and configure the Oracle RAC binaries:

See "[About preparing to install Oracle RAC](#)" on page 175.

Note: OCR and Vote disk must be on non-replicated shared storage.

After successful Oracle RAC installation and configuration, verify that CRS daemons and resources are up on all nodes.

`$ crs_stat -t`

Setting up the database for the secondary site

To set up the database for the secondary site

- 1 Do not create the database. The database will be replicated from the primary site.
 - If you are using hardware-based replication, the database, disk group, and volumes will be replicated from the primary site.
Create the directory for the CFS mount point which will host the database data and control files.
 - If you are using VVR for replication, create an identical disk group and volumes for the replicated content with the same names and size as listed on the primary site.

Create the directories for the CFS mount points as they are on the primary site. These will be used to host the database and control files when the failover occurs and the secondary is promoted to become the primary site.

- 2 Copy the init\$ORACLE_SID.ora file from \$ORACLE_HOME/dbs at the primary to \$ORACLE_HOME/dbs at the secondary.
- 3 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:

```
$ 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 configure the listener.ora and tnsnames.ora files on the secondary site.
- 5 On the secondary site, register the database using the `srvctl` command as Oracle user.

```
$ srvctl add database -d database_name -o oracle_home -p sp_file
```

To prevent automatic database instance restart, change the Management policy for the database (automatic, manual) to MANUAL using the `srvctl` command:

```
$ srvctl modify database -d database_name -y manual
```

Configuring replication on clusters at both sites

You must configure replication for the database files. Once replication is configured, make sure it is functioning correctly by testing before proceeding.

To configure replication at both sites

- 1 If you are using hardware-based replication, make sure that the replication management software for managing replication is installed on all nodes in both clusters.
- 2 At both sites, identify the disks on which the database resides at the primary site and associate them with the corresponding disks at the secondary site.

For VVR:

See "[Setting up replication using VVR on the primary site](#)" on page 244.

For Hardware-based replication:

See your hardware documentation for details on setting up replication between the two sites.

- 3 Start replication between the sites.

See "[Starting replication of Oracle RAC database volume](#)" on page 251.

Modifying the ClusterService group for global clusters

You have configured VCS service groups for SFCFS RAC on each cluster. Each cluster requires an additional virtual IP address associated with the cluster for cross-cluster communication. The VCS installation and creation of the ClusterService group typically involves defining this IP address.

Configure a global cluster by setting:

- Heartbeat
- Wide area cluster (wac)
- GCO IP (gcoip)
- remote cluster resources

See the *Veritas Cluster Server User's Guide* for complete details on global clustering.

Modifying the global clustering configuration using the wizard

The global clustering wizard completes the following tasks:

- Validates the ability of the current configuration to support a global cluster environment.
- Creates the components that enable the separate clusters, each of which contains a different set of GAB memberships, to connect and operate as a single unit.
- Creates the ClusterService group, or updates an existing ClusterService group.

Run the global clustering configuration wizard on each of the clusters; you must have the global clustering license in place on each node in the cluster.

To modify the ClusterService group for global clusters using the global clustering wizard

- 1 On the primary cluster, start the GCO Configuration wizard:

```
# /opt/VRTSvcs/bin/gcoconfig
```

- 2 The wizard discovers the NIC devices on the local system and prompts you to enter the device to be used for the global cluster. Specify the name of the device and press Enter.
- 3 If you do not have NIC resources in your configuration, the wizard asks you whether the specified NIC will be the public NIC used by all the systems. Enter y if it is the public NIC; otherwise enter n. If you entered n, the wizard prompts you to enter the names of NICs on all systems.

- 4 Enter the virtual IP address for the local cluster.
- 5 If you do not have IP resources in your configuration, the wizard prompts you for the netmask associated with the virtual IP. The wizard detects the netmask; you can accept the suggested value or enter another one.

The wizard starts running commands to create or update the ClusterService group. Various messages indicate the status of these commands. After running these commands, the wizard brings the ClusterService group online.

Defining the remote cluster and heartbeat objects

After configuring global clustering, add the remote cluster object to define the IP address of the cluster on the secondary site, and the heartbeat object to define the cluster-to-cluster heartbeat.

Heartbeats monitor the health of remote clusters. VCS can communicate with the remote cluster only after you set up the heartbeat resource on both clusters.

To define the remote cluster and heartbeat

- 1 On the primary site, enable write access to the configuration:

```
# haconf -makerw
```

- 2 Define the remote cluster and its virtual IP address.

In this example, the remote cluster is rac_cluster102 and its IP address is 10.11.10.102:

```
# haclus -add rac_cluster102 10.11.10.102
```

- 3 Complete step 1 and step 2 on the secondary site using the name and IP address of the primary cluster.

In this example, the primary cluster is rac_cluster101 and its IP address is 10.10.10.101:

```
# haclus -add rac_cluster101 10.10.10.101
```

- 4 On the primary site, add the heartbeat object for the cluster. In this example, the heartbeat method is ICMP ping.

```
# hahb -add Icmp
```

- 5 Define the following attributes for the heartbeat resource:
 - ClusterList lists the remote cluster.
 - Arguments enables you to define the virtual IP address for the remote cluster.

For example:

```
# hahb -modify Icmp ClusterList rac_cluster102
# hahb -modify Icmp Arguments 10.11.10.102 -clus rac_cluster102
```

- 6 Save the configuration and change the access to read-only on the local cluster:

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

- 7 Complete step 4-6 on the secondary site using appropriate values to define the cluster on the primary site and its IP as the remote cluster for the secondary cluster.
- 8 Verify cluster status with the `hastatus -sum` command on both clusters.

```
# hastatus -sum
```

9 Display the global setup by executing haclus -list command.

```
# haclus -list
    rac_cluster101
    rac_cluster102
```

Example of heartbeat additions to the main.cf file on the primary site:

```
.
.
.
remotecluster rac_cluster102 (
    Cluster Address = "10.11.10.102"
)
heartbeat Icmp (
    ClusterList = { rac_cluster102 }
    Arguments @rac_cluster102 = { "10.11.10.102" }
)

system galaxy (
)
.
.
```

Example heartbeat additions to the main.cf file on the secondary site:

```
.
.
.
remotecluster rac_cluster101 (
    Cluster Address = "10.10.10.101"
)
heartbeat Icmp (
    ClusterList = { rac_cluster101 }
    Arguments @rac_cluster101 = { "10.10.10.101" }
)

system mercury (
)
.
.
```

See the *Veritas Cluster Server User's Guide* for details for configuring the required and optional attributes of the heartbeat object.

Configuring the VCS service groups for global clusters

To configure VCS service groups for global clusters

- 1 Configure and enable global groups for databases and resources.
 - Configure VCS service groups at both sites.
 - Configure the replication agent at both sites.
 - For example:
For example:
See “[Migrating the role of primary site to the secondary site](#)” on page 264.
See “[Migrating the role of new primary site back to the original primary site](#)” on page 265.
- 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 264.
See “[Migrating the role of new primary site back to the original primary site](#)” on page 265.
- 3 Upon successful testing, bring the environment into production.
For complete details on VVR in a shared disk environment:
See the *Veritas Volume Replicator Administrator’s Guide*.

Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring global clustering using VVR](#)
- [Setting up replication using VVR on the primary site](#)
- [Setting up replication using VVR on the secondary site](#)
- [Starting replication of Oracle RAC database volume](#)
- [Configuring VCS to replicate the database volume using VVR](#)
- [Using VCS commands on SFCFS RAC global clusters](#)
- [Using VVR commands on SFCFS RAC global clusters](#)

About configuring global clustering using VVR

Before configuring clusters for global clustering, make sure both clusters have product and database software installed and configured.

Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster and VVR enabled.

See “[About Veritas product licensing](#)” on page 47.

See “[Enabling a keyless license for SFCFS RAC with VVR for a global cluster](#)” on page 230.

After setting up two clusters running SFCFS 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 SFCFS 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 SFCFS RAC environment](#)” on page 231.
- Setting up replication for clusters at both sites.
See “[Setting up replication using VVR on the primary site](#)” on page 244.
See “[Setting up replication using VVR on the secondary site](#)” on page 247.
- Starting replication of the database.
See “[Starting replication of Oracle RAC database volume](#)” on page 251.
- Configuring VCS for replication on clusters at both sites.
See “[Configuring VCS to replicate the database volume using VVR](#)” on page 253.

Setting up replication using VVR on the primary site

Setting up replication with VVR in a global cluster environment involves the following tasks:

- If you have not already done so, create a disk group on the storage on the primary site.
- Creating the Storage Replicator Log (SRL) in the disk group for the database.
See “[Creating the SRL volume on the primary site](#)” on page 244.
- Creating the Replicated Volume Group (RVG) on the primary site.
See “[Setting up the Replicated Volume Group \(RVG\) on the primary site](#)” on page 245.

Creating the SRL volume on the primary site

Create the SRL. The SRL is a volume in the RVG. The RVG also holds the data volumes for replication.

- The data volume on the secondary site has the same name and the same size as the data volume on the primary site.
- The SRL on the secondary site has the same name and the same size as the SRL on the primary site.
- The data volume and SRL volume should exist in the same disk group.
- If possible, create SRLs on disks without other volumes.

- Mirror SRLs and data volumes in the absence of hardware-based mirroring.

To create the SRL volume on the primary site

- 1 On the primary site, determine the size of the SRL volume based on the configuration and amount of use.

See the Veritas Volume Replicator documentation for details.

- 2 Using the following command, determine whether a node is the master or the slave:

```
# vxrdctl -c mode
```

- 3 From the master node, issue the following command:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk5
```

- 4 Using the following command, start the SRL volume by starting all volumes in the disk group:

```
# vxvol -g oradatadg startall
```

Setting up the Replicated Volume Group (RVG) on the primary site

Before creating the RVG on the primary site, make sure the volumes and CVM group are active and online.

