

Veritas Storage Foundation™ for Sybase ASE CE Installation and Configuration Guide

Linux

6.0.1

Veritas Storage Foundation™ for Sybase ASE CE Installation and Configuration Guide

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

Product version: 6.0.1

Document version: 6.0.1 Rev 0

Legal Notice

Copyright © 2012 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 [supportolutions@symantec.com](mailto: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 Web site.

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

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Technical Support	4
Section 1 Installation overview and planning	17
Chapter 1 Introducing Veritas Storage Foundation for Sybase ASE CE	19
About Veritas Storage Foundation for Sybase ASE CE	19
Benefits of SF Sybase CE	20
About SF Sybase CE components	21
About SF Sybase CE optional features	22
About VCS notifications	22
About global clusters	22
About Veritas Volume Replicator	23
About I/O fencing	23
About Cluster Manager (Java Console)	24
About Veritas Operations Manager	24
About Symantec Operations Readiness Tools	25
SF Sybase CE cluster setup models	25
Typical configuration of four-node SF Sybase CE cluster	26
Typical configuration of SF Sybase CE clusters in secure mode	27
Typical configuration of VOM-managed SF Sybase CE clusters	28
Typical configuration of SF Sybase CE global clusters for disaster recovery	29
Chapter 2 System requirements	31
Important preinstallation information	31
Hardware requirements	32
Supported operating systems	32
Coordinator disk requirements for I/O fencing	32
Supported database software	33
Supported SF Sybase CE configurations	33
Supported replication technologies for global clusters	33

	Discovering product versions and various requirement information	34
Chapter 3	Planning to install SF Sybase CE	35
	Planning your network configuration	35
	Planning the public network configuration for Sybase ASE CE	35
	Planning the private network configuration for Sybase ASE CE	36
	Planning the storage	36
	Planning the storage for SF Sybase CE	36
	Planning the storage for Sybase ASE CE	37
	Planning volume layout	38
	About planning to configure I/O fencing	38
	Typical SF Sybase CE cluster configuration with disk-based I/O fencing	39
	Planning for cluster management	40
Chapter 4	Licensing SF Sybase CE	43
	About Veritas product licensing	43
	About SF Sybase CE licenses	44
	Setting or changing the product level for keyless licensing	46
	Installing Veritas product license keys	47
Section 2	Preparing to install SF Sybase CE	49
Chapter 5	Preparing to install SF Sybase CE	51
	About preparing to install and configure SF Sybase CE	52
	Synchronizing time settings on cluster nodes	53
	Setting up inter-system communication	53
	Setting up ssh on cluster systems	53
	Mounting the product disc	55
	Setting up shared storage	55
	Setting the environment variables for SF Sybase CE	56
	Setting the kernel.panic tunable	56
	Configuring the I/O scheduler	57
	Configuring the SLES10 network	57
	Optimizing LLT media speed settings on private NICs	57
	Guidelines for setting the media speed of the LLT interconnects	57
	Verifying the systems before installation	58

	About installation and configuration methods	58
	About the Veritas installer	59
Section 3	Installation of SF Sybase CE using the script-based installer	63
Chapter 6	Installing SF Sybase CE	65
	Installing SF Sybase CE using the Veritas script-based installation program	65
Chapter 7	Configuring SF Sybase CE	69
	About configuring SF Sybase CE	69
	Configuring the SF Sybase CE components using the script-based installer	70
	Configuring the SF Sybase CE cluster	72
Chapter 8	Configuring SF Sybase CE clusters for data integrity	89
	Setting up disk-based I/O fencing using installsfsybasece	89
	Initializing disks as VxVM disks	89
	Identifying disks to use as coordinator disks	90
	Checking shared disks for I/O fencing	90
	Configuring disk-based I/O fencing using installsfsybasece	95
Section 4	Post-installation tasks	99
Chapter 9	Verifying the installation	101
	Performing a postcheck on a node	101
	Verifying SF Sybase CE installation using VCS configuration file	101
	Verifying LLT, GAB, and cluster operation	102
	Verifying LLT	102
	Verifying GAB	104
	Verifying the cluster	106
	Verifying the cluster nodes	107

Chapter 10	Performing additional post-installation and configuration tasks	111
	About enabling LDAP authentication for clusters that run in secure mode	111
	Enabling LDAP authentication for clusters that run in secure mode	113
	Configuring Veritas Volume Replicator	118
	Running SORT Data Collector to collect configuration information	119
Section 5	Upgrade of SF Sybase CE	121
Chapter 11	Planning to upgrade SF Sybase CE	123
	About types of upgrade	123
	Supported upgrade paths	124
Chapter 12	Performing a full upgrade of SF Sybase CE using the product installer	125
	About full upgrades	125
	Preparing to perform a full upgrade to SF Sybase CE 6.0.1	126
	Upgrading to SF Sybase CE 6.0.1	128
	Upgrading SF Sybase CE using the Veritas script-based installation program	130
Chapter 13	Performing an automated full upgrade of SF Sybase CE using response files	135
	Upgrading SF Sybase CE using a response file	135
	Response file variables to upgrade Veritas Storage Foundation for Sybase ASE CE	136
	Sample response file for upgrading SF Sybase CE	138
Chapter 14	Performing a phased upgrade of SF Sybase CE	139
	About phased upgrade	139
	Performing phased upgrade of SF Sybase CE from version 5.0 and later releases	140
	Step 1: Performing pre-upgrade tasks on the first half of the cluster	141
	Step 2: Upgrading the first half of the cluster	143

	Step 3: Performing pre-upgrade tasks on the second half of the cluster	145
	Step 4: Performing post-upgrade tasks on the first half of the cluster	147
	Step 5: Upgrading the second half of the cluster	148
	Step 6: Performing post-upgrade tasks on the second half of the cluster	149
Chapter 15	Performing a rolling upgrade of SF Sybase CE	151
	About rolling upgrades	151
	Supported rolling upgrade paths	154
	Preparing to perform a rolling upgrade to SF Sybase CE 6.0.1	154
	Performing a rolling upgrade using the installer	156
	Performing a rolling upgrade using the script-based installer	156
Chapter 16	Performing post-upgrade tasks	159
	Re-joining the backup boot disk group into the current disk group	159
	Reverting to the backup boot disk group after an unsuccessful upgrade	160
	Setting or changing the product license level	160
	Upgrading disk layout versions	161
	Upgrading CVM protocol version and VxVM disk group version	161
	Verifying the cluster	162
Section 6	Installation and upgrade of Sybase ASE CE	165
Chapter 17	Installing, configuring, and upgrading Sybase ASE CE	167
	Before installing Sybase ASE CE	167
	Preparing for local mount point on VxFS for Sybase ASE CE binary installation	168
	Preparing for shared mount point on CFS for Sybase ASE CE binary installation	169
	Installing Sybase ASE CE software	170
	Preparing to create a Sybase ASE CE cluster	170
	Creating the Sybase ASE CE cluster	172
	Preparing to configure the Sybase instances under VCS control	172

	Language settings for the Sybase agent	173
	Configuring Sybase for detail monitoring	173
	Encrypting passwords for Sybase	175
	About setting up detail monitoring for the agentfor Sybase	175
	Configuring a Sybase ASE CE cluster under VCS control using the SF Sybase CE installer	178
	Upgrading Sybase ASE CE	185
Section 7	Automated installation using response files	187
Chapter 18	About reponse files	189
	About response files	189
	Response file syntax	190
	Guidelines for creating the SF Sybase CE response file	191
	Installation scenarios for response files	192
Chapter 19	Installing and configuring SF Sybase CE using a response file	193
	Installing and configuring SF Sybase CE	193
	Response file variables to install Veritas Storage Foundation for Sybase ASE CE	195
	Response file variables to configure Veritas Storage Foundation for Sybase ASE CE	197
	Sample response files for installing and configuring SF Sybase CE	207
Chapter 20	Performing an automated I/O fencing configuration using response files	209
	Configuring I/O fencing using response files	209
	Response file variables to configure disk-based I/O fencing	210
	Sample response file for configuring disk-based I/O fencing	213
Chapter 21	Configuring a cluster under VCS control using a response file	215
	Configuring a Sybase cluster under VCS control with a response file	215
	Response file variables to configure SF Sybase CE in VCS	216

Section 8	Adding and removing nodes	219
Chapter 22	Adding a node to SF Sybase CE clusters	221
	About adding a node to a cluster	221
	Before adding a node to a cluster	222
	Adding the node to a cluster manually	225
	Starting Veritas Volume Manager (VxVM) on the new node	225
	Configuring cluster processes on the new node	226
	Setting up the node to run in secure mode	228
	Starting fencing on the new node	231
	Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node	231
	After adding the new node	233
	Configuring the ClusterService group for the new node	233
	Adding a node to a cluster using the SF Sybase CE installer	234
	Adding the new instance to the Sybase ASE CE cluster	237
	Creating Sybase user and groups	238
	Preparing the mount point for Sybase resources on the new node	238
	Adding a new Sybase ASE CE instance to the Sybase ASE CE cluster	238
	Bringing the new Sybase ASE CE instance under VCS control	239
Chapter 23	Removing a node from SF Sybase CE clusters	241
	About removing a node from a cluster	241
	Removing a node from a cluster	242
	Modifying the VCS configuration files on existing nodes	243
	Removing security credentials from the leaving node	246
Section 9	Configuration of disaster recovery environments	247
Chapter 24	Configuring disaster recovery environments	249
	Disaster recovery options for SF Sybase CE	249
	About setting up a global cluster environment for SF Sybase CE	250
	About configuring a parallel global cluster using Veritas Volume Replicator (VVR) for replication	250

Section 10	Uninstallation of SF Sybase CE	253
Chapter 25	Preparing to uninstall SF Sybase CE from a cluster	255
	About uninstalling SF Sybase CE from a cluster	255
	Options for uninstalling SF Sybase CE	256
	Preparing to uninstall SF Sybase CE from a cluster	257
	Stopping applications that use the Sybase database	258
	Stopping Sybase instances	258
	Backing up the Sybase database	259
	Uninstalling Sybase ASE CE (optional)	259
	Removing root disk encapsulation	260
	Stopping the applications that use CVM or CFS (outside of VCS control)	261
	Unmounting CFS file systems (outside of VCS control)	261
	Stopping VCS	262
	Stopping the applications that use VxVM or VxFS (outside of VCS control)	262
	Unmounting VxFS file systems (outside of VCS control)	263
Chapter 26	Uninstalling SF Sybase CE using the product installer	265
	Uninstalling SF Sybase CE with the script-based installer	265
	Removing the SF Sybase CE RPMs	265
	Removing other configuration files (optional)	267
Chapter 27	Performing an automated uninstallation of SF Sybase CE using response files	269
	Uninstalling SF Sybase CE using a response file	269
	Response file variables to uninstall Veritas Storage Foundation for Sybase ASE CE	270
	Sample response file for uninstalling SF Sybase CE	271
Section 11	Installation reference	273
Appendix A	SF Sybase CE installation RPMs	275
	SF Sybase CE installation RPMs	275

Appendix B	Installation scripts	279
	Installation script options	279
	About using the postcheck option	284
Appendix C	Sample installation and configuration values	287
	SF Sybase CE installation and configuration information	287
	SF Sybase CE worksheet	287
Appendix D	Tunable files for installation	293
	About setting tunable parameters using the installer or a response file	293
	Setting tunables for an installation, configuration, or upgrade	294
	Setting tunables with no other installer-related operations	295
	Setting tunables with an un-integrated response file	296
	Preparing the tunables file	297
	Setting parameters for the tunables file	297
	Tunables value parameter definitions	298
Appendix E	Configuration files	305
	About sample main.cf files	305
	Sample main.cf files for Sybase ASE CE configurations	305
	Sample main.cf for a basic Sybase ASE CE cluster configuration under VCS control with shared mount point on CFS for Sybase binary installation	306
	Sample main.cf for a basic Sybase ASE CE cluster configuration with local mount point on VxFS for Sybase binary installation	310
	Sample main.cf for a primary CVM VVR site	314
	Sample main.cf for a secondary CVM VVR site	321
Appendix F	High availability agent information	329
	About agents	329
	VCS agents included within SF Sybase CE	330
	VCS agent for Sybase included within SF Sybase CE	330
	CVMCluster agent	331
	Entry points for CVMCluster agent	331
	Attribute definition for CVMCluster agent	331
	CVMCluster agent type definition	332
	CVMCluster agent sample configuration	333
	CVMVxconfig agent	333

Entry points for CVMVxconfigd agent	333
Attribute definition for CVMVxconfigd agent	334
CVMVxconfigd agent type definition	335
CVMVxconfigd agent sample configuration	336
CVMVolDg agent	336
Entry points for CVMVolDg agent	337
Attribute definition for CVMVolDg agent	337
CVMVolDg agent type definition	338
CVMVolDg agent sample configuration	339
CFSMount agent	339
Entry points for CFSMount agent	339
Attribute definition for CFSMount agent	340
CFSMount agent type definition	342
CFSMount agent sample configuration	343
Process agent	343
Agent functions	343
State definitions	343
Attributes	344
Resource type definition	345
Sample configurations	346
Monitoring options for the Sybase agent	346
Sybase resource type	346
Type definition for the Sybase agent	347
Attribute definitions for the Sybase agent	347

Appendix G	Compatibility issues when installing Storage Foundation for Sybase ASE CE with other products	357
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	357
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	358
	Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present	358
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	359
Index		361

Installation overview and planning

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

Introducing Veritas Storage Foundation for Sybase ASE CE

This chapter includes the following topics:

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

About Veritas Storage Foundation for Sybase ASE CE

Veritas Storage Foundation™ for Sybase® Adaptive Server Enterprise Cluster Edition (SF Sybase CE) by Symantec leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of Sybase ASE CE on UNIX platforms. The solution uses 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.

SF Sybase CE integrates existing Symantec storage management and clustering technologies into a flexible solution which administrators can use to:

- Create a standard toward application and database management in data centers. SF Sybase CE provides flexible support for many types of applications and databases.
- Set up an infrastructure for Sybase ASE CE that simplifies database management while fully integrating with Sybase clustering solution.
- Apply existing expertise of Symantec technologies toward this product.

The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Benefits of SF Sybase CE

SF Sybase CE provides the following benefits:

- Use of a generic clustered file system (CFS) technology or a local file system (VxFS) technology for storing and managing Sybase ASE CE installation binaries.
- Support for file system-based management. SF Sybase CE provides a generic clustered file system technology for storing and managing Sybase ASE CE data files as well as other application data.
- Use of Cluster File System (CFS) for the Sybase ASE CE quorum device.
- Support for a standardized approach toward application and database management. A single-vendor solution for the complete SF Sybase CE software stack lets you devise a standardized approach toward application and database management. Further, administrators can apply existing expertise of Veritas technologies toward SF Sybase CE.
- Easy administration and monitoring of SF Sybase CE clusters using Veritas Operations Manager.
- Enhanced scalability and availability with access to multiple Sybase ASE CE instances per database in a cluster.
- Prevention of data corruption in split-brain scenarios with robust SCSI-3 Persistent Reservation (PR) based I/O fencing.
- Support for sharing all types of files, in addition to Sybase ASE CE database files, across nodes.
- Increased availability and performance using Veritas Dynamic Multi-Pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the Host Bus Adapters (HBAs) and Storage Area Network (SAN) switches.

- Fast disaster recovery with minimal downtime and interruption to users. Users can transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites. If a node fails, clients that are attached to the failed node can reconnect to a surviving node and resume access to the shared database. Recovery after failure in the SF Sybase CE environment is far quicker than recovery for a failover database.
- Support for block-level replication using VVR.

For more information on the SF Sybase CE components, see the following documents:

Veritas Storage Foundation for Sybase ASE CE Installation and Configuration Guide

Veritas Storage Foundation for Sybase ASE CE Administrator's Guide

About SF Sybase CE components

SF Sybase CE manages database instances running in parallel on multiple nodes using the following architecture and communication mechanisms to provide the infrastructure for Sybase ASE CE.

Table 1-1 SF Sybase CE component products

Component product	Description
Cluster Volume Manager (CVM)	Enables simultaneous access to shared volumes based on technology from Veritas Volume Manager (VxVM).
Cluster File System (CFS)	Enables simultaneous access to shared file systems based on technology from Veritas File System (VxFS).
Cluster Server (VCS)	Uses technology from Veritas Cluster Server to manage Sybase ASE CE databases and infrastructure components.
VXFEN	The VCS module prevents cluster corruption through the use of SCSI3 I/O fencing, where the vxfen mode is set to sybase.
VXFEND	The VXFEN daemon communicates directly with VCMP and relays membership modification messages.
VCMP	VCMP provides interface between Sybase cluster and the SF Sybase CE components.
QRMUTIL	QRMUTIL provides Sybase instance status.

Table 1-1 SF Sybase CE component products (*continued*)

Component product	Description
Sybase agent	The VCS agent is responsible for bringing Sybase ASE online, taking it offline, and monitoring it.. It obtains status by checking for processes, performing SQL queries on a running database, and interacting with QRMUTIL.

See “[About Veritas Storage Foundation for Sybase ASE CE](#) ” on page 19.

About SF Sybase CE optional features

You can configure the following optional features in an SF Sybase CE cluster:

- VCS notifications
See “[About VCS notifications](#)” on page 22.
- Global clusters
See “[About global clusters](#)” on page 22.
- Veritas Volume Replicator
See “[About Veritas Volume Replicator](#)” on page 23.

About VCS notifications

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

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

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

About 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.

You are required to have a separate license to configure global clusters. You may add this license during the installation or at any time after the installation completes.

About Veritas Volume Replicator

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

VVR replicates the application writes on the volumes at the source location to one or more remote locations across any distance. It provides a consistent copy of application data at the remote locations. If a disaster occurs at the source location, you can use the copy of the application data at the remote location and restart the application at the remote location. The host at the source location on which the application is running is known as the Primary host. The host at the target location is known as the Secondary host. You can have up to 32 Secondary hosts in a VVR environment. VVR provides several methods to initialize the application data between the primary location and the remote location. Some of the methods include using the network, using the tape backup, and moving the disks physically.

About I/O fencing

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

The fencing operation determines the following:

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

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

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

See [“About planning to configure I/O fencing”](#) on page 38.

See the *Veritas Storage Foundation for Sybase ASE CE Administrator's Guide*.

About preferred fencing

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

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy for the following:

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

See the *Veritas Storage Foundation for Sybase ASE CE Administrator's Guide* for more details.

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types. You cannot manage the new features of releases 6.0 and later using the Java Console.

See *Veritas Cluster Server Administrator's Guide*.

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

About Veritas Operations Manager

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.

Symantec recommends using Veritas Operations Manager (VOM) to manage Storage Foundation and Cluster Server environments.

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 Symantec Operations Readiness Tools

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

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

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

To access SORT, go to:

<https://sort.symantec.com>

SF Sybase CE cluster setup models

SF Sybase CE supports a variety of cluster configurations.


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

- Basic setup
See “[Typical configuration of four-node SF Sybase CE cluster](#)” on page 26.
- Secure setup
See “[Typical configuration of SF Sybase CE clusters in secure mode](#)” on page 27.
- Central management setup
See “[Typical configuration of VOM-managed SF Sybase CE clusters](#)” on page 28.
- Global cluster setup
See “[Typical configuration of SF Sybase CE global clusters for disaster recovery](#)” on page 29.

Typical configuration of four-node SF Sybase CE cluster

Figure 1-1 depicts a high-level view of a basic SF Sybase CE configuration for a four-node cluster.

Figure 1-1 Sample four-node SF Sybase CE 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, using similar network devices and matching port numbers.

For example, if you use eth1 on one end of a link, it is recommended that you use eth1 on the other end too.

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.
- The quorum and Sybase datafile disks configured on the shared storage that is available to each node.
Disks for Sybase ASE CE binary can be configured either on shared storage or on local storage.
For shared storage:
See [“Preparing for shared mount point on CFS for Sybase ASE CE binary installation”](#) on page 169.
For local storage:
See [“Preparing for local mount point on VxFS for Sybase ASE CE binary installation”](#) on page 168.
- VCS manages the resources that are required by Sybase ASE CE. The resources must run in parallel on each node.

Typical configuration of SF Sybase CE clusters in secure mode


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

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


For information about how to configure secure clusters, see your product installation guide.

Figure 1-2 Typical configuration of SF Sybase CE clusters in secure mode

Multiple clusters



Single cluster




Typical configuration of VOM-managed SF Sybase CE clusters

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

See [“About Veritas Operations Manager”](#) on page 24.

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

Figure 1-3 Typical configuration of VOM-managed clusters




Typical configuration of SF Sybase CE global clusters for disaster recovery

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

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

Figure 1-4 Global clusters



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

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

For information on supported replication technologies:

See [“Supported replication technologies for global clusters”](#) on page 33.

System requirements

This chapter includes the following topics:

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

Important preinstallation information

Before you install SF Sybase CE, make sure you have reviewed the following information:

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

Hardware requirements

Table 2-1 lists the hardware requirements for SF Sybase CE.

Table 2-1 Hardware requirements for basic clusters

Item	Description
SF Sybase CE systems	Two to four systems with two or more CPUs. For details on the additional requirements for Sybase, see the Sybase documentation.
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: # <code>./installsfbasece -precheck node_name</code> For details on the additional space that is required for Sybase, see the Sybase documentation.
RAM	Each SF Sybase CE system requires at least 2 GB.
Network	Two or more private links and one public link. Links must be 100BaseT or gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication. These links must be of the same type; you cannot mix 100BaseT and gigabit. Symantec recommends gigabit Ethernet using enterprise-class switches for the private links.
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 operating systems

For information on supported operating systems, see the *Veritas Storage Foundation for Sybase ASE CE Release Notes*.

Coordinator disk requirements for I/O fencing

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

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

Supported database software

SF Sybase CE supports Sybase ASE CE 15.5 only at time of publication.

For the latest information on the supported Sybase ASE CE database versions, see the following Technical Support TechNote:

<http://www.symantec.com/docs/DOC4848>

See the Sybase ASE CE documentation for more information.

Supported SF Sybase CE configurations

The following Sybase configuration options are required in an SF Sybase CE environment:

- Set SF Sybase CE fencing to "sybase" mode.
- Configure Sybase private networks on LLT links
- Set Sybase cluster membership to "vcs" mode.
- Configure Sybase instances under VCS control.

Supported replication technologies for global clusters

SF Sybase CE supports the software replication technology Veritas Volume Replicator (VVR) for global cluster configurations.

Discovering product versions and various requirement information

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

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

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

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

To run the version checker

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

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

Planning to install SF Sybase CE

This chapter includes the following topics:

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

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 Sybase ASE CE

Public interconnects are used by the clients to connect to Sybase ASE CE database. The public networks must be physically separated from the private networks.

See Sybase ASE CE documentation for more information on recommendations for public network configurations.

Planning the private network configuration for Sybase ASE CE

Private interconnect is an essential component of a shared disk cluster installation. It is a physical connection that allows inter-node communication. Symantec recommends that these interconnects and LLT links must be the same. You must have the IP addresses configured on these interconnects, persistent after reboot. You must use solutions specific to the operating System.

See Sybase ASE CE documentation for more information on recommendations for private network configurations.

Planning the storage

The following recommendations ensure better performance and availability of storage.

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

Planning the storage for SF Sybase CE

[Table 3-1](#) lists the type of storage required for SF Sybase CE.