To review the status of replication objects on the primary site

- 1 Verify the volumes you intend to include in the group are active.
- 2 Review the output of the `hagrp -state cvm` command to verify that the CVM group is online.
- 3 On each site, verify vradmin is running:

```
# ps -ef |grep vradmin
root 536594 598036 0 12:31:25 0 0:00 grep vradmin
```

If vradmin is not running start it:

```
# vxstart_vvr
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmind: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
# ps -ef |grep vradmin
root 536782 1 0 12:32:47 - 0:00 /usr/sbin/vradmind
root 1048622 598036 0 12:32:55 0 0:00 grep vradmin
# netstat -an |grep 4145
tcp4 0 0 *.4145 *.* LISTEN
udp4 0 0 *.4145 *.*
```

To create the RVG

The command to create the primary RVG takes the form:

```
vradmin -g disk_group createpri rvg_name data_volume srl_volume
```

where:

- `disk_group` is the name of the disk group containing the database
- `rvg_name` is the name for the RVG
- `data_volume` is the volume that VVR replicates
- `srl_volume` is the volume for the SRL

For example, to create the `rac1_rvg` RVG, enter:

```
# vradmin -g oradatadg createpri rac1_rvg rac1_vol rac1_srl
```

The command creates the RVG on the primary site and adds a Data Change Map (DCM) for each data volume. In this case, a DCM exists for `rac1_vol`.

Setting up replication using VVR on the secondary site

To create objects for replication on the secondary site, use the `vradmin` command with the `addsec` option. To set up replication on the secondary site, perform the following tasks:

- If you have not already done so, create a disk group to hold data volume, SRL, and RVG on the storage on the secondary site.
- Create volumes for the database and SRL on the secondary site.
See “[Creating the data and SRL volumes on the secondary site](#)” on page 247.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.
See “[Editing the /etc/vx/vras/.rdg files](#)” on page 248.
- 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 248.
- Create the replication objects on the secondary site.
See “[Setting up the disk group on secondary site for replication](#)” on page 249.

Creating the data and SRL volumes on the secondary site

Note the following when creating volumes for the data and SRL:

- The sizes and names of the volumes must reflect the sizes and names of the corresponding volumes in the primary site.
- Create the data and SRL volumes on different disks in the disk group. Use the `vxdisk -g diskgroup list` command to list the disks in the disk group.
- Mirror the volumes.

To create the data and SRL volumes on the secondary site

- 1 In the disk group created for the Oracle RAC database, create a data volume of same size as that in primary for data; in this case, the `rac_voll` volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_voll 6600M nmirror=2 disk1 disk2
```

- 2** Create the volume for the SRL, using the same name and size of the equivalent volume on the primary site. Create the volume on different disks from the disks for the database volume, but on the same disk group that has the data volume:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk6
```

Editing the /etc/vx/vras/.rdg files

Editing the /etc/vx/vras/.rdg file on the secondary site enables VVR to replicate the disk group from the primary site to the secondary site. On each node, VVR uses the /etc/vx/vras/.rdg file to check the authorization to replicate the RVG on the primary site to the secondary site. The file on each node in the secondary site must contain the primary disk group ID, and likewise, the file on each primary system must contain the secondary disk group ID.

To edit the /etc/vx/vras/.rdg files

- 1** On a node in the primary site, display the primary disk group ID:

```
# vxprint -l diskgroup
```

.....

- 2** On each node in the secondary site, edit the /etc/vx/vras/.rdg file and enter the primary disk group ID on a single line.
- 3** On each cluster node of the primary cluster, edit the /etc/vx/vras/.rdg file and enter the secondary disk group ID on a single line.

Setting up IP addresses for RLINKs on each cluster

Creating objects with the vradmin command requires resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.

To set up IP addresses for RLINKS on each cluster

- 1** For each RVG running on each cluster, set up a virtual IP address on one of the nodes of the cluster. These IP addresses are part of the RLINK.

The example assumes for the cluster on the primary site:

- The public 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
- 2 Use the same commands with appropriate values for the interface, IP address, and net mask on the secondary site.
- The example assumes for the secondary site:
- The public network interface is eth0:1
 - virtual IP address is 10.11.9.102
 - net mask is 255.255.255.0
- 3 Define the virtual IP addresses to correspond to a virtual cluster host name on the primary site and a virtual cluster host name on the secondary site.
- Update the /etc/hosts file on all the nodes on both the primary and secondary sites.
- The examples assume:
- rac_cluster101 has IP address 10.10.9.101
 - rac_cluster102 has IP address 10.11.9.102
- 4 Use the ping command to verify the links are functional.

Setting up the disk group on secondary site for replication

Create the replication objects on the secondary site from the master node of the primary site, using the `vradmin` command.

To set up the disk group on the secondary site for replication

- 1 Issue the command in the following format from the cluster on the primary site:

```
# vradmin -g dg_pri addsec rvg_pri pri_host sec_host
```

where:

- `dg_pri` is the disk group on the primary site that VVR will replicate. For example: `rac1_vol`
- `rvg_pri` is the RVG on the primary site. For example: `rac1_rvg`
- `pri_host` is the virtual IP address or resolvable virtual host name of the cluster on the primary site.
For example: `rac_cluster101_1`

- sec_host is the virtual IP address or resolvable virtual host name of the cluster on the secondary site.

For example: rac_cluster102_1

For example, the command to add the cluster on the primary site to the Replicated Data Set (RDS) is:

```
vradmin -g oradatadg addsec rac1_rvg \
rac_cluster101_1
rac_cluster102_1
```

On the secondary site, the above command performs the following tasks:

- Creates an RVG within the specified disk group using the same name as the one for the primary site
- Associates the data and SRL volumes that have the same names as the ones on the primary site with the specified RVG
- Adds a data change map (DCM) for the data volume
- Creates cluster RLINKS for the primary and secondary sites with the default names

2 Verify the list of RVGs in the RDS by executing the following command.

```
# vradmin -g oradatadg -l printrvg
```

For example:

```
Replicated Data Set: rac1_rvg
Primary:
HostName: 10.180.88.187 <localhost>
RvgName: rac1_rvg
DgName: rac1_vol
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_rac_cluster102_1_rac1_rvg, detached=on,
synchronous=off
Secondary:
HostName: 10.190.99.197
RvgName: rac1_rvg
DgName: oradatadg
datavol_cnt: 1
```

```
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_rac_cluster101_1_rac1_rvg, detached=on,
synchronous=off
```

Note: Once the replication is started the value of the detached flag will change the status from OFF to ON.

Starting replication of Oracle RAC database volume

When you have both the primary and secondary sites set up for replication, you can start replication from the primary site to the secondary site.

Start with the default replication settings:

- Mode of replication: synchronous=off
- Latency Protection: latencyprot=off
- SRL overflow protection: srlprot_autodcm
- Packet size: packet_size=8400
- Network protocol: protocol=UDP

Method of initial synchronization:

- Automatic synchronization
- Full synchronization with Checkpoint

For guidelines on modifying these settings and information on choosing the method of replication for the initial synchronization:

See the *Veritas Volume Replicator Administrator's Guide*

Starting replication using automatic synchronization

Use the `vradmin` command to start replication or the transfer of data from the primary site to the secondary site over the network. Because the cluster on the secondary site uses only one host name, the command does not require the `sec_host` argument.