Table 3-1 Type of storage required for SF Sybase CE

SF Sybase CE files	Type of storage
SF Sybase CE binaries	Local
SF Sybase CE fencing coordinator disks	Shared

Planning the storage for Sybase ASE CE

Review the storage options and guidelines for Sybase ASE CE:

- Storage options for Sybase ASE CE binaries
 See “[Planning the storage for Sybase ASE CE binaries](#)” on page 37.
- Storage options for Sybase ASE CE datafiles and quorum device
 See “[Planning the storage for Sybase ASE CE datafiles and quorum device](#)” on page 37.

Note: Symantec strongly recommends retaining the default setting (global) for the CVM diskgroup disk detach policy, for Sybase ASE CE binaries, datafiles, and quorum device. For other disk detach policy options, see the Veritas Storage Foundation Administrator’s Guide.

Planning the storage for Sybase ASE CE binaries

The Sybase ASE CE binaries can be stored on a local or shared storage, depending on your high availability requirements.

Note: Symantec recommends that you install the Sybase ASE CE binaries on a shared storage on CFS.

Consider the following points while planning the installation:

- CFS installation provides a single Sybase ASE CE installation, regardless of the number of nodes. This scenario offers a reduction in the storage requirements and easy addition of nodes.
- Local installation provides improved protection against a single point of failure and also allows for applying the Sybase ASE CE patches in a rolling fashion.

Planning the storage for Sybase ASE CE datafiles and quorum device

Storage for Sybase ASE CE datafiles and quorum device has to be configured on shared storage on CFS.

Table 3-2 Type of storage required for Sybase ASE CE

Sybase ASE CE files	Type of storage
Sybase ASE CE binaries	Shared or local
Sybase ASE CE datafiles	Shared

Table 3-2 Type of storage required for Sybase ASE CE (*continued*)

Sybase ASE CE files	Type of storage
Quorum device	Shared

Note: Refer to the Sybase ASE CE documentation for Sybase's recommendation on the required disk space for Sybase ASE CE binaries, Sybase ASE CE datafiles and quorum device.

Planning volume layout

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

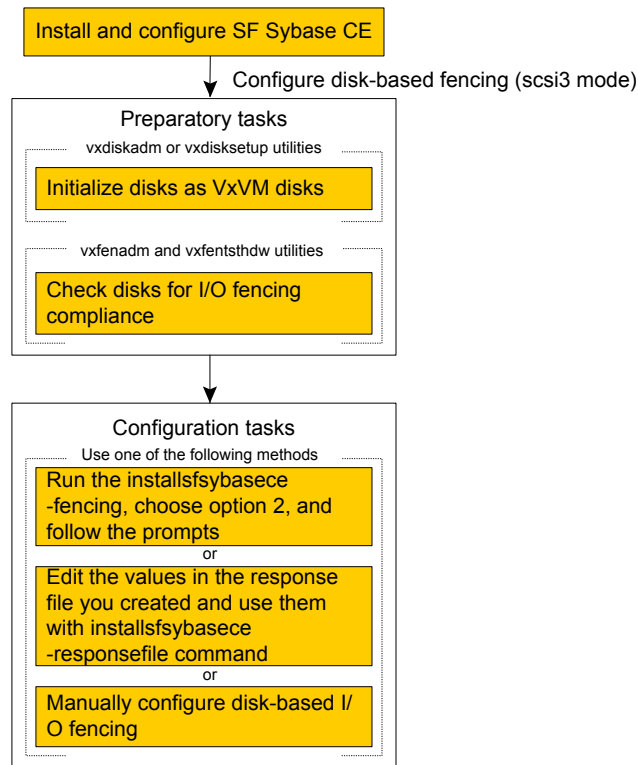
- Mirror the volumes across two or more storage arrays, if using VxVM mirrors. Keep the Fast Mirror Resync regionsize equal to the database block size to reduce the copy-on-write (COW) overhead. Reducing the regionsize increases the amount of Cache Object allocations leading to performance overheads.
- Distribute the I/O load uniformly on all Cache Objects when you create multiple Cache Objects.
- Implement zoning on SAN switch to control access to shared storage. Be aware that physical disks may be shared by multiple servers or applications and must therefore be protected from accidental access.
- Choose DMP I/O policy based on the storage network topology and the application I/O pattern.
- Exploit thin provisioning for better return on investment.

About planning to configure I/O fencing

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

[Figure 3-1](#) illustrates a high-level flowchart to configure I/O fencing for the SF Sybase CE cluster.

Figure 3-1 Workflow to configure I/O fencing



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

Using the
installsfbasece

See [“Setting up disk-based I/O fencing using installsfbasece”](#) on page 89.

Using response files

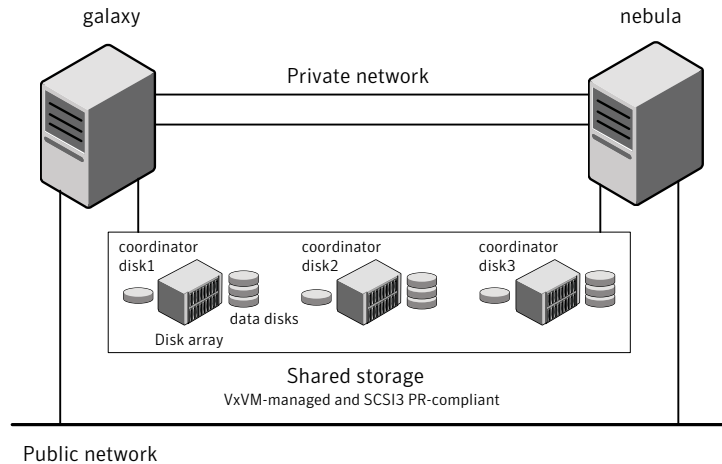
See [“Response file variables to configure disk-based I/O fencing”](#) on page 210.
 See [“Configuring I/O fencing using response files”](#) on page 209.

Manually editing
configuration files

Typical SF Sybase CE cluster configuration with disk-based I/O fencing

[Figure 3-2](#) displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.

Figure 3-2 Typical SF Sybase CE cluster configuration with disk-based I/O fencing



Planning for cluster management

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

Table 3-3 List of agents

Agent	Description
VCS agent for Sybase	<p>Sybase database management</p> <p>The VCS Sybase agent is recommended for managing Sybase databases. VCS controls the Sybase database in this configuration. In the basic monitoring mode, the agent detects an application failure if a configured Sybase server process is not running.</p>
VCS agents for CVM	<p>Volume management</p> <p>An SF Sybase CE installation automatically configures the CVMCluster resource and the CVMVxconfigd resource. You must configure the CVMVolDg agent for each shared disk group.</p>
VCS agents for CFS	<p>File system management</p> <p>If the database uses cluster file systems, configure the CFMount agent for each volume in the disk group.</p>

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

Agent	Description
VCS process agent for vxfsend	vxfsend process/daemon management

Licensing SF Sybase CE

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. If you do not comply with the above terms, continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

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

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

If you upgrade to this release from a prior release of the Veritas software, the installer asks whether you want to upgrade the key to the new version. 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 46.
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 47.
See the `vxlicinst (1m)` manual page.

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

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

About SF Sybase CE licenses

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

Table 4-1 SF Sybase CE license levels (keyless licensing)

License	Description	Features enabled
SFSYBASECE	SF Sybase CE Enterprise Edition	The license enables the following features: <ul style="list-style-type: none"> ■ Veritas Volume Manager ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services
SFSYBASECE_VR	SF Sybase CE Enterprise Edition with VR (Veritas Replicator)	The license enables the following features: <ul style="list-style-type: none"> ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services
SFSYBASECE_VFR	SF Sybase CE Enterprise Edition with VFR (Veritas File Replicator)	The license enables the following features: <ul style="list-style-type: none"> ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas File System ■ Veritas Cluster Server ■ Veritas Mapping Services
SFSYBASECE_GCO	SF Sybase CE Enterprise Edition with GCO (Global Cluster Option)	The license enables the following features: <ul style="list-style-type: none"> ■ Veritas Volume Manager ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services
SFSYBASECE_VR_GCO	SF Sybase CE Enterprise Edition with VR and GCO	The license enables the following features: <ul style="list-style-type: none"> ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services

Table 4-1 SF Sybase CE license levels (keyless licensing) *(continued)*

License	Description	Features enabled
SFSYBASECE_VFR_GCO	SF Sybase CE Enterprise Edition with VFR and GCO	<p>The license enables the following features:</p> <ul style="list-style-type: none"> ■ Veritas Volume Manager Veritas Volume Replicator is enabled. ■ Veritas File System ■ Veritas Cluster Server Global Cluster Option is enabled. ■ Veritas Mapping Services

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.

For more information to use keyless licensing 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 Change your current working directory:

```
# cd /opt/VRTSvlic/bin
```

- 2 View the current setting for the product level.

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

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

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

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

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas product license keys

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

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

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

To install a new license

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

```
# cd /opt/VRTS/bin  
  
# ./vxlicinst -k license key
```

To see a list of your vxkeyless keys, enter the following command:

```
# ./vxkeyless display
```

After you upgrade from a previous release, the output you see when you run the `vxkeyless display` command includes the previous release's vxkeyless keys. Each vxkeyless key name includes the suffix `_<previous_release_version>`. For example, `DMP_6.0`, or `SFENT_VR_5.1SP1`, or `VCS_GCO_5.1`. During the upgrade process, the CPI installer prompts you to update the vxkeyless keys to the current release level. If you update the vxkeyless keys during the upgrade process, you no longer see the `_<previous_release_number>` suffix after the keys are updated.

Preparing to install SF Sybase CE

- [Chapter 5. Preparing to install SF Sybase CE](#)

Preparing to install SF Sybase CE

This chapter includes the following topics:

- [About preparing to install and configure SF Sybase CE](#)
- [Synchronizing time settings on cluster nodes](#)
- [Setting up inter-system communication](#)
- [Mounting the product disc](#)
- [Setting up shared storage](#)
- [Setting the environment variables for SF Sybase CE](#)
- [Setting the kernel.panic tunable](#)
- [Configuring the I/O scheduler](#)
- [Configuring the SLES10 network](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Verifying the systems before installation](#)
- [About installation and configuration methods](#)
- [About the Veritas installer](#)

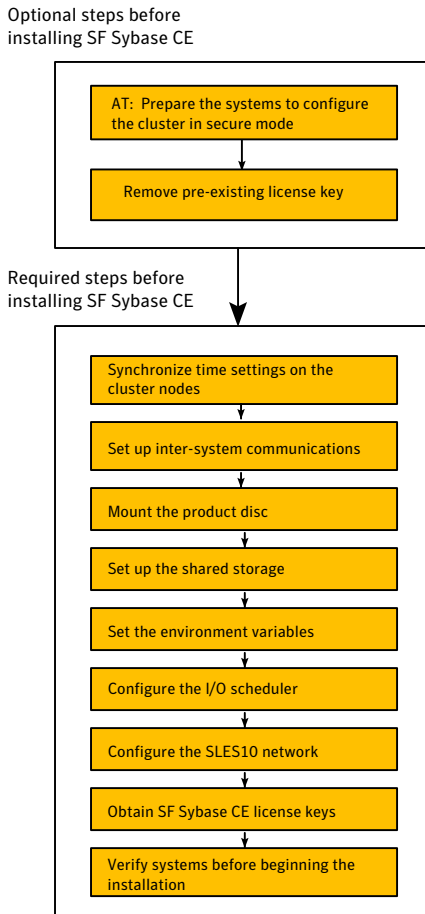
About preparing to install and configure SF Sybase CE

Before you install and configure SF Sybase CE, you need to perform some preinstallation tasks for the required and optional components of SF Sybase CE.

If you do not want to configure the optional components and features, proceed directly to the mandatory pre-installation tasks:

Figure 5-1 illustrates an overview of the mandatory and optional pre-installation steps for SF Sybase CE. The optional tasks are performed only for optional components or features that you plan to use.

Figure 5-1 SF Sybase CE pre-installation tasks



Synchronizing time settings on cluster nodes

Symantec recommends that the time settings on all cluster nodes be synchronized by running the Network Time Protocol (NTP) daemon.

The installer provides the option for automatic NTP synchronization.

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

- On SLES10, use the `yast` command:

```
# yast ntp-client
```

Setting up inter-system communication

When you install SF Sybase CE using the `installsfybasece`, make sure that communication between systems exists. By default the installer uses `ssh`. You must have root privileges for the system where you run `installsfybasece`. This privilege facilitates to issue `ssh` or `rsh` commands on all systems in the cluster. If `ssh` is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases. Similarly, `rsh` must be configured in such a way to not prompt for passwords.

If system communication is not possible between systems using `ssh` or `rsh`, you have recourse.

Warning: The `rsh` and `ssh` commands to the remote systems, where SF Sybase CE is to be installed, must not print any extraneous characters.

Setting up ssh on cluster systems

Use the Secure Shell (`ssh`) to install SF Sybase CE on all systems in a cluster from a system outside of the cluster. Before you start the installation process, verify that `ssh` is configured correctly.

Use Secure Shell (`ssh`) to do the following:

- Log on to another system over a network
- Execute commands on a remote system
- Copy files from one system to another

The `ssh` shell provides strong authentication and secure communications over channels. It is intended to replace `rlogin`, `rsh`, and `rcp`.

The Remote Shell (rsh) is disabled by default to provide better security. Use ssh for remote command execution.

Configuring ssh

The procedure to configure ssh uses OpenSSH example file names and commands.

Note: You can configure ssh in other ways. Regardless of how ssh is configured, complete the last step in the example to verify the configuration.

To configure ssh

- 1 Log on to the system from which you want to install SF Sybase CE.
- 2 Generate a DSA key pair on this system by running the following command:

```
# ssh-keygen -t dsa
```

- 3 Accept the default location of ~/.ssh/id_dsa.
- 4 When the command prompts, enter a passphrase and confirm it.
- 5 Change the permissions of the .ssh directory by typing:

```
# chmod 755 ~/.ssh
```

- 6 The file ~/.ssh/id_dsa.pub contains a line that begins with ssh_dss and ends with the name of the system on which it was created. Copy this line to the /root/.ssh/authorized_keys2 file on all systems where you plan to install SF Sybase CE.

If the local system is part of the cluster, make sure to edit the authorized_keys2 file on that system.

- 7 Run the following commands on the system where you are installing:

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

This step is shell-specific and is valid for the duration the shell is alive.

- 8 When the command prompts, enter your DSA passphrase.

To avoid running the ssh-agent on each shell, run the X-Window system and configure it so that you are not prompted for the passphrase. Refer to the documentation for your distribution for more information.

- 9 To verify that you can connect to the systems where you plan to install SF Sybase CE, type:

```
# ssh -x -l root sys1 ls
# ssh -x -l root sys2 ifconfig
```

The commands should execute on the remote system without having to enter a passphrase or password.

Mounting the product disc

You must have superuser (root) privileges to load the SF Sybase CE software.

You can unmount the product disc after completing the SF Sybase CE installation and configuration.

To mount the product disc

- 1 Log in as the superuser to a cluster node or a remote node in the same subnet as the cluster nodes.
- 2 Insert the product disc with the SF Sybase CE 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/cdrom /mnt/cdrom
```

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 for SF Sybase CE

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
```

Setting the kernel.panic tunable

By default, the kernel.panic tunable is set to zero. Therefore the kernel does not reboot automatically if a node panics. To ensure that the node reboots automatically after it panics, this tunable must be set to a non zero value.

To set the kernel.panic tunable

- 1 Set the kernel.panic tunable to a desired value in the /etc/sysctl.conf file.

For example, kernel.panic = 10, will assign a value 10 seconds to the kernel.panic tunable. This step makes the change persistent across reboots.

- 2 Run the command:

```
sysctl -w kernel.panic=10
```

In case of a panic, the node will reboot after 10 seconds.

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.

To determine whether a system uses the deadline scheduler

- ◆ Look for "elevator=deadline" in `/proc/cmdline`.

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

Configuring the SLES10 network

Before you install SF Sybase CE on SLES10, some network configuration is required on SLES10.

See the SLES10-specific documentation for configuring persistent network interface names, so that a specific network interface name is always assigned to a network hardware address associated to a specific network port.

See the following document for details:

```
#/usr/share/doc/packages/sysconfig/README.Persistent_Interface_Names
```

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 or the operating system manual for more information.

Verifying the systems before installation

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

- Option 1: Run Symantec Operations Readiness Tools (SORT).
For information on downloading and running SORT:
<https://sort.symantec.com>
- Option 2: Run the `installsfsybasece` with the `"-precheck"` option as follows:
Navigate to the directory that contains the `installsfsybasece` program.
The program is located in the `storage_foundation_for_sybase_ce` directory.
Start the preinstallation check:

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

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

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

About installation and configuration methods

You can use one of the following methods to install and configure SF Sybase CE.

SF Sybase CE does not support installation and configuration using the Web installer.

Table 5-1 Installation and configuration methods

Method	Description
<p>Interactive installation and configuration using the script-based installer</p> <p>Note: If you obtained SF Sybase CE from an electronic download site, you must use the <code>installsfbasece</code> script instead of the <code>installer</code> script.</p>	<p>You can use one of the following script-based installers:</p> <ul style="list-style-type: none"> ■ Common product installer script: <code>installer</code> The common product installer script provides a menu that simplifies the selection of installation and configuration options. ■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsfbasece</code> script is identical to specifying SF Sybase CE from the <code>installer</code> script. Use this method to install or configure only SF Sybase CE.
<p>Silent installation using the response file</p>	<p>The response file automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. You can use the script-based installers with the response file to install silently on one or more systems.</p>

About the Veritas installer

To install your Veritas product, use one of the following methods:

- The general product installer. The general product installer enables you to install and configure the product, verify preinstallation requirements, and view the product’s description. You perform the installation from a disc, and you are prompted to choose a product to install.
See [“Installing SF Sybase CE using the Veritas script-based installation program”](#) on page 65.
- Product-specific installation scripts. If you obtained a standalone Veritas product from an electronic download site, the single product download files do not contain the general product installer. Use the product installation script to install the individual products. You can find these scripts at the root of the product media in the scripts directory. These scripts are also installed with the product.

[Table 5-2](#) lists all the SFHA Solutions product installation scripts. The list of product installation scripts that you find on your system depends on the product that you install on your system.

Note: The name of the script is different depending on whether you run the script from the install media or from a system on which the product software is installed.

Table 5-2 Product installation scripts

Veritas product name	Product installation script (When running the script from the install media)	Product installation script (When running the script from a system on which the SFHA Solutions product is installed)
Veritas Cluster Server (VCS)	<code>installvcs</code>	<code>installvcs<version></code>
Veritas Storage Foundation (SF)	<code>installsf</code>	<code>installsf<version></code>
Veritas Storage Foundation and High Availability (SFHA)	<code>installsfha</code>	<code>installsfha<version></code>
Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)	<code>installsfcfsha</code>	<code>installsfcfsha<version></code>
Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)	<code>installsfrac</code>	<code>installsfrac<version></code>
Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)	<code>installsfsybasece</code>	<code>installsfsybasece<version></code>
Veritas Dynamic Multi-Pathing	<code>installdmp</code>	<code>installdmp<version></code>
Symantec VirtualStore	<code>installsvs</code>	<code>installsvs<version></code>

The scripts that are installed on the system include the product version in the script name. For example, to install the SF Sybase CE script from the install media, run the `installsfsybasece` command. However, to run the script from the installed binaries, run the `installsfsybasece<version>` command.

For example, for the 6.0.1 version:

```
# /opt/VRTS/install/installsfsybasece601 -configure
```

Note: Do not include the release version if you use the general product installer to install the product.

At most points during the installation you can type the following characters for different actions:

- Use `b` (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions.
- Use `Control+c` to stop and exit the program if an installation procedure hangs. After a short delay, the script exits.
- Use `q` to quit the installer.
- Use `?` to display help information.
- Use the Enter button to accept a default response.

See [“Installation script options”](#) on page 279.

Installation of SF Sybase CE using the script-based installer

- [Chapter 6. Installing SF Sybase CE](#)
- [Chapter 7. Configuring SF Sybase CE](#)
- [Chapter 8. Configuring SF Sybase CE clusters for data integrity](#)

Installing SF Sybase CE

This chapter includes the following topics:

- [Installing SF Sybase CE using the Veritas script-based installation program](#)

Installing SF Sybase CE 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, swap space, and available volume space.

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

- Installs the SF Sybase CE 6.0.1 RPMs.

The following sample procedure installs SF Sybase CE on two systems—`sys1` and `sys2`.

To install SF Sybase CE

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

SF Sybase CE Run the program:

installer

```
# ./installsfbasece sys1 sys2
```

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

```
# ./installer sys1 sys2
```

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

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.

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

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

Enter the superuser password for system sys1:

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

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

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 installation of SF Sybase CE on sys1.

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

5 Enter y to agree to the End User License Agreement (EULA).

6 Select the type of installation—Minimal, Recommended, All. Each option displays the disk space that is required for installation.

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

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

Select the RPMs to be installed on all systems?

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

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

The installer installs the SF Sybase CE RPMs and patches.

7 Select the appropriate license option.

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

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

```
Enter a SF Sybase CE license key:
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X
```

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

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

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

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

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

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

If you enable Veritas Volume Replicator, the installer displays the following prompt:

```
Would you like to enable the
Veritas File Replicator? [y,n,q] (n)
```

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

The installer registers the license.

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

```
Would you like to configure SF Sybase CE on  
sys1 sys2 [y,n,q] (n) y
```

Note: If you had quit the installer before registering the sfsybasece license key, make sure the license key is registered on the system before starting the SF Sybase CE configuration. To register the license key, use the `installer -license` command.

- 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
```

- 11 Enter **y** if you want to view the summary file.

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

Configuring SF Sybase CE

This chapter includes the following topics:

- [About configuring SF Sybase CE](#)
- [Configuring the SF Sybase CE components using the script-based installer](#)

About configuring SF Sybase CE

You need to configure SF Sybase CE when:

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

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

SF Sybase CE configuration involves the following high-level tasks:

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

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

- Verifies the cluster information.
- Stops SF Sybase CE processes.
- Creates SF Sybase CE configuration files.
- Starts SF Sybase CE processes.

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

Configuring the SF Sybase CE components using the script-based installer

After installation, log in to the product installer to configure SF Sybase CE components. No configuration changes are made to the systems until all configuration questions are completed and confirmed.

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

Start the `installsfbasece` or `installer` program if you quit the installer after installation. If running the program from the `/opt/VRTS/install` directory, use `installsfbaseceversion`, where *version* is the specific release version.

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

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

Note: If you want to reconfigure SF Sybase CE, 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. You must also unmount the all VxFS/CFS mounts that are not configured under VCS.

To configure the SF Sybase CE components

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

SF Sybase CE installer Run the program:

```
# cd /opt/VRTS/install
# ./installsfbaseceversion \
  -configure sys1 sys2
```

Where *version* is the specific release version.
See [“About the Veritas installer”](#) on page 59.

Common product installer Run the program:

```
# ./installer -configure sys1 sys2
```

Choose **Veritas Storage Foundation for Sybase ASE CE** to configure SF Sybase CE.

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

- 3 Enter **1** to select the option **Configure SF Sybase CE sub-components**.