To start replication using automatic synchronization

- ◆ From the primary site, use the following command to automatically synchronize the RVG on the secondary site:

```
vradmin -g disk_group -a startrep pri_rvg sec_host
```

where:

- disk_group is the disk group on the primary site that VVR will replicate
- pri_rvg is the name of the RVG on the primary site
- sec_host is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -a startrep rac1_rvg
    rac_cluster102
```

Starting replication using full synchronization with Checkpoint

Use the `vradmin` command with the Checkpoint option to start replication using full synchronization with Checkpoint.

To start replication using full synchronization with Checkpoint

- 1 From the primary site, synchronize the RVG on the secondary site with full synchronization (using the `-c checkpoint` option):

```
vradmin -g disk_group -full -c ckpt_name syncrvg pri_rvg sec_host
```

where:

- disk_group is the disk group on the primary site that VVR will replicate
- ckpt_name is the name of the checkpoint on the primary site
- pri_rvg is the name of the RVG on the primary site
- sec_host is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -c rac1_ckpt syncrvg rac1_rvg
    rac_cluster102
```

- 2 To start replication after full synchronization, enter the following command:

```
# vradmin -g oradatadg -c rac1_ckpt startrep rac1_rvg
    rac_cluster102
```

Verifying replication status

Verify that replication is properly functioning.

To verify replication status

- 1 Check the status of VVR replication:

```
# vradmin -g disk_group_name repstatus rvg_name
```

- 2 Review the `flags` output for the status. The output may appear as `connected` and `consistent`. For example:

```
# vxprint -g oradatadg -l rlk_rac_cluster102_1_rac1_rvg
Rlink: rlk_rac_cluster102_1_rac1_rvg
info: timeout=500 packet_size=8400 rid=0.1078
      latency_high_mark=10000 latency_low_mark=9950
      bandwidth_limit=none
state: state=ACTIVE
      synchronous=off latencyprot=off srlprot=autodcm
.
.
.
protocol: UDP/IP
checkpoint: rac1_ckpt
flags: write enabled attached consistent connected
asynchronous
```

Configuring VCS to replicate the database volume using VVR

After configuring both clusters for global clustering and setting up the Oracle RAC database for replication, configure VCS to provide high availability for the database. Specifically, configure VCS agents to control the cluster resources, including the replication resources.

To view the sample main.cf files on your system:

```
# cat /etc/VRTSvcs/conf/config/main.cf
```

About modifying the VCS configuration for replication

The following resources must be configured or modified for replication:

- Log owner group

- RVG group
- CVMVolDg resource
- RVGSharedPri resource
- Oracle RAC database service group

For detailed examples of service group modification:

For more information on service replication resources:

See the *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*.

Log owner group

Create a log owner group including the RVGLogowner resources. The RVGLogowner resources are used by:

- RLINKs for the RVG
- RVGLogowner resource. The RVG and its associated disk group are defined as attributes for the RVGLogowner resource.

The RVG log owner service group has an online local firm dependency on the service group containing the RVG.

The VCS uses the following agents to control the following resources:

- RVGLogowner agent to control the RVGLogowner resource
- RVGShared agent to control the RVGShared resource

RVG group

Create an RVG group that includes the RVGShared resource replication objects. Define the RVGShared resource and CVMVolDg resource together within a parallel service group. The group is defined as parallel because it may be online at the same time on all cluster nodes.

CVMVolDg resource

The CVMVolDg resource does not have volumes specified for the CVMVolume attribute; the volumes are contained in the RVG resource. The CVMVolume attribute for the CVMVolDg resource is empty because all volumes in the RVG are defined by the RVG attribute of the RVGShared resource. The RVG service group has an online local firm dependency on the CVM service group.

For a detailed description of the CVMVolDg agent in this guide:

See “[CVMVolDg agent](#)” on page 325.

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.

Modifying the VCS Configuration on the Primary Site

The following are the procedural highlights required to modify the existing VCS configuration on the primary site:

- Configure two service groups:
 - A log owner group including the RVGLogowner resource.
 - An RVG group including the RVGShared resource replication objects.
- Add the RVGSharedPri resource to the existing Oracle RAC database service group and define this group as a global group by setting the ClusterList and ClusterFailOverPolicy attributes.
- Move the CVMVolDg resource from the existing Oracle RAC database service group to the newly created RVG group.

To modify VCS on the primary site

- 1 Log into one of the nodes on the primary cluster.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while you make changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Review the sample configuration file after the SFCFS 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
    NetworkHosts = "10.10.8.1"
)

RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)
requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

- 5 Add the RVG service group using the appropriate values for your cluster and nodes.

Example `RVGgroup` service group:

```
group RVGgroup (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy,nebula }
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)
CVMVolDg racdata_voldg (
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)
requires group cvm online local firm
racdata_rvg requires racdata_voldg
```

- 6 Modify the Oracle RAC service group using the appropriate values for your cluster and nodes:

- Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute. See the **bolded** attribute in the example that follows.
- Add the ClusterFailOverPolicy cluster attribute. Symantec recommends using the **Manual** value. See the **bolded** attribute in the example.
- Add the `RVGSharedPri` resource to the group configuration.
- Remove the `CVMVolDg` resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
- Specify the service group (**online, local, firm**) to depend on the RVG service group.
- Remove the existing dependency of the Database service group on the `CVM` service group. Remove the line:

```
requires group CVM online local firm
```

- Remove the existing dependency between the `CFSMount` for the database and the `CVMVoldg` for the Oracle RAC database. Remove the line:

Configuring VCS to replicate the database volume using VVR

```
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, nebulia = 1 }
    ClusterList = { rac_cluster101 = 0, rac_cluster102 = 1 }
    Parallel = 1
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { galaxy,nebulia }
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/rac1_vol"
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
```

- 7 Save and close the main.cf file.
- 8 Use the following command to verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file:

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

- 9 Stop and restart VCS.

```
# hastop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hastart
```

Modifying the VCS Configuration on the Secondary Site

The following are highlights of the procedure to modify the existing VCS configuration on the secondary site:

- Add the log owner and RVG service groups.
- Add a service group to manage the Oracle RAC database and the supporting resources.
- Define the replication objects and agents, such that the cluster at the secondary site can function as a companion to the primary cluster.

The following steps are similar to those performed on the primary site.

To modify VCS on the secondary site

- 1 Log into one of the nodes on the secondary site as root.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while making changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Edit the CVM group on the secondary site.

Review the sample configuration file after the SFCFS 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:

Configuring VCS to replicate the database volume using VVR

```
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.
- 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
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri

```

- 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.

```
# hastop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hastart
```

11 Verify that VCS brings all resources online. On one node, enter the following command:

```
# hagrp -display
```

The Oracle RAC, RVG, and CVM groups are online on both nodes of the primary site. The RVGLogOwner group is online on one node of the cluster. If either the RVG group or the RVGLogOwner group is partially online, manually bring the groups online using the `hagrp -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:

```
# hagrp -display
```

The Oracle RAC service group is offline on the secondary site, but the CVM, RVG log owner, and RVG groups are online.

This completes the setup for an SFCFS RAC global cluster using VVR for replication. Symantec recommends testing a global cluster before putting it into production.

Using VCS commands on SFCFS RAC global clusters

For information on the VCS commands for global clusters:

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

Using VVR commands on SFCFS RAC global clusters

If you have two SFCFS RAC clusters configured to use VVR for replication, the following administrative functions are available:

- Migration of the role of the primary site to the remote site
- Takeover of the primary site role by the secondary site

About migration and takeover of the primary site role

Migration is a planned transfer of the role of primary replication host from one cluster to a remote cluster. This transfer enables the application on the remote cluster to actively use the replicated data. The former primary cluster becomes free for maintenance or other activity.

Takeover occurs when an unplanned event (such as a disaster) causes a failure, making it necessary for the applications using the replicated data to be brought online on the remote cluster.

Migrating the role of primary site to the secondary site

After configuring the replication objects within VCS, you can use VCS commands to migrate the role of the cluster on the primary site to the remote cluster. In the procedure below, VCS takes the replicated database service group, *database_grp*, offline on the primary site and brings it online on the secondary site; the secondary site now assumes the role of the primary site.

Note: The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

To migrate the role of primary site to the remote site

- 1 From the primary site, use the following command to take the Oracle service group offline on all nodes.

```
# hagrp -offline database_grp -any
```

Wait for VCS to take all Oracle service groups offline on the primary site.