```
1) Configure Cluster File System
2) Configure I/O Fencing in Sybase Mode
3) Configure Sybase ASE CE Instance in VCS
4) Exit SFSYBASECE Configuration
Choose option: [1-4,q] (1)
```

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

```
The installer has discovered an existing installer process.
The process exited while performing configure of
SF Sybase CE on sys1.
Do you want to resume this process? [y,n,q,?] (y) n
```

- 5 Configure the Veritas Cluster Server component to set up the SF Sybase CE cluster.

See [“Configuring the SF Sybase CE cluster”](#) on page 72.

- 6 Add VCS users.
See [“Adding VCS users”](#) on page 81.
- 7 Configure SMTP email notification.
See [“Configuring SMTP email notification”](#) on page 82.
- 8 Configure SNMP trap notification.
See [“Configuring SNMP trap notification”](#) on page 84.

Configuring the SF Sybase CE cluster

Configure the systems on which you installed SF Sybase CE to be part of your cluster.

To configure a cluster for SF Sybase CE

- 1 Log in to the installer.
See [“Configuring the SF Sybase CE components using the script-based installer”](#) on page 70.
- 2 Select the **Configure Cluster File System** option from the main menu.
Press Return to continue.
If there are any SF Sybase CE processes running, these processes are stopped.
Press Return to continue.
- 3 VCS configuration includes configuring the cluster, users, secure mode if required, and notification.
To configure a cluster:
 - Configure the cluster name.
See [“Configuring the cluster name”](#) on page 72.
 - Configure private heartbeat links.
See [“Configuring private heartbeat links”](#) on page 73.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

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

```
Enter the unique cluster name: [q,?] clus1
```


Configuring private heartbeat links

You now configure the private heartbeat links 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 for high performance, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

The following procedure helps you configure LLT over Ethernet.

To configure private heartbeat links

- 1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or UDP.
 - Option 1: LLT over Ethernet (answer installer questions)
Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.
Skip to step 2.
 - Option 2: LLT over UDP (answer installer questions)
Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
Skip to step 3.
 - Option 3: Automatically detect configuration for LLT over Ethernet
Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.
Make sure that you activated the NICs for the installer to be able to detect and automatically configure the heartbeat links.
Skip to step 5.

Note: Option 3 is not available when the configuration is a single node configuration.

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

The installer discovers and lists the network interface cards.

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

Enter the NIC for the first private heartbeat link on sys1:

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

eth1 has an IP address configured on it. It could be a public NIC on sys1.

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 sys1:

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

eth2 has an IP address configured on it. It could be a public NIC on sys1.

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

Would you like to configure a third private heartbeat link?

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

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

```

Enter the NIC for the first private heartbeat
link on sys1: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on sys1: [b,q,?] (50000) ?
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat
link on sys1: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on sys1: [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 sys1: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on sys1: [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 sys1, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

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

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

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

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

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

- 6 Enter a unique cluster ID:

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

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

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

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

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

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

Do one of the following:

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

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[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)
```

If you want to set up trust relationships for your secure cluster, refer to the following topics:

See [“Configuring a secure cluster node by node”](#) on page 78.

Configuring Veritas Storage Foundation for Sybase ASE CE in secure mode

Configuring SF Sybase CE in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. SF Sybase CE user names and passwords are not used when a cluster is running in secure mode. You can select the secure mode to be FIPS compliant while configuring the secure mode.

To configure SF Sybase CE in secure mode

1 Enter appropriate choices when the installer prompts you:

```
Would you like to configure the VCS cluster in
secure mode [y,n,q] (n) y
1. Configure the cluster in secure mode without FIPS
2. Configure the cluster in secure mode with FIPS
3. Back to previous menu
Select the option you would like to perform [1-2,b,q] (1) 2
```

2 To verify the cluster is in secure mode after configuration, run the command:

```
# haclus -<value> SecureClus
```

The command returns 1 if cluster is in secure mode, else returns 0.

Configuring a secure cluster node by node

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

[Table 7-1](#) lists the tasks that you must perform to configure a secure cluster.

Table 7-1 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 78.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 79.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 80.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

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

```
# /opt/VRTS/install/installsfbasece<version> -securityonnode
```

Where *<version>* is the specific release version.

See [“About the Veritas installer”](#) on page 59.

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

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

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

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

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

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

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

Configuring the remaining nodes

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

To configure security on each remaining node

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

```
# /opt/VRTS/install/installsfsybasece<version> -securityonnode
```

Where *<version>* is the specific release version.

See [“About the Veritas installer”](#) on page 59.

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

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

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

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

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

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

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

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

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

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

```
# /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent
```

```
# /opt/VRTSvcs/bin/haconf -dump -makero
```

- 2 On the first node, stop the VCS engine.

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


- 3 On all nodes, stop the CmdServer.


```
# /opt/VRTSvcs/bin/CmdServer -stop
```
- 4 On the first node, edit the `/etc/VRTSvcs/conf/config/main.cf` file to resemble the following:


```
cluster clus1 (
SecureClus = 1
)
```
- 5 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.


```
# touch /etc/VRTSvcs/conf/config/.secure
```
- 6 On the first node, start VCS. Then start VCS on the remaining nodes.


```
# /opt/VRTSvcs/bin/hastart
```
- 7 On all nodes, start CmdServer.


```
# /opt/VRTSvcs/bin/CmdServer
```
- 8 On the first node, unfreeze the service groups.


```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hagrps -list Frozen=1
# /opt/VRTSvcs/bin/hagrps -unfreeze groupname -persistent
# /opt/VRTSvcs/bin/haconf -dump -makero
```

Adding VCS users

If you have enabled a secure VCS cluster, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

To add VCS users

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

```
Do you wish to accept the default cluster credentials of  
'admin/password'? [y,n,q] (y) n  
Enter the user name: [b,q,?] (admin)  
Enter the password:  
Enter again:
```

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

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

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

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

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

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

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

Configuring SMTP email notification

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

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

To configure SMTP email notification

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

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

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

See [“Configuring SNMP trap notification”](#) on page 84.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for the VCS Notifier to use on sys1:
[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
```

See “[Configuring global clusters](#)” on page 86.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for the VCS Notifier to use on sys1:
[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,?] sys5
```

- 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 sys5 [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,?] sys4
Enter the minimum severity of events for which SNMP traps
should be sent to sys4 [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: sys5 receives SNMP traps for Error or  
higher events
```

```
Console: sys4 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 SF Sybase CE global clusters.

To configure the global cluster option

- 1 Review the required information to configure the global cluster option.
- 2 Specify whether you want to configure the global cluster option.

```
Do you want to configure the Global Cluster Option? [y,n,q] (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)
```

For complete details for configuring your disaster recovery environment once clusters are installed and configured:

See the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*.

Configuring SF Sybase CE clusters for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using `installsfsybasece`](#)

Setting up disk-based I/O fencing using `installsfsybasece`

You can configure I/O fencing using the `-fencing` option of the `installsfsybasece`.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

- 1 List the new external disks or the LUNs as recognized by the operating system. On each node, enter:

```
# fdisk -l
```

- 2 To initialize the disks as VxVM disks, use one of the following methods:
 - Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks. For more information see the *Veritas Storage Foundation Administrator's Guide*.
 - Use the `vxdisksetup` command to initialize a disk as a VxVM disk.

```
# vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i sdr
```

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

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 89.

Review the following procedure to identify disks to use as coordinator disks.

To identify the coordinator disks

- 1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# vxdisk -o alldgs list
```

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

See [“Checking shared disks for I/O fencing”](#) on page 90.

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure SF Sybase CE meets the I/O fencing requirements. You can test the shared disks using the `vxfsentsthdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfsenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

Make sure to test the disks that serve as coordinator disks.

The `vxfsentsthdw` utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

See the *Veritas Storage Foundation for Sybase ASE CE Administrator's Guide*.

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

- Verifying the Array Support Library (ASL)

See [“Verifying Array Support Library \(ASL\)”](#) on page 91.

- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 92.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfcntl utility”](#) on page 93.

Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

To verify Array Support Library (ASL)

- 1 If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

The ASL for the supported storage device that you add is available from the disk array vendor or Symantec technical support.

- 2 Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

The following output is a sample:

```
# vxddladm listsupport all
```

```
LIBNAME          VID          PID
=====
libvxhitachi.so  HITACHI      DF350, DF400, DF400F,
                  DF500, DF500F
libvxxp1281024.so HP           All
libvxxp12k.so    HP           All
libvxdds2a.so    DDN          S2A 9550, S2A 9900,
                  S2A 9700
libvxpurple.so   SUN          T300
libvxxiotechE5k.so XIOTECH      ISE1400
libvxcopan.so   COPANSYS     8814, 8818
libvxibm8k.so    IBM          2107
```

- 3 Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

```
# vxdisk scandisks
```

See the Veritas Volume Manager documentation for details on how to add and configure disks.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the `vxfcntlshdw` utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SF Sybase CE.
- 2 Ensure that both nodes are connected to the same disk during the testing. Use the `vxfenadm` command to verify the disk serial number.

```
# vxfenadm -i diskpath
```

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

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

From node A, enter:

```
# vxfenadm -i /dev/sdx
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

```
Vendor id : EMC  
Product id : SYMMETRIX  
Revision : 5567  
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/sdy` path.

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

```
# vxfenadm -i /dev/sdz
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

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

Testing the disks using `vxfcntlsthdw` utility

This procedure uses the `/dev/sdx` disk in the steps.

If the utility does not show a message that states a disk is ready, the verification has failed. Failure of verification can be the result of an improperly configured disk array. The failure can also be due to a bad disk.

If the failure is due to a bad disk, remove and replace it. The `vxfentsthdw` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/sdx is ready to be configured for I/O Fencing on
node sys1
```

For more information on how to replace coordinator disks, refer to the *Veritas Storage Foundation for Sybase ASE CE Administrator's Guide*.

To test the disks using `vxfentsthdw` utility

- 1 Make sure system-to-system communication functions properly.
- 2 From one node, start the utility.
Run the utility with the `-n` option if you use `rsh` for communication.

```
# vxfentsthdw [-n]
```

- 3 The script warns that the tests overwrite data on the disks. After you review the overview and the warning, confirm to continue the process and enter the node names.

Warning: The tests overwrite and destroy data on the disks unless you use the `-r` option.

```
***** WARNING!!!!!!!!!! *****
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: sys1
Enter the second node of the cluster: sys2
```

- 4 Enter the names of the disks that you want to check. Each node may know the same disk by a different name:

```
Enter the disk name to be checked for SCSI-3 PGR on node
sys1 in the format:
for dmp: /dev/vx/rdmp/sdx
for raw: /dev/sdx
Make sure it's the same disk as seen by nodes sys1 and sys2
/dev/sdx
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
sys2 in the format:
for dmp: /dev/vx/rdmp/sdx
for raw: /dev/sdx
Make sure it's the same disk as seen by nodes sys1 and sys2
/dev/sdx
```

If the serial numbers of the disks are not identical, then the test terminates.

- 5 Review the output as the utility performs the checks and reports its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node sys1.

```
The disk is now ready to be configured for I/O Fencing on node
sys1
```

```
ALL tests on the disk /dev/sdx have PASSED
The disk is now ready to be configured for I/O Fencing on node
sys1
```

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

Configuring disk-based I/O fencing using installsfybasece

Note: The installer stops and starts SF Sybase CE to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop SF Sybase CE.

To set up disk-based I/O fencing using the `installsfsybasece`

- 1 Start the `installsfsybasece` with `-fencing` option.

```
# /opt/VRTS/install/installsfsybasece<version> -fencing
```

Where `<version>` is the specific release version.

See [“About the Veritas installer”](#) on page 59.

The `installsfsybasece` starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether SF Sybase CE 6.0.1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **1** to configure fencing in Sybase mode.

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-3,b,q] 1
```

- 4 Review the output as the configuration program checks whether VxVM is already started and is running.

- If the check fails, configure and enable VxVM before you repeat this procedure.
- If the check passes, then the program prompts you for the coordinator disk group information.

- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.
The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
 - Enter the number corresponding to the **Create a new disk group** option.

The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.

Symantec recommends that you use three disks as coordination points for disk-based I/O fencing.

If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
- Enter the disk group name.

- 6** Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfsntsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 90.

- 7** After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 8** Enter the I/O fencing disk policy that you chose to use. For example:

```
Enter disk policy for the disk(s) (raw/dmp): [b,q,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
 - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9** Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 10** Review the output as the configuration program does the following:
- Stops VCS and I/O fencing on each node.
 - Configures disk-based I/O fencing and starts the I/O fencing process.
 - Updates the VCS configuration file `main.cf` if necessary.
 - Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Updates the I/O fencing configuration file `/etc/vxfenmode`.

- Starts VCS on each node to make sure that the SF Sybase CE is cleanly configured to use the I/O fencing feature.

11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.

12 Configure the Coordination Point Agent.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)
```

13 Enter a name for the service group for the Coordination Point Agent.

```
Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen) vxfen
```

14 Set the level two monitor frequency.

```
Do you want to set LevelTwoMonitorFreq? [y,n,q] (y)
```

15 Decide the value of the level two monitor frequency.

```
Enter the value of the LevelTwoMonitorFreq attribute: [b,q,?] (5)
```

Installer adds Coordination Point Agent and updates the main configuration file.

Post-installation tasks

- [Chapter 9. Verifying the installation](#)
- [Chapter 10. Performing additional post-installation and configuration tasks](#)

Verifying the installation

This chapter includes the following topics:

- [Performing a postcheck on a node](#)
- [Verifying SF Sybase CE installation using VCS configuration file](#)
- [Verifying LLT, GAB, and cluster operation](#)

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See [“About using the postcheck option”](#) on page 284.

To run the postcheck command on a node

- 1 Run the installer with the `-postcheck` option.

```
# ./installer -postcheck system_name
```

- 2 Review the output for installation-related information.

Verifying SF Sybase CE installation using VCS configuration file

The configuration file, `main.cf`, is created on each node at `/etc/VRTSvcs/conf/config/`. Review the `main.cf` configuration file after the SF Sybase CE installation and before the Sybase installation.

Verify the following information in the `main.cf` file:

- The cluster definition within the `main.cf` includes the cluster information that was provided during the configuration. The cluster information includes the cluster name, cluster address, and the names of cluster users and administrators.
- The `UseFence = SCSI3` attribute is present in the file.
- If you configured the cluster in secure mode, the “`SecureClus = 1`” cluster attribute is set.

For more information on the configuration file:

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 102.
- 4 Verify GAB operation.
See “[Verifying GAB](#)” on page 104.
- 5 Verify the cluster operation.
See “[Verifying the cluster](#)” on page 106.

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 `sys1`.
- 2 Run the `lltstat` command on the node `sys1` to view the status of LLT.

```
lltstat -n
```



```

1 sys2          OPEN

                eth1 UP      08:00:20:8F:D1:F2
                eth2 DOWN

```

The command reports the status on the two active nodes in the cluster, sys1 and sys2.

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

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

However, the output in the example shows different details for the node sys2. 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 sys1 in a two-node cluster:

```
lltstat -p
```

The output resembles:

```

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

```

Verifying GAB

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

[Table 9-1](#) lists the different ports that the software configures for different functions.

Table 9-1 GAB port description

Port	Function
a	GAB
b	I/O fencing
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 . Run the <code>vxctl protocolversion</code> command to check the protocol version.
v	Cluster Volume Manager (CVM)
w	vxconfigd (module for CVM)
y	Cluster Volume Manager (CVM) I/O shipping

For more information on GAB, refer to the *Veritas Cluster Server Administrator's Guide*.

To verify GAB

- ◆ To verify the GAB operation, type the following command on each node:

```
# /sbin/gabconfig -a
```

For example, the command returns the following output:

```
GAB Port Memberships
=====
Port a gen ada401 membership 01
Port b gen ada40d membership 01
Port d gen ada409 membership 01
Port f gen ada41c membership 01
Port h gen ada40f membership 01
Port o gen ada406 membership 01
Port u gen ada41a membership 01
Port v gen ada416 membership 01
Port w gen ada418 membership 01
Port y gen ada42a membership 01
```

Verifying the cluster

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

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

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

To verify the cluster

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

```
# hastatus -summary
```

The output resembles:

```

-- SYSTEM STATE
-- System                State                Frozen

A sys1                   RUNNING            0
A sys2                   RUNNING            0

-- GROUP STATE
-- Group                 System             Probed  AutoDisabled  State

```

- 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.

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 `sys1`. The list continues with similar information for `sys2` (not shown) and any other nodes in the cluster.

```

#System  Attribute                Value
sys1     AgentsStopped            0
sys1     AvailableCapacity        100

```

```

sys1      CPUThresholdLevel    Critical 90 Warning 80 Note 70
Info 60

sys1      CPUUsage           0

sys1      CPUUsageMonitoring  Enabled 0 ActionThreshold 0
ActionTimeLimit 0 Action NONE
NotifyThreshold 0 NotifyTimeLimit 0

sys1      Capacity          100

sys1      ConfigBlockCount    293

sys1      ConfigChecksum      37283

sys1      ConfigDiskState     CURRENT

sys1      ConfigFile          /etc/VRTSvcs/conf/config

sys1      ConfigInfoCnt       0

sys1      ConfigModDate      Mon Sep 03 07:14:23 CDT 2012

sys1      ConnectorState     Up

sys1      CurrentLimits

sys1      DiskHbStatus

sys1      DynamicLoad          0

sys1      EngineRestarted     0

sys1      EngineVersion        6.0.10.0

sys1      FencingWeight       0

sys1      Frozen              0

sys1      GUIIPAddr

sys1      HostUtilization      CPU 0 Swap 0

sys1      LLTNodeId            0

sys1      LicenseType          PERMANENT_SITE

sys1      Limits

sys1      LinkHbStatus         eth1 UP eth2 UP

```

```

sys1      LoadTimeCounter      0
sys1      LoadTimeThreshold    600
sys1      LoadWarningLevel     80
sys1      NoAutoDisable        0
sys1      NodeId               0
sys1      OnGrpCnt             7
sys1      PhysicalServer
sys1      ShutdownTimeout      600
sys1      SourceFile           ./main.cf
sys1      SwapThresholdLevel    Critical 90 Warning 80 Note 70
                                   Info 60
sys1      SysName              sys1
sys1      SysState             RUNNING
sys1      SystemLocation
sys1      SystemOwner
sys1      SystemRecipients
sys1      TFrozen              0
sys1      TRSE                 0
sys1      UpDownState          Up
sys1      UserInt              0
sys1      UserStr
sys1      VCSFeatures          DR
sys1      VCSMode

```


Performing additional post-installation and configuration tasks

This chapter includes the following topics:

- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Configuring Veritas Volume Replicator](#)
- [Running SORT Data Collector to collect configuration information](#)

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 OpenLDAP and Windows Active Directory.

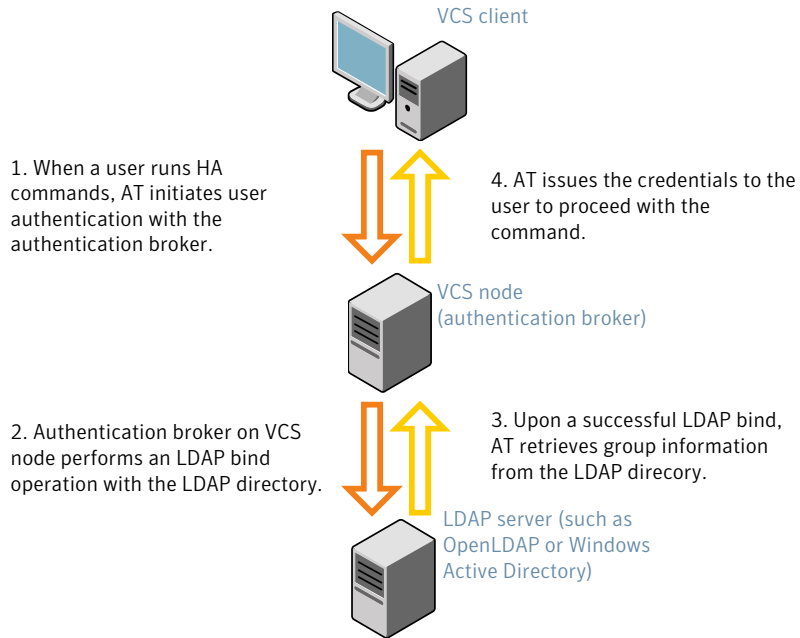
For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

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

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

[Figure 10-1](#) depicts the SF Sybase CE cluster communication with the LDAP servers when clusters run in secure mode.

Figure 10-1 Client communication with LDAP servers



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

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

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is `posixAccount`)
 - UserObject Attribute (the default is `uid`)
 - User Group Attribute (the default is `gidNumber`)
 - Group Object Class (the default is `posixGroup`)
 - GroupObject Attribute (the default is `cn`)
 - Group GID Attribute (the default is `gidNumber`)
 - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, `UserBaseDN=ou=people,dc=comp,dc=com`)

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

Enabling LDAP authentication for clusters that run in secure mode

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

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

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.6.0
```

To enable OpenLDAP authentication for clusters that run in secure mode

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

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-d -s domain_controller_name_or_ipaddress -u domain_user
```

Attribute list file name not provided, using AttributeList.txt

Attribute file created.

You can use the `cat` command to view the entries in the attributes file.

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

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-c -d windows_domain_name
```

Attribute list file not provided, using default AttributeList.txt

CLI file name not provided, using default CLI.txt

CLI for addldapdomain generated.

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

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x
```

```
Using default broker port 2821
```

```
CLI file not provided, using default CLI.txt
```

```
Looking for AT installation...
```

```
AT found installed at ./vssat
```

```
Successfully added LDAP domain.
```

- 4 Check the AT version and list the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.12.0

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains
Domain Name : mydomain.com
Server URL : ldap://192.168.20.32:389
SSL Enabled : No
User Base DN : CN=people,DC=mydomain,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 Check the other domains in the cluster.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showdomains -p vx
```

The command output lists the number of domains that are found, with the domain names and domain types.

6 Generate credentials for the user.

```
# unset EAT_LOG

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate \
-d ldap:windows_domain_name -p user_name -s user_password -b \
localhost:14149
```

7 Add non-root users as applicable.

```
# useradd user1

# passwd pw1

Changing password for "user1"

user1's New password:

Re-enter user1's new password:

# su user1

# bash

# id

uid=204(user1) gid=1(staff)

# pwd

# mkdir /home/user1

# chown user1 /home/ user1
```

8 Log in as non-root user and run `ha` commands as LDAP user.

```
# cd /home/user1

# ls

# cat .vcspwd

101 localhost mpise LDAP_SERVER ldap

# unset VCS_DOMAINTYPE

# unset VCS_DOMAIN

# /opt/VRTSvcs/bin/hasys -state

#System      Attribute      Value
cluster1:sysA  SysState      FAULTED
cluster1:sysB  SysState      FAULTED
cluster2:sysC  SysState      RUNNING
cluster2:sysD  SysState      RUNNING
```

Configuring Veritas Volume Replicator

Perform this step only if you have not already configured VVR during the installation.

By default, the installer installs the required VVR configuration files irrespective of whether or not you choose to enable VVR. To configure VVR manually in SF Sybase CE, simply start VVR using the `vxstart_vvr` command. The command starts the VVR daemons and configures the ports. You may change the default settings at any time.

For instructions on changing the default settings, see the *Veritas Volume Replicator Administrator's Guide*.

To configure VVR

- 1 Log into each node in the cluster as the root user.
- 2 Start VVR:

```
# vxstart_vvr start
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmind: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
VxVM VVR V-5-2-5942 Starting Communication daemon: [OK]
```

Running SORT Data Collector to collect configuration information

SORT Data Collector now supersedes the VRTSexplorer utility. Run the Data Collector with the `VxExplorer` option to gather information about the system.