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the vxrlink -g command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
# vxrlink -g data_disk_group status rlk_rac_cluster101_priv_rac1_rvg
```

Where rlk_rac_cluster101_priv_rac1_rvg is the RLINK.

- 3 On the secondary site, which is now the new primary site, bring the Oracle service group online on all nodes:

```
# hagrp -online database_grp -any
```

Migrating the role of new primary site back to the original primary site

After migrating the role of the primary site to the secondary site, you can use VCS commands to migrate the role of the cluster on the new primary site to the original primary site. In the procedure below, VCS takes the replicated database service group, *database_grp*, offline on the new primary (former secondary) site and brings it online on the original primary site; the original primary site now resumes the role of the primary site.

Note: The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

To migrate the role of new primary site back to the original primary site

- 1 Make sure that all CRS resources are online, and switch back the group *database_grp* to the original primary site.

Issue the following command on the remote site:

```
# hagrp -offline database_grp -any
```

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxrlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
# vxrlink -g data_disk_group status rlk_rac_cluster101_priv_rac1_rvg
```

Where `rlk_rac_cluster101_priv_rac1_rvg` is the RLINK.

- 3 Make sure that *database_grp* is offline on the new primary site. Then, execute the following command on the original primary site to bring the *database_grp* online:

```
# hagrp -online database_grp -any
```

Taking over the primary role by the remote cluster

Takeover occurs when the remote cluster on the secondary site starts the application that uses replicated data. This situation may occur if the secondary site perceives the primary site as dead, or when the primary site becomes inaccessible (perhaps for a known reason). For a detailed description of concepts of taking over the primary role:

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

Before enabling the secondary site to take over the primary role, the administrator on the secondary site must "declare" the type of failure at the remote (primary, in this case) site and designate the failure type using one of the options for the `haclus` command.

Takeover options are:

- [Disaster](#)
- [Outage](#)

- [Disconnect](#)
- [Replica](#)

Disaster

When the cluster on the primary site is inaccessible and appears dead, the administrator declares the failure type as "disaster." For example, fire may destroy a data center, including the primary site and all data in the volumes. After making this declaration, the administrator can bring the service group online on the secondary site, which now has the role as "primary" site.

Outage

When the administrator of a secondary site knows the primary site is inaccessible for a known reason, such as a temporary power outage, the administrator may declare the failure as an "outage." Typically, an administrator expects the primary site to return to its original state.

After the declaration for an outage occurs, the RVGSharedPri agent enables DCM logging while the secondary site maintains the primary replication role. After the original primary site becomes alive and returns to its original state, DCM logging makes it possible to use fast fail back resynchronization when data is resynchronized to the original cluster.

Before attempting to resynchronize the data using the fast fail back option from the current primary site to the original primary site, take the precaution at the original primary site of making a snapshot of the original data. This action provides a valid copy of data at the original primary site for use in the case the current primary site fails before the resynchronization is complete.

See "[Examples for takeover and resynchronization](#)" on page 268.

See "[Replica](#)" on page 268.

Disconnect

When both clusters are functioning properly and the heartbeat link between the clusters fails, a split-brain condition exists. In this case, the administrator can declare the failure as "disconnect," which means no attempt will occur to take over the role of the primary site at the secondary site. This declaration is merely advisory, generating a message in the VCS log indicating the failure results from a network outage rather than a server outage.

Replica

In the rare case where the current primary site becomes inaccessible while data is resynchronized from that site to the original primary site using the fast fail back method, the administrator at the original primary site may resort to using a data snapshot (if it exists) taken before the start of the fast fail back operation. In this case, the failure type is designated as "replica".

Examples for takeover and resynchronization

The examples illustrate the steps required for an outage takeover and resynchronization.

To take over after an outage

- 1 From any node of the secondary site, issue the `haclus` command:

```
# haclus -declare outage -clus rac_cluster101
```

- 2 After declaring the state of the remote cluster, bring the Oracle service group online on the secondary site. For example:

```
# hagrp -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 Volume Replicator Administrator's Guide* for details on RVG snapshots.

- 2 Resynchronize the RVG. From the CVM master node of the current primary site, issue the *hares* command and the *-action* option with the *fbsync* action token to resynchronize the RVGSharedPri resource. For example:

```
# hares -action ora_vvr_shpri fbsync -sys mercury
```

To determine which node is the CVM master node, type:

```
# vxdctl -c mode
```

- 3 Perform one of the following commands, depending on whether the resynchronization of data from the current primary site to the original primary site is successful:
 - If the resynchronization of data is successful, use the *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:

```
# vxrvrg -g data_disk_group snaprestore rac1_rvg
```

Troubleshooting CVM and VVR components of SFCFS RAC

The following topic is useful for troubleshooting the VVR component of SFCFS RAC.

Updating the rlink

If the rlink is not up to date, use the `hares -action` command with the `resync` action token to synchronize the RVG.

The following command example is issued on any node (`galaxy`, in this case) in the primary cluster, specifying the RVGSharedPri resource, `ora_vvr_shpri`:

```
# haresh -action ora_vvr_shpri resync -sys galaxy
```

VCS agents to manage wide-area failover

VCS agents now manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

VCS provides agents for other array-based or application-based solutions. This section covers the replication agents that is bundled with VVR. See the VCS replication agent documentation for more details.

Note: See the Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide for more information about the RVG and RVGPrimary agents.

Note: The RVGSnapshot agent is not supported for SFCFS RAC.

DNS agent

The DNS agent updates the canonical name-mapping in the domain name server after a wide-area failover. See the Veritas Cluster Server Bundled Agents Reference Guide for more information about the agent.

RVG agent

The RVG agent manages the Replicated Volume Group (RVG). Specifically, it brings the RVG online, monitors read-write access to the RVG, and takes the RVG offline. Use this agent when using VVR for replication.

RVGPrimary agent

The RVGPrimary agent attempts to migrate or take over a Secondary to a Primary following an application failover. The agent has no actions associated with the offline and monitor routines.

7

Section

Uninstallation of SFCFS RAC

- [Chapter 21. Preparing to uninstall SFCFS RAC from a cluster](#)
- [Chapter 22. Uninstalling SFCFS RAC from a cluster](#)

Preparing to uninstall SFCFS RAC from a cluster

This chapter includes the following topics:

- [About uninstalling SFCFS RAC from a cluster](#)
- [Options for uninstalling SFCFS RAC](#)
- [Preparing to uninstall SFCFS RAC from a cluster](#)

About uninstalling SFCFS RAC from a cluster

You can uninstall SFCFS RAC using the `uninstallsfcfsrac` program.

Note: After you uninstall SFCFS 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 SFCFS RAC.

Options for uninstalling SFCFS RAC

[Table 21-1](#) lists the available options for uninstalling SFCFS RAC:

Table 21-1 Options for uninstalling SFCFS RAC

Options	Description
SFCFS RAC uninstallation program	Use the <code>uninstallsfcfsrac</code> program to uninstall SFCFS RAC. See “ Preparing to uninstall SFCFS RAC from a cluster ” on page 276.

Table 21-1 Options for uninstalling SFCFS RAC (*continued*)

Options	Description
Response file	<p>Use a response file to automate or perform an unattended uninstallation of SFCFS RAC.</p> <p>See “Uninstalling SFCFS RAC using a response file” on page 281.</p>

Preparing to uninstall SFCFS RAC from a cluster

Perform the steps in the following procedure before you uninstall SFCFS RAC from a cluster.

To prepare to uninstall SFCFS RAC from a cluster

- 1 Log in as the root user on any node in the cluster.
- 2 Verify that the following directories are set in your PATH environment variable in order to execute the necessary commands:

```
/opt/VRTS/bin  
/opt/VRTSvcs/bin
```

- 3 Back up the following configuration files:

```
# mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`  
# mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
```

- 4 On all the nodes, stop the CFS-dependant applications that are not under VCS control using application specific commands.

For example, to stop Oracle Clusterware:

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

- 5 Stop VCS:

```
# hastop -all
```

- 6 Verify that port h is not open:

```
# gabconfig -a
```

- 7 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

- 8 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# umount /filesystem
```

- 9 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g disk_group stopall
```

To verify that no volumes are open:

```
# vxprint -Aht -e v_open
```

Preparing to uninstall SFCFS RAC from a cluster

Uninstalling SFCFS RAC from a cluster

This chapter includes the following topics:

- [Uninstalling SFCFS RAC from a cluster using the script-based installer](#)
- [Uninstalling SFCFS RAC with the Veritas Web-based installer](#)
- [Uninstalling SFCFS RAC using a response file](#)

Uninstalling SFCFS RAC from a cluster using the script-based installer

Perform the steps in the following procedure to uninstall SFCFS RAC from a cluster.

To uninstall SFCFS RAC from a cluster manually

- 1 Log in as the root user on any node in the cluster.
- 2 Navigate to the directory that contains the uninstallation program:

```
# cd /opt/VRTS/install
```

- 3 Start the uninstallation program:

```
# ./uninstallsfcfsrac galaxy nebula
```

The installer stops the SFCFS RAC processes and uninstalls the RPMs.

After the uninstallation completes, the installer displays the location of the log and summary files. You may view the files to confirm the uninstallation.

Uninstalling SFCFS RAC with the Veritas Web-based installer

This section describes how to uninstall with the Veritas Web-based installer.

To uninstall SFCFS RAC

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.

See “[Starting the Veritas Web-based installer](#)” on page 97.

- 4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 5 Select **Storage Foundation Cluster File System for Oracle RAC** from the Product drop-down list, and click **Next**.
- 6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Next** to uninstall SFCFS RAC on the selected system.
- 8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 9 After the installer stops the processes, the installer removes the products from the specified system.

Click **Next**.

- 10 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 11 Click **Finish**.

The Web-based installer prompts you for another task.

Uninstalling SFCFS RAC using a response file

Perform the steps in the following procedure to uninstall SFCFS RAC using a response file.

To uninstall SFCFS 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 105.

Note: You must replace the host names in the response file with that of the systems from which you want to uninstall SFCFS RAC.

For a sample response file:

See “[Sample response file for uninstalling SFCFS RAC](#)” on page 282.

3 Navigate to the directory containing the SFCFS RAC uninstallation program:

```
# cd /opt/vrts/install
```

4 Start the uninstallation:

```
# ./uninstallsfcfsrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

5 Reboot the nodes:

```
# shutdown -r now
```

Response file variables to uninstall SFCFS RAC

[Table 22-1](#) lists the response file variables that you can define to uninstall SFCFS RAC.

Table 22-1 Response file variables specific to uninstalling SFCFS RAC

Variable	List or Scalar	Description
CFG{opt}{uninstall}	Scalar	Uninstalls SFCFS RAC RPMs. (Required)

Table 22-1 Response file variables specific to uninstalling SFCFS RAC
(continued)

Variable	List or Scalar	Description
CFG{systems}	List	List of systems on which the product is to be uninstalled. (Required)
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 ssh 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)

Sample response file for uninstalling SFCFS RAC

The following sample response file uninstalls SFCFS RAC from nodes, galaxy and nebula.

```
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{prod}="SFCFSRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
```

1;

8

Section

Installation reference

- [Appendix A. SFCFS RAC installation RPMs](#)
- [Appendix B. Installation scripts](#)
- [Appendix C. Configuration files](#)
- [Appendix D. Automatic Storage Management](#)
- [Appendix E. High availability agent information](#)

Appendix

SFCFS RAC installation RPMs

This appendix includes the following topics:

- [SFCFS RAC installation RPMs](#)

SFCFS RAC installation RPMs

[Table A-1](#) lists the RPM name and contents for each SFCFS RAC RPM.

Table A-1 List of SFCFS RAC RPMs

RPM	Content	Configuration
VRTSatClient	Required to use Symantec Product Authentication Service.	Minimum
VRTSatServer	Required to use Symantec Product Authentication Service.	Minimum
VRTScps	Required for Coordination Point Server (CPS). Contains the binaries for the Veritas Coordination Point Server. Depends on VRTSvxfen.	Recommended
VRTSamf	Contains the binaries for the Veritas Minimum Asynchronous Monitoring Framework kernel driver functionality for the process and mount based agents.	Minimum
VRTSgab	Contains the binaries for Veritas Cluster Server group membership and atomic broadcast services. Depends on VRTSllt.	Minimum

Table A-1 List of SFCFS RAC RPMs (*continued*)

RPM	Content	Configuration
VRTSllt	Contains the binaries for Veritas Cluster Server low-latency transport.	Minimum
VRTSperl	Contains Perl for Veritas.	Minimum
VRTSspt	Contains the binaries for Veritas Software Support Tools.	Recommended
VRTSvcs	<p>Contains the following components:</p> <ul style="list-style-type: none"> ■ Contains the binaries for Veritas Cluster Server. ■ Contains the binaries for Veritas Cluster Server manual pages. ■ Contains the binaries for Veritas Cluster Server English message catalogs. ■ Contains the binaries for Veritas Cluster Server utilities. These utilities include security services. <p>Depends on VRTSvxifen, VRTSgab, and VRTSllt.</p>	Minimum
VRTSvcsag	<p>Contains the binaries for Veritas Cluster Server bundled agents.</p> <p>Depends on VRTSvcs.</p>	Minimum
VRTSvcsea	<p>Required for VCS with the high availability agent for Oracle.</p> <p>VRTSvcsea contains the binaries for Veritas high availability agents for Oracle.</p>	Minimum
VRTSvllic	Contains the binaries for Symantec License Utilities.	Minimum
VRTSvxifen	<p>Contains the binaries for Veritas I/O fencing.</p> <p>Depends on VRTSgab.</p>	Minimum
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System	Minimum
VRTSfssdk	<p>Veritas File System Software Developer Kit</p> <p>For VxFS APIs, the package contains the public Software Developer Kit (SDK), which includes headers, libraries, and sample code. The SDK is required if some user programs use VxFS APIs.</p>	All
VRTSglm	Veritas Group Lock Manager for Storage Foundation Cluster File System	Minimum

Table A-1 List of SFCFS RAC RPMs (*continued*)

RPM	Content	Configuration
VRTSgms	Veritas Group Messaging Services for Storage Foundation Cluster File System	Minimum
VRTSob	Veritas Enterprise Administrator	Recommended
VRTSodm	ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed for Oracle RAC. Oracle Disk Manager enables Oracle to improve performance and manage system bandwidth.	Minimum
VRTSvxfs	Veritas File System binaries	Minimum
VRTSvxvm	Veritas Volume Manager binaries	Minimum
VRTSaslapm	Required for the support and compatibility of various storage arrays Veritas Array Support Library (ASL) and Minimum Array Policy Module (APM) binaries	Minimum
VRTSlvmconv	Tool for conversion of LVM configuration to Veritas Volume Manager Configuration	All
VRTSsfmh	Veritas Storage Foundation Managed Host Recommended Discovers configuration information on a Storage Foundation managed host. This information is stored on a central database, which is not part of this release. You must download the database separately at: http://www.symantec.com/business/storage-foundation-manager	Recommended
VRTSvcsdr	Veritas Cluster Server Disk Reservation Modules and Utilities	Recommended

Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 provides several installation scripts.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the installer, use the appropriate product installation script.

The following product installation scripts are available:

Veritas Cluster Server (VCS)	<code>installvcs</code>
Veritas Storage Foundation (SF)	<code>installsf</code>
Veritas Storage Foundation and High Availability (SFHA)	<code>installsfha</code>
Veritas Storage Foundation Cluster File System (SFCFS)	<code>installsfcfs</code>
Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC)	<code>installsfcfsrac</code>
Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)	<code>installsfrac</code>
Symantec Product Authentication Service (AT)	<code>installat</code>

Veritas Dynamic Multi-pathing	installldmp
Symantec VirtualStore	installsvs

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
# ./installer -stop
```

To start the processes

- ◆ Use the `-start` option to start the product installation script.

For example, to start the product's processes, enter the following command:

```
# ./installer -start
```

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

Installation program has improved failure handling

The product installer has improved ability to recover from failed installations, as follows:

- A recovery file is created if an installation fails due to a failed network connection. This file enables the install program to resume from the point where the installation failed.

- New options are available to start or stop the Veritas processes without requiring a full installation or configuration.

Installation script options

Table B-1 shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See “[About installation scripts](#)” on page 291.

Table B-1 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.
<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 ssh or rsh configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of ssh or rsh are abruptly terminated.
<code>-configure</code>	Configures the product after installation.

Table B-1 Available command line options (*continued*)

Command Line Option	Function
-copyinstallscripts	<p>Use this option when you manually install products and want to use the installation scripts that are stored on the system to perform product configuration, uninstallation, and licensing tasks without the product media.</p> <p>Use this option to copy the installation scripts to an alternate rootpath when you use it with the <code>-rootpath</code> option.</p> <p>The following examples demonstrate the usage for this option:</p> <ul style="list-style-type: none"> ■ <code>./installer -copyinstallscripts</code> Copies the installation and uninstallation scripts for all products in the release to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code> ■ <code>./installproduct_name -copyinstallscripts</code> Copies the installation and uninstallation scripts for the specified product and any subset products for the product to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code> ■ <code>./installer -copyinstallscripts -rootpath alt_root_path</code> The path <code>alt_root_path</code> can be a directory like <code>/rdisk2</code>. In that case, this command copies installation and uninstallation scripts for all the products in the release to <code>/rdisk2/opt/VRTS/install</code>. CPI perl libraries are copied to <code>/rdisk2/opt/VRTSperl/lib/site_perl/release_name</code>, where the <code>release_name</code> is a string that starts with UXRT and includes the release version with no punctuation.
-fencing	Configures I/O fencing in a running cluster.

Table B-1 Available command line options (*continued*)

Command Line Option	Function
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-ignorepatchreqs</code>	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the prerequisite RPMs or patches are missed on the system.
<code>-install</code>	The <code>-install</code> option is used to install products on systems.
<code>-installallpkgs</code>	Specifies that all 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 <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
<code>-license</code>	Registers or updates product licenses on the specified systems.
<code>-listpatches</code>	The <code>-listpatches</code> option displays product patches in correct installation order.
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-makeresponsefile</code>	The <code>-makeresponsefile</code> generates a response file without doing an actual installation. Text displaying install, uninstall, start, and stop actions are simulations. These actions are not being performed on the system.
<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.

Table B-1 Available command line options (*continued*)

Command Line Option	Function
<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.
<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 patches, file system space, and other system requirements in order to install the product.

Table B-1 Available command line options (*continued*)

Command Line Option	Function
<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>-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>-security</code>	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. You can specify this option with the <code>installvcs</code> , <code>installsfha</code> or <code>installsfefs</code> scripts. For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i> .
<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>-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.

Table B-1 Available command line options (*continued*)

Command Line Option	Function
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-upgrade_kernelpkgs	The -upgrade_kernelpkgs option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version
-upgrade_nonkernelpkgs	The -upgrade_nonkernelpkgs option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."
-version	Checks and reports the installed products and their versions. Identifies the installed and missing 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.

Configuration files

This appendix includes the following topics:

- [About the LLT and GAB configuration files](#)
- [Sample main.cf file for configuring a volume and file system under VCS](#)

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

[Table C-1](#) lists the LLT configuration files and the information that these files contain.

Table C-1 LLT configuration files

File	Description
<code>/etc/sysconfig/llt</code>	<p>This file stores the start and stop environment variables for LLT:</p> <ul style="list-style-type: none">■ <code>LLT_START</code>—Defines the startup behavior for the LLT module after a system reboot. Valid values include:<ul style="list-style-type: none">1—Indicates that LLT is enabled to start up.0—Indicates that LLT is disabled to start up.■ <code>LLT_STOP</code>—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include:<ul style="list-style-type: none">1—Indicates that LLT is enabled to shut down.0—Indicates that LLT is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFCFS RAC configuration.</p>

Table C-1 LLT configuration files (*continued*)

File	Description
/etc/llthosts	<p>The file <code>llthosts</code> is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.</p> <p>For example, the file <code>/etc/llthosts</code> contains the entries that resemble:</p> <pre>0 galaxy 1 nebula</pre>
/etc/llttab	<p>The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the LLT network links that correspond to the specific system.</p> <pre>set-node galaxy set-cluster 2 link eth1 eth1 - ether -- link eth2 eth2 - ether --</pre> <p>For example, the file <code>/etc/llttab</code> contains the entries that resemble:</p> <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> <p>If you use aggregated interfaces, then the file contains the aggregated interface name instead of the <code>eth-MAC_address</code>.</p> <p>The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the <code>link</code> command. These lines identify the two network cards that the LLT protocol uses.</p> <p>If you configured a low priority link under LLT, the file also includes a "link-lowpri" line.</p> <p>Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.</p>

[Table C-2](#) lists the GAB configuration files and the information that these files contain.

Table C-2

GAB configuration files

File	Description
/etc/sysconfig/gab	<p>This file stores the start and stop environment variables for GAB:</p> <ul style="list-style-type: none"> ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFCFS RAC configuration.</p>
/etc/gabtab	<p>After you install SFCFS RAC, the file /etc/gabtab contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file /etc/gabtab contains a line that resembles:</p> <pre>/sbin/gabconfig -c -nN</pre> <p>The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <code>N</code> nodes are ready to form the cluster. Symantec recommends that you set <code>N</code> to be the total number of nodes in the cluster.</p> <p>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.</p>

Sample main.cf file for configuring a volume and file system under VCS

The main.cf file is located in the folder /etc/VRTSvcs/conf/config.

To configure a volume and file system under VCS, update the VCS configuration file, main.cf, given below:

```
// Sample main.cf assumes the following configuration:
// CRS_HOME on local VxFS file system
// ORACLE_HOME on CFS
// OCR and Vote-Disk on CVM
// Oracle database tablespaces on CVM or CFS
```

Sample main.cf file for configuring a volume and file system under VCS

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

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

system galaxy (
)

system nebula (
)

// CRS_HOME on galaxy on local VxFS file system
group crshome_grp_galaxy (
    SystemList = { galaxy = 0 }
    AutoFailOver = 0
    AutoStartList = { galaxy }
)

DiskGroup crshome_voldg_galaxy (
    DiskGroup = crsbindg_galaxy
)

Mount crshome_mnt_galaxy (
    MountPoint = "/oracle/crsbin"
    BlockDevice = "/dev/vx/dsk/crsbindg_galaxy/crsbinvol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-y"
)

Volume crshome_vol_galaxy (
    DiskGroup = crsbindg_galaxy
    Volume = crsbinvol
)

crshome_mnt_galaxy requires crshome_vol_galaxy
crshome_vol_galaxy requires crshome_voldg_galaxy

// CRS_HOME on nebula on local VxFS file system
```

Sample main.cf file for configuring a volume and file system under VCS

```

group crshome_grp_nebula (
    SystemList = { nebula = 0 }
    AutoFailOver = 0
    AutoStartList = { nebula }
)

DiskGroup crshome_voldg_nebula (
    DiskGroup = crsbindg_nebula
)

Mount crshome_mnt_nebula (
    MountPoint = "/oracle/crsbin"
    BlockDevice = "/dev/vx/dsk/crsbindg_nebula/crsbinvol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-y"
)

Volume crshome_vol_nebula (
    DiskGroup = crsbindg_nebula
    Volume = crsbinvol
)

crshome_mnt_nebula requires crshome_vol_nebula
crshome_vol_nebula requires crshome_voldg_nebula

// CVM group for:
//      OCR and Vote-Disk on CVM
//      Oracle database tablespaces on CVM or CFS
group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

// OCR and Vote-disk on CVM
CVMVolDg ocrvvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvol, vdvol }
    CVMActivation = sw
)

```

Sample main.cf file for configuring a volume and file system under VCS

```
// Oracle database tablespaces on CFS (Not to be used for
// Oracle database on CVM)
CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

// Oracle database tablespaces on CFS or CVM
CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group orahome_grp online local firm

oradata_mnt requires oradata_voldg

// ORACLE_HOME group
group orahome_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

// ORACLE_HOME on CFS
CFSMount orabin_mnt (
    Critical = 0
    MountPoint = "/oracle/orabin"
    BlockDevice = "/dev/vx/dsk/orabindg/orabinvol"
)

CVMVolDg orabin_voldg (
    Critical = 0
    CVMDiskGroup = orabindg
    CVMVolume = { orabinvol }
    CVMActivation = sw
)
```