Visit the SORT Website and download the UNIX Data Collector appropriate for your operating system.

<https://sort.symantec.com>

For more information:

<https://sort.symantec.com/public/help/wwhelp/wwhimpl/js/html/wwhelp.htm>

Upgrade of SF Sybase CE

- Chapter 11. Planning to upgrade SF Sybase CE
- Chapter 12. Performing a full upgrade of SF Sybase CE using the product installer
- Chapter 13. Performing an automated full upgrade of SF Sybase CE using response files
- Chapter 14. Performing a phased upgrade of SF Sybase CE
- Chapter 15. Performing a rolling upgrade of SF Sybase CE
- Chapter 16. Performing post-upgrade tasks

Planning to upgrade SF Sybase CE

This chapter includes the following topics:

- [About types of upgrade](#)
- [Supported upgrade paths](#)

About types of upgrade

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

SF Sybase CE does not support upgrades using the Web installer.

[Table 11-1](#) lists the supported types of upgrade.

Table 11-1 Types of upgrade

Type of upgrade	Method of upgrade	Procedures
Full upgrade	Veritas script-based installation programs <ul style="list-style-type: none"> ■ Interactive mode ■ Non-interactive mode using response files 	Complete the following steps: <ul style="list-style-type: none"> ■ Preparing to upgrade ■ Upgrading to SF Sybase CE 6.0.1 See the chapter <i>Performing a full upgrade to SF Sybase CE 6.0.1</i>. ■ Completing post-upgrade tasks See the chapter <i>Performing post-upgrade tasks</i>.

Table 11-1 Types of upgrade (*continued*)

Type of upgrade	Method of upgrade	Procedures
Phased upgrade	Combination of manual steps and the Veritas script-based installation programs	Complete the steps in the chapter <i>Performing a phased upgrade to SF Sybase CE 6.0.1</i> .
Rolling upgrade	Veritas script-based installation programs	Complete the steps in the chapter <i>Performing a rolling upgrade to SF Sybase CE 6.0.1</i> .

Supported upgrade paths

The upgrade paths are categorized based on the combination of upgrades you plan to perform.

[Table 11-2](#) lists the supported upgrade paths for upgrades on SLES.

Table 11-2 Supported upgrade paths on SLES

From product version	From OS version	To SF Sybase CE version	To OS version	Supported upgrade type
SF Sybase CE 5.0	SLES 10 SP2 SLES 10	SF Sybase CE 6.0.1	SLES 10 SP4	Full or phased upgrade
SF Sybase CE 5.0 P1	SLES 10 SP2 SLES 10	SF Sybase CE 6.0.1	SLES 10 SP4	Full or phased upgrade
SF Sybase CE 5.0 PR1	SLES 10 SP3 SLES 10	SF Sybase CE 6.0.1	SLES 10 SP4	Full or phased upgrade
5.0 PR1 P1	SLES 10 SP3 SLES 10	SF Sybase CE 6.0.1	SLES 10 SP4	Full or phased upgrade
6.0	SLES 10 SP4	SF Sybase CE 6.0.1	SLES 10 SP4	Full, rolling, or phased upgrade

Performing a full upgrade of SF Sybase CE using the product installer

This chapter includes the following topics:

- [About full upgrades](#)
- [Preparing to perform a full upgrade to SF Sybase CE 6.0.1](#)
- [Upgrading to SF Sybase CE 6.0.1](#)

About full upgrades

A full upgrade involves upgrading all the nodes in the cluster at the same time. All components are upgraded during the process. The cluster remains unavailable for the duration of the upgrade.

Note: You can not roll back the upgrade to a previous version after you upgrade to version 6.0.1.

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

- Common product installer (`installer`)
The common product installer provides menu options for installing and configuring multiple Veritas products.
- SF Sybase CE installation programs (`installsfsybasece`)

The SF Sybase CE installation program provide menu options for installing and configuring SF Sybase CE.

Note: If you obtained SF Sybase CE from an electronic download site, you must use the product installer (`installsfbasece`) instead of the common product installer (`installer`).

You can also perform a full upgrade using a response file. You can create a response file by using the response file template or by customizing a response file that is generated by the script-based installer.

For more information about response files:

See [“About response files”](#) on page 189.

Preparing to perform a full upgrade to SF Sybase CE 6.0.1

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

To prepare to upgrade SF Sybase CE

- 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`, `SybaseTypes.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.save
# cp /etc/VRTSvcs/conf/config/SybaseTypes.cf \
/var/tmp/SybaseTypes.cf.save
```

- 3 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

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

If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 5 ■ If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

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

- 6 Stop the Sybase Binaries service group (binmnt group).

```
# hagrps -offline binmnt -any
```

- 7 ■ If the Sybase database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify sybase_group AutoStart 0
# haconf -dump -makero
```

- 8 Stop VCS on all nodes:

```
# hastop -all
```

One way to check whether or not the configuration is valid is to check the `main.cf` file as follows:

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

However, this method can not verify whether all configurations are valid. If SF Sybase CE was running properly before the upgrade, the configurations are valid.

- 9 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
```

- 10 If you plan to upgrade the operating system, stop all ports.

Upgrading to SF Sybase CE 6.0.1

This section provides instructions for the following upgrade scenarios:

SF Sybase CE and
minor operating
system upgrade

Perform the steps in the following procedure if you plan to perform a minor upgrade of the operating system, for example from SLES 10 SP2 to SLES 10 SP4, along with SF Sybase CE.

See [“To upgrade SF Sybase CE and operating system \(minor OS upgrade\)”](#) on page 128.

To upgrade SF Sybase CE and operating system (minor OS upgrade)

- 1 If you want to upgrade the operating system, perform the following steps:

- Rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- Upgrade the operating system on all nodes in the cluster.
For instructions, see the operating system documentation.

Note: If reboot is required, use `shutdown -r now` command to reboot the nodes.

- After the system restarts, restore the `/etc/llttab` file to its original name:


```
# mv /etc/llttab.save /etc/llttab
```

- 2 Upgrade to SF Sybase CE 6.0.1 using the script-based installer.

See [“Upgrading SF Sybase CE using the Veritas script-based installation program”](#) on page 130.

You can also perform a silent upgrade:

See [“Upgrading SF Sybase CE using a response file”](#) on page 135.

- 3 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 4 Bring the sybasece resource group online.

```
# hagr -online sybasece -sys node_name
```

- 5 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 6 ■ If the Sybase database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the Sybase Binaries service group (binmnt) and sybasece service group online automatically when VCS starts:

```
# haconf -makerw
# hagr -modify sybasece AutoStart 1
# haconf -dump -makero
```

- 7 Complete other post-upgrade steps.

For instructions, see the chapter *Performing post-upgrade tasks* in this document.

- See [“Re-joining the backup boot disk group into the current disk group”](#) on page 159.
- See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 160.
- See [“Setting or changing the product license level”](#) on page 160.
- See [“Upgrading disk layout versions”](#) on page 161.
- See [“Upgrading CVM protocol version and VxVM disk group version ”](#) on page 161.
-

8 Find out which node is the CVM master. Enter the following:

```
# vxdtl -c mode
```

9 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdtl upgrade
```

Upgrading SF Sybase CE using the Veritas script-based installation program

Use one of the following Veritas script-based installation programs to upgrade SF Sybase CE: `installer` or `installsfybasece`

The installer performs the following tasks to upgrade SF Sybase CE:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Sybase CE processes before the upgrade.
- Uninstalls SF Sybase CE.
- Installs the SF Sybase CE 6.0.1 RPMs on the nodes.
- Starts SF Sybase CE 6.0.1 on all the nodes.
- Displays the location of the log files, summary file, and response file.

To upgrade to SF Sybase CE 6.0.1 using the `installsfsybasece` program

1 Start the installation program using one of the following ways:

SF Sybase CE installer	Navigate to the product directory on the installation media that contains the installation program.
------------------------	---

The program is located in the `storage_foundation_for_sybase_ce` directory.

Run the program:

```
# ./installsfsybasece sys1 sys2
```

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

Run the program:

```
# ./installer sys1 sys2
```

From the opening Selection Menu, choose **G** for "**Upgrade a Product.**"

Select the option **Full Upgrade.**"

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

The installer verifies the systems for compatibility.

Note: If `had` is stopped before upgrade, the installer displays the following warning:

VCS is not running before upgrade. Please make sure all the configurations are valid before upgrade.

If the configuration files are valid, you may ignore the message.

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

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

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

Enter **y** to agree to the End User License Agreement (EULA).

The installer displays the list of RPMs that will be uninstalled. Press **Enter** to view the list of RPMs that will be upgraded.

The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's boot disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.

- 3 Enter the name of the backup boot disk group when prompted. Press **Enter** to accept the default.

You are prompted to start the split operation.

- 4 Enter **y** to continue with the split operation.

The split operation can take some time to complete.

Note: Verify the boot device from which the system is set to boot. Make sure that the system is set to start from the boot device with the required version of SF Sybase CE.

5 Enter y to stop the SF Sybase CE processes.

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

The installer stops the processes and uninstalls SF Sybase CE. After the uninstallation, the installer installs SF Sybase CE 6.0.1 and starts SF Sybase CE 6.0.1 on all the nodes.

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFSYBASECE is licensed on the systems
```

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

6 Complete the remaining tasks to finish the upgrade:

See [“Upgrading to SF Sybase CE 6.0.1”](#) on page 128.

Performing an automated full upgrade of SF Sybase CE using response files

This chapter includes the following topics:

- [Upgrading SF Sybase CE using a response file](#)
- [Response file variables to upgrade Veritas Storage Foundation for Sybase ASE CE](#)
- [Sample response file for upgrading SF Sybase CE](#)

Upgrading SF Sybase CE using a response file

Perform the steps in the following procedure to upgrade to SF Sybase CE 6.0.1 using a response file.

To upgrade SF Sybase CE 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 189.

For response file variable definitions:

See [“Response file variables to upgrade Veritas Storage Foundation for Sybase ASE CE”](#) on page 136.

For a sample response file:

See [“Sample response file for upgrading SF Sybase CE”](#) on page 138.

- 3 Navigate to the product directory on the installation media that contains the SF Sybase CE installation program.
- 4 Start the installation:

```
# ./installsfbasece -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 Veritas Storage Foundation for Sybase ASE CE

[Table 13-1](#) lists the response file variables that you can define to configure SF Sybase CE.

Table 13-1 Response file variables for upgrading SF Sybase CE

Variable	Description
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
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}{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{mirrordgname}{system}	If the root dg is encapsulated and you select split mirror is selected: Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional

Table 13-1 Response file variables for upgrading SF Sybase CE (*continued*)

Variable	Description
CFG{splitmirror}{system}	If the root dg is encapsulated and you select split mirror is selected: Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Sample response file for upgrading SF Sybase CE

The following sample response file performs a full upgrade on the system sys1.

```
our %CFG;  
  
$CFG{accepteula}=1;  
$CFG{opt}{gco}=1;  
$CFG{opt}{upgrade}=1;  
$CFG{opt}{vr}=1;  
$CFG{systems}=[ qw(sys1) ];  
$CFG{vcs_allowcomms}=1;
```

Performing a phased upgrade of SF Sybase CE

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing phased upgrade of SF Sybase CE from version 5.0 and later releases](#)

About phased upgrade

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

For supported upgrade paths:

See “[Supported upgrade paths](#)” on page 124.

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 examples in the procedures assume a four-node SF Sybase CE cluster with the nodes *sys1* and *sys2* constituting the first half of the cluster and the nodes *sys3* and *sys4* constituting the second half of the cluster.




Performing phased upgrade of SF Sybase CE from version 5.0 and later releases

Table 14-1 illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

Table 14-1 Summary of phased upgrade

First half of the cluster	Second half of the cluster
SF Sybase CE cluster before the upgrade:	
<p>STEP 1: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> ■ Switch failover applications. ■ Stop all parallel applications. <p>See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 141.</p> <p>STEP 2: Upgrade SF Sybase CE.</p> <p>See “Step 2: Upgrading the first half of the cluster” on page 143.</p>	<p>The second half of the cluster is up and running.</p>
<p>The first half of the cluster is not running.</p>	<p>STEP 3: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> ■ Stop all parallel and failover applications. ■ Stop SF Sybase CE. <p>See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 145.</p> <p>The downtime starts now.</p>

Table 14-1 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
<p>STEP 4: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> ■ Start SF Sybase CE. ■ Start all applications. <p>See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 147.</p> <p>The downtime ends here.</p>	<p>The second half of the cluster is not running.</p> 
<p>The first half of the cluster is up and running.</p> 	<p>STEP 5: Upgrade SF Sybase CE.</p> <p>See “Step 5: Upgrading the second half of the cluster” on page 148.</p> <p>STEP 6: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> ■ Start SF Sybase CE. ■ Start all applications. <p>See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 149.</p>
<p>The phased upgrade is complete and both the first and the second half of the cluster are running.</p> 	

Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `SybaseTypes.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

- 2 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

- 3 Stop all applications that are not configured under VCS but dependent on Sybase ASE CE or resources controlled by VCS. Use native application commands to stop the application.
- 4 Stop the applications configured under VCS. Take the Sybase database group offline.

```
# hagr -offline sybase_group -sys sys1  
# hagr -offline sybase_group -sys sys2
```

- 5 Stop the Sybase Binaries service group (binmnt group).

```
# hagr -offline binmnt -sys sys1  
  
# hagr -offline binmnt -sys sys2
```

- 6 If the Sybase database is managed by VCS, set the `AutoStart` value to 0 to prevent the service group from starting automatically when VCS starts:

```
# haconf -makerw  
# hagr -modify sybasece AutoStart 0  
# haconf -dump -makero
```

- 7 Unmount the CFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:


```
# mount | grep vxfs | grep cluster
```

```
# fuser -cu /mount_point
```
 - Unmount the non-system CFS file system:


```
# umount /mount_point
```
- 8 Stop the parallel service groups and switch over failover service groups on each of the nodes in the first half of the cluster:


```
# hastop -local
```
- 9 Unmount the VxFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:


```
# mount | grep vxfs
```

```
# fuser -cu /mount_point
```
 - Unmount the non-system VxFS file system:


```
# umount /mount_point
```
- 10 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.
- 11 If you plan to upgrade the operating system, stop all ports on first half of the cluster sys1, sys2.


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

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

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

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```


- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 Upgrade SF Sybase CE. Navigate to the product directory on the installation media. When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd product folder
```

For SLES 10 (x86_64)

```
# cd /dvd_mount/sles10_x86_64/\
storage_foundation_for_sybase_ce
```

```
# ./installsfybasece -upgrade sys1 sys2
```

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.

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFSYBASECE is licensed on the systems
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Sybase ASE CE or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop the applications configured under VCS. Take the Sybase database group offline.

```
# hagr -offline sybase_group -sys sys3
```

```
# hagr -offline sybase_group -sys sys4
```

- 3 Unmount the CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster
```

```
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 4 Stop VCS on each of the nodes in the second half of the cluster:

```
# hastop -local
```

- 5 Unmount the VxFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
```

```
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

6 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

7 Stop all ports.

While upgrading from 5.0 or later:

```
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

While upgrading from 6.0:

```
# /opt/VRTS/install/installsfsybasece -stop sys3 sys4
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
# /etc/init.d/gab start

# gabconfig -x
```

2 On the first half of the cluster, start SF Sybase CE:

```
# cd /opt/VRTS/install
# ./installsfsybasece<version> -start sys1 sys2
```

Where *<version>* is the specific release version.

See [“About the Veritas installer”](#) on page 59.

3 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.

- 4 Bring the sybasece group online.

```
# hagrps -online sybasece -sys sys1
```

```
# hagrps -online sybasece -sys sys2
```

Note: The downtime ends here.

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

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 On the second half of the cluster, upgrade SF Sybase CE. Navigate to the product directory on the installation media.

Invoke the SF Sybase CE installer with the `-upgrade` option. The installer upgrades the second half of the cluster.

For SLES 10 (x86_64)

```
# cd /dvd_mount/sles10_x86_64/\
storage_foundation_cluster_file_system_for_sybase_ce

# ./installsfybasece -upgrade sys3 sys4
```

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFSYBASECE license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.

SFSYBASECE is licensed on the systems
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 On the second half of the cluster, start SF Sybase CE:

```
# cd /opt/VRTS/install

# ./installsfybasece<version> -start sys3 sys4
```

Where `<version>` is the specific release version.

See [“About the Veritas installer”](#) on page 59.

- 3 Upgrade VxVM disk group version.
See [“Upgrading CVM protocol version and VxVM disk group version”](#) on page 161.
- 4 Upgrade disk layout version.
See [“Upgrading disk layout versions”](#) on page 161.
- 5 Bring the sybasece group online.

```
# hagrps -online sybasece_group -sys sys3  
  
# hagrps -online sybasece_group -sys sys4
```
- 6 If the Sybase database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw  
# hagrps -modify sybasece AutoStart 1  
# haconf -dump -makero
```
- 7 Set or change the product license level, if required.
See [“Setting or changing the product license level”](#) on page 160.

Note: In case of Sybase ASE CE version prior to 15.5 ASE CE, upgrade the database to 15.5 ASE CE or 15.5 ASE CE latest ESD, after upgrading the cluster.

See [“Upgrading Sybase ASE CE”](#) on page 185.

Performing a rolling upgrade of SF Sybase CE

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Supported rolling upgrade paths](#)
- [Preparing to perform a rolling upgrade to SF Sybase CE 6.0.1](#)
- [Performing a rolling upgrade using the installer](#)

About rolling upgrades

The rolling upgrade minimizes downtime for highly available clusters to the amount of time that it takes to perform a service group failover. The rolling upgrade has two main phases where the installer upgrades kernel RPMs in phase 1 and VCS agent RPMs in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

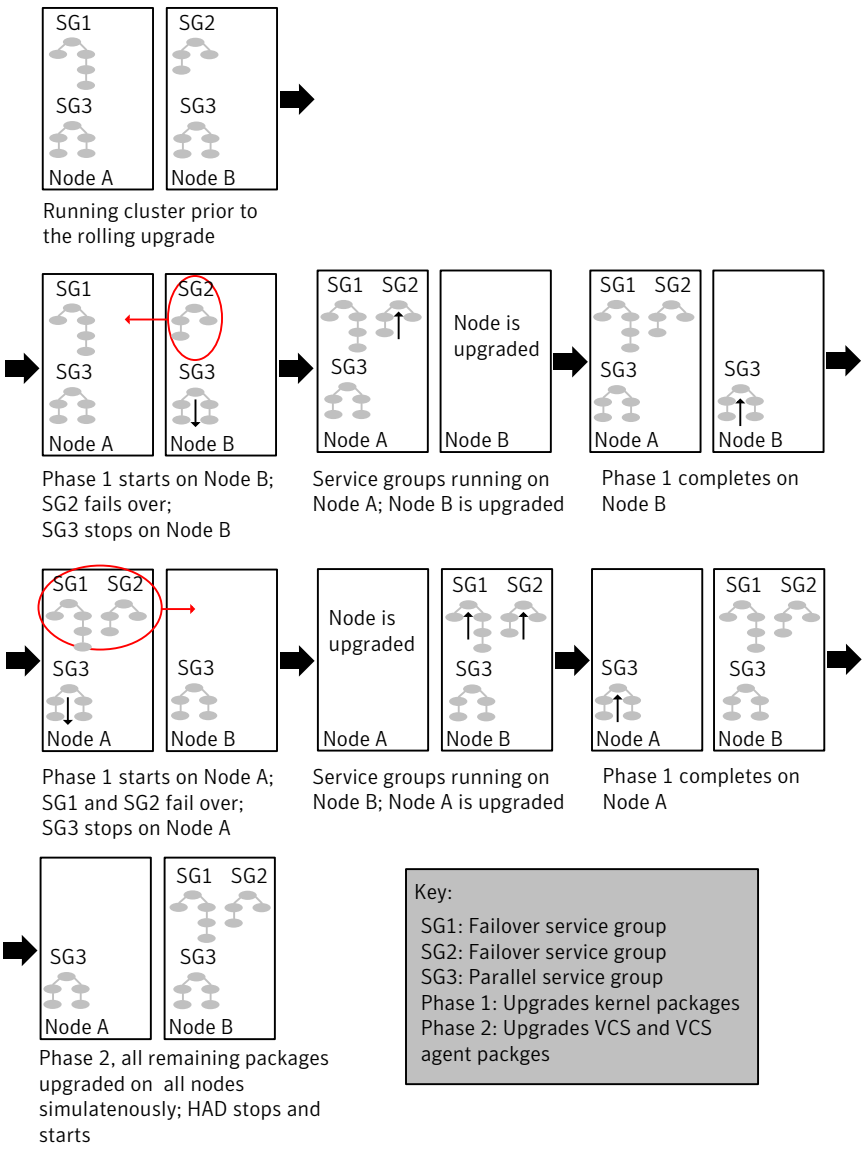
The following is an overview of the flow for a rolling upgrade:

1. The installer performs prechecks on the cluster.
2. Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over. The downtime is limited to the applications that are failed over and not the entire cluster.

3. The installer performs the second phase of the upgrade on all of the nodes in the cluster. The second phase of the upgrade includes downtime of the Veritas Cluster Server (VCS) engine HAD, but does not include application downtime.

[Figure 15-1](#) illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 15-1 Example of the installer performing a rolling upgrade



The following limitations apply to rolling upgrades:

- Rolling upgrades are not compatible with phased upgrades. Do not mix rolling upgrades and phased upgrades.

Supported rolling upgrade paths

You can perform a rolling upgrade of SF Sybase CE with the script-based installer. The rolling upgrade procedures support minor operating system upgrades.

Table 15-1 shows the versions of SF Sybase CE for which you can perform a rolling upgrade to SF Sybase CE 6.0.1.

Table 15-1 Supported rolling upgrade paths

Platform	SF Sybase CE version
Linux	6.0, 6.0RP1
SLES 10 SP4	

Preparing to perform a rolling upgrade to SF Sybase CE 6.0.1

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: If you plan to upgrade the operating system, make sure that you upgrade all nodes before you start rolling upgrade of SF Sybase CE.

To prepare to upgrade SF Sybase CE

Perform the steps on the first subcluster.

- 1 Log in as superuser to one of the nodes in the subcluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `SybaseTypes.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

- 3 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

Warning: Backup `/etc/vx/cbr/bk` directory.

- 4 Switch over all failover service groups to the nodes in the other subcluster:

```
# hagr -switch grp_name -to sys_name
```

- 5 Stop the applications configured under VCS. Take the Sybase database group offline.

```
# hagr -offline sybase_group -sys sys1
```

```
# hagr -offline sybase_group -sys sys2
```

- 6 Stop the Sybase Binaries service group (binmnt group).

```
# hagr -offline binmnt -sys sys1
```

```
# hagr -offline binmnt -sys sys2
```

- 7 If the Sybase database is managed by VCS, set the AutoStart value to 0 to prevent the service group from starting automatically when VCS starts:

```
# haconf -makerw
```

```
# hagr -modify sybasece AutoStart 0
```

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

- 8 Unmount all the CFS file system which is not under VCS control.

```
# mount |grep vxfs | grep cluster
```

```
# fuser -m /mount_point
```

```
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 9 Take all the parallel VCS service groups offline on each of the nodes in the current subcluster