Sample main.cf file for configuring a volume and file system under VCS

```
CFSfsckd vxfsckd (
    )

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 1, nebula = 2 }
    CVMTransport = gab
    CVMTimeout = 200
    )

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

orabin_mnt requires orabin_voldg
orabin_mnt requires vxfsckd
orabin_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```


Automatic Storage Management

This appendix includes the following topics:

- [About ASM in SFCFS RAC environments](#)
- [ASM configuration with SFCFS RAC](#)
- [Configuring ASM in SFCFS RAC environments](#)

About ASM in SFCFS RAC environments

ASM is an integrated storage management solution from Oracle RAC that combines file system and volume management capabilities. It provides storage for data files, control files, online redo logs, and archive log files, and backup files. ASM can be configured with Cluster Volume Manager (CVM) for better performance and availability. CVM mirrored volumes with dynamic multipathing improves data access performance and offers continuous data availability in large heterogeneous SAN environments. You can create CVM disk groups and volumes for use as ASM 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.

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

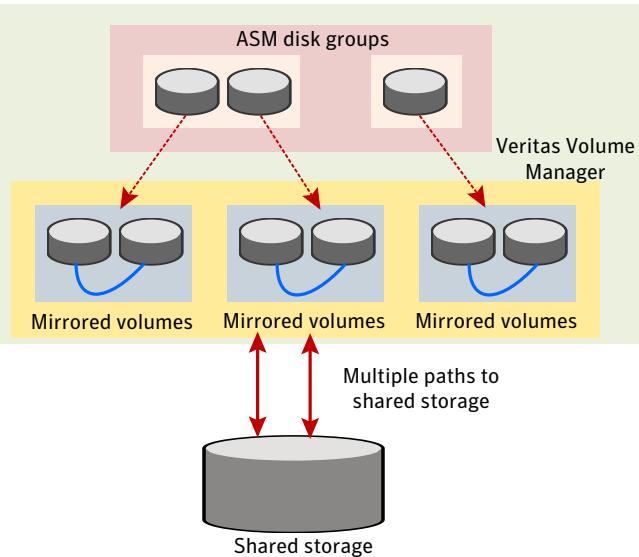
ASM configuration with SFCFS RAC

Configure ASM disk groups over CVM volumes in SFCFS RAC environments. The CVM volumes are mirrored for high availability of data and leverage dynamic multipathing to access the shared storage.

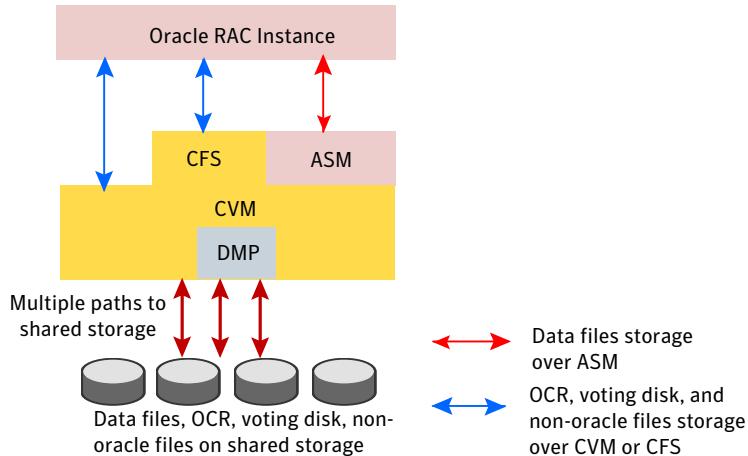
[Figure D-1](#) illustrates the configuration of ASM disk groups over CVM.

Figure D-1

ASM disk groups over CVM



[Figure D-2](#) illustrates ASM in an SFCFS RAC environment.

Figure D-2 Supported Oracle RAC configuration

The figure illustrates the following configuration:

- The ASM disk groups are configured over CVM volumes.
- The Oracle Clusterware and database binaries are stored locally.
- The Oracle database files are stored on ASM configured over CVM. The Oracle databases are managed by Oracle Clusterware.
- The Oracle Cluster Registry and voting disk can be stored on CVM raw volumes or CFS.

Configuring ASM in SFCFS RAC environments

The following procedure provides a summary of the tasks required for configuring ASM in SFCFS RAC installations:

Note: Make sure that you have installed SFCFS RAC and Oracle RAC before installing and configuring ASM.

Before you install, review the planning guidelines for ASM:

See “[Planning for Oracle ASM over CVM](#)” on page 43.

To configure ASM in SFCFS RAC installations

- 1 Set up ASM in a home directory separate from the ORACLE_HOME directory.
See “[Setting up Automatic Storage Management](#)” on page 310.
- 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 311.
- 3 Create ASM disk groups and instances.
See “[Creating ASM disk groups and instances](#)” on page 312.
- 4 Verify the ASM installation.
See “[Verify the ASM installation](#)” on page 314.
- 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 314.

Setting up Automatic Storage Management

Set up ASM in a directory distinct from the Oracle database home (ORACLE_HOME) directory. A separate ASM home directory enables you to upgrade or uninstall the software independent of the Oracle database home.

The following procedure describes how to set up ASM using the Oracle Universal Installer. Symantec recommends that you set up ASM locally on each node.

To set up ASM using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the DISPLAY variable.
 - For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0 export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

- 2 Start the Oracle Universal Installer.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3 Enter the following information when prompted by the Oracle Universal Installer:

Select installation type

Select **Enterprise Edition**.

Specify home details

Enter the full path of the ASM home directory.

Note: Oracle recommends that the ASM home directory be different from the ORACLE_HOME directory used for installing the Oracle database.

Oracle Home location: The installation destination (\$ORACLE_HOME). Depending on where you install Oracle binaries, this location is either on shared storage or an identical location on each of the local cluster nodes.

Specify Hardware Cluster Installation Mode

Select **Cluster Installation**.

Select the nodes on which you want to install the Oracle RAC database software.

- 4 Click **Next**.

The Oracle Universal Installer runs a product-specific prerequisite check. Any items that are flagged must be manually checked and configured.

- 5 On the **Select Configuration Option** screen, select the option **Install database software only**.

Note: Do not select the option **Configure Automatic Storage Management (ASM)**. Symantec recommends that you configure ASM later.

- 6 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the installation.
- 7 Run the root.sh script as prompted by the Oracle Universal Installer.

Creating database storage on ASM

To create the storage for Oracle databases on ASM, you need to 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 node.

To determine the CVM master node:

```
# vxrctl -c mode
```

- 2 Initialize the disks as VxVM disks:

```
# vxdisksetup -i sdx
```

- 3 Create the CVM disk group and volume:

```
# vxrdg -s init ora_asm_dg sdx
```

```
# vxassist -g ora_asm_dg make ora_asm_vol 2000M sdx
```

- 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
```

As Oracle user configure ASM using either DBCA or EM. If you do not use either GUI, you will need to manually create the ASM instances and disk groups.

Creating ASM disk groups and instances

You can create ASM disk groups using one of the following options:

- Database Configuration Assistant (DBCA)
- Oracle Enterprise Manager (OEM)
- 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 +ASMn 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 (asm_diskstring), with the disk group name.

For ASM instances that use an spfile, the parameter is updated automatically by Oracle DBCA.

In either case, dependencies must be managed manually. Symantec recommends the use of the ASMDG agent for easier management of ASM disk groups.

The sample procedure describes the creation of ASM disk groups using Oracle DBCA.

To create ASM disk groups and instances using Oracle DBCA

- 1 Log in as the Oracle user to one of the nodes in the cluster.
- 2 Change to the \$ORACLE_HOME directory for ASM:

```
$ cd $ORACLE_HOME
```

For example, if the ASM home directory is /app/asmhome:

```
$ cd /app/asmhome
```

- 3 Start the Database Configuration Assistant (DBCA) utility.

```
$ dbca
```
- 4 On the Welcome screen, select **Oracle Real Application Clusters database** as the database type and click **Next**.
- 5 Select the option **Configure Automatic Storage Management** and click **Next**.
- 6 Select the nodes for which you want to manage disk groups and click **Next**.
- 7 Enter the ASM SYS password and click **OK**.
Select the **Create server parameter file** option and enter the name of the server parameter file.
- 8 Click **Create New** to create new disk groups.
- 9 On the Create Disk Group screen, click **Change Disk Discovery Path** to select the path that contains the ASM disks.

- 10** On the Change Disk Discovery Path screen, enter the full path that contains the VxVM volumes created for ASM, for example
`/dev/vx/dsk/oradatadg/asmvol.`

The list of disks now display on the Create Disk Group screen.

- 11** On the Create Disk Group screen, enter the following information:

Disk Group Name	Enter a name for the disk group, for example, ASMDG01
Redundancy	Select the option External . Mirroring is performed by CVM.
Select member disks	Select the Show All option to view all 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.
	Click OK .

- 12** Verify the disk group status using the following SQL command:

```
SQL>select name, state from v$asm_diskgroup;
```

Verify the ASM installation

Verify that the database services for ASM are up and running after the installation.

To verify the ASM installation

- ◆ Verify the status of ASM on all the nodes in the cluster:

```
# cd /app/crshome/bin
# ./srvctl status asm
ASM instance +ASM1 is running on node galaxy
ASM instance +ASM1 is running on node nebula
```

The example 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.

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:

```
# hagrp -add dbgrp
```

- 3 Modify the attributes of the service group:

```
# hagrp -modify dbgrp Parallel 1  
  
# hagrp -modify dbgrp SystemList galaxy 0 nebula 1  
  
# hagrp -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 rac_dg  
# hares -modify oradata_voldg CVMActivation sw  
# hares -modify oradata_voldg CVMVolume oradatavol
```

- 6 Add the ASMDG resource for the service group:

```
# hares -add asm_dg 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, which is distinct from the ORACLE_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 oracle
```

- 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 asmdgoradata_voldg
```

- 11 Set the dependencies between the Oracle resource and the ASMDG resource for the Oracle service group:

```
# hares -link asmdbasmdg
```

- 12 Create an online local firm dependency between the dbgrp service group and the cvm service group:

```
# hagrp -link dbgrpcvm online local firm
```

- 13 Enable the Oracle service group:

```
# hagrp -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:

```
# hagrp -online dbgrp -any
```

High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount 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 SFCFS RAC agent are described in this appendix.

VCS agents included within SFCFS RAC

SFCFS RAC includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDg agent
- CFSMount agent

An SFCFS RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

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*

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 E-1](#) describes the entry points used by the CVMCluster agent.

Table E-1 CVMCluster agent entry points

Entry Point	Description
Online	Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by automatically importing the shared disk groups.
Offline	Removes a node from the CVM cluster port.
Monitor	Monitors the node's CVM cluster membership state.

Attribute definition for CVMCluster agent

[Table E-2](#) describes the user-modifiable attributes of the CVMCluster resource type.

Table E-2 CVMCluster agent attributes

Attribute	Dimension	Description
CVMClustName	string-scalar	Name of the cluster.
CVMNodeAddr	string-association	List of host names and IP addresses.
CVMNodeId	string-association	Associative list. The first part names the system; the second part contains the LLT ID number for the system.
CVMTransport	string-scalar	Specifies the cluster messaging mechanism. Default = gab Note: Do not change this value.
PortConfigd	integer-scalar	The port number that is used by CVM for vxconfigd-level communication.
PortKmsgd	integer-scalar	The port number that is used by CVM for kernel-level communication.
CVMTtimeout	integer-scalar	Timeout in seconds used for CVM cluster reconfiguration. Default = 200

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```

type CVMCluster (
    static int InfoTimeout = 0
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName,
        CVMNodeAddr, CVMNodeId, PortConfigd, PortKmsgd,
        CVMTtimeout }
    NameRule = ""
    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
    int PortConfigd
    int PortKmsgd
    int CVMTtimeout
)

```

Note: The attributes `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` are not used in an SFCFS 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
    CVMTtimeout = 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 SFCFS 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 E-3 describes the entry points for the CVMVxconfigd agent.

Table E-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 E-4 describes the modifiable attributes of the CVMVxconfigd resource type.

Table E-4 CVMVxconfigd agent attribute

Attribute	Dimension	Description
CVMVxconfigdArgs	keylist	<p>List of the arguments that are sent to the <code>online</code> entry point.</p> <p>Symantec recommends always specifying the <code>syslog</code> option.</p>

Table E-4 CVMVxconfigd agent attribute (*continued*)

Attribute	Dimension	Description
IMF	integer-association	

Table E-4 CVMVxconfigd agent attribute (*continued*)

Attribute	Dimension	Description
		<p>This resource-type level attribute determines whether the CVMVxconfigd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. <p>Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources <p>Default: 0</p> <ul style="list-style-type: none"> ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> ■ 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 oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the

Table E-4 CVMVxconfigd agent attribute (*continued*)

Attribute	Dimension	Description
		<p>RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.</p> <p>Default: 3.</p> <p>For more details of IMF attribute for the agent type, refer to the <i>Veritas Cluster Server Administrator's Guide</i>.</p>

CVMVxconfigd agent type definition

The following type definition is included in the CVMTypes.cf file:

```
type CVMVxconfigd (
    static int FaultOnMonitorTimeouts = 2
    static int RestartLimit = 5
    static str ArgList[] { CVMVxconfigdArgs }
    static str Operations = OnOnly
    keylist CVMVxconfigdArgs
)
```

CVMVxconfigd agent sample configuration

The following is an example definition for the `CVMVxconfigd` resource in the CVM service group:

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

CVMVolDg agent

The CVMVolDg agent manages the CVM disk groups and CVM volumes and volume sets within the disk groups by performing the following functions:

- Imports the shared disk group from the CVM master node
- Starts the volumes and volume sets in the disk group
- Monitors the disk group, volumes, and volume sets
- Optionally, deports the disk group when the dependent applications are taken offline. The agent deports the disk group only if the appropriate attribute is set.

Configure the CVMVolDg agent for each disk group used by a Oracle service group. A disk group must be configured to only one Oracle service group. If cluster file systems are used for the database, configure the CFSMount agent for each volume or volume set in the disk group.

Entry points for CVMVolDg agent

Table E-5 describes the entry points used by the CVMVolDg agent.

Table E-5 CVMVolDg agent entry points

Entry Point	Description
Online	<p>Imports the shared disk group from the CVM master node, if the disk group is not already imported.</p> <p>Starts all volumes and volume sets in the shared disk group specified by the CVMVolume attribute.</p> <p>Sets the disk group activation mode to shared-write if the value of the CVMActivation attribute is sw. You can set the activation mode on both slave and master systems.</p>
Offline	<p>Removes the temporary files created by the online entry point.</p> <p>If the CVMDepartOnOffline attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.</p>

Table E-5 CVMVolDg agent entry points (*continued*)

Entry Point	Description
Monitor	<p>Determines whether the disk group, the volumes, and the volume sets are online.</p> <p>The agent takes a volume set offline if the file system metadata volume of a volume set is discovered to be offline in a monitor cycle.</p> <p>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 CVMVolDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.</p>
Clean	Removes the temporary files created by the online entry point.

Attribute definition for CVMVolDg agent

[Table E-6](#) describes the user-modifiable attributes of the CVMVolDg resource type.

Table E-6 CVMVolDg agent attributes

Attribute	Dimension	Description
CVMDiskGroup (required)	string-scalar	Shared disk group name.
CVMVolume (required)	string-keylist	Name of shared volumes or volume sets. This list is used to check that the volumes or volume sets are in the correct state before allowing the resource to come online, and that the volumes remain in an enabled state.
CVMActivation (required)	string-scalar	Activation mode for the disk group. Default = sw (shared-write) This is a localized attribute.

Table E-6 CVMVolDg agent attributes (*continued*)

Attribute	Dimension	Description
CVMVolumeIoTest(optional)	string-keylist	<p>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.</p>
CVMDepotOnOffline (optional)	integer-scalar	<p>Indicates whether or not the shared disk group must be deported when the last online CVMVolDg resource for a disk group is taken offline.</p> <p>The value 1 indicates that the agent will deport the shared disk group from the CVM master node, if not already deported, when the last online CVMVolDg resource for the disk group is taken offline.</p> <p>The value 0 indicates that the agent will not deport the shared disk group when the CVMVolDg resource is taken offline.</p> <p>The default value is set to 0.</p> <p>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.</p> <p>The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the CVMDepotOnOffline attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.</p> <p>The deport operation fails if the shared disk group contains open volumes.</p>

CVMVolDg agent type definition

The CVMTypes.cf file includes the CVMVolDg type definition:

```
type CVMVolDg (
    static keylist RegList = { CVMActivation, CVMVolume }
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMDiskGroup, CVMVolume, CVMActivation,
                           CVMVolumeIoTest, CVMDGAction, CVMDeportOnOffline }
    str CVMDiskGroup
    str CVMDGAction
    keylist CVMVolume
    str CVMActivation
    keylist CVMVolumeIoTest
    int CVMDeportOnOffline
    temp int voldg_stat
)
```

CVMVolDg agent sample configuration

Each Oracle service group requires a CVMVolDg resource type to be defined. The following is a sample configuration:

```
CVMVolDg ora_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradata1, oradata2 }
    CVMActivation = sw
)
```

CFSMount agent

The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point.

The agent executable is located in /opt/VRTSvcs/bin/CFSMount/CFSMountAgent.

The CFSMount type definition is described in the /etc/VRTSvcs/conf/config/CFSTypes.cf file.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CFSMount agent

[Table E-7](#) provides the entry points for the CFSMount agent.

Table E-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 E-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

Table E-8 CFSMount Agent attributes

Attribute	Dimension	Description
MountPoint	string-scalar	Directory for the mount point.
BlockDevice	string-scalar	Block device for the mount point.
NodeList	string-keylist	List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list.

Table E-8 CFSMount Agent attributes (*continued*)

Attribute	Dimension	Description
IMF	integer-association	

Table E-8CFSMount Agent attributes (*continued*)

Attribute	Dimension	Description
		<p>Resource-type level attribute that determines whether the CFSMount agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. <p>Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 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 <p>Default: 0</p> <ul style="list-style-type: none"> ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources

Table E-8 CFSMount Agent attributes (*continued*)

Attribute	Dimension	Description
		<ul style="list-style-type: none"> ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the 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.
MountOpt (optional)	string-scalar	<p>Options for the mount command. To create a valid MountOpt attribute string:</p> <ul style="list-style-type: none"> ■ Use the VxFS type-specific options only. ■ Do not use the -o flag to specify the VxFS-specific options. ■ Do not use the -t vxfs 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>
Policy (optional)	string-scalar	List of nodes to assume the primaryship of the cluster file system if the primary fails. If set to NULL or if none of the hosts specified in the list is active when the primary fails, a node is randomly selected from the set of active nodes to assume primaryship.

CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition:

CFSMount agent

```
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;
)
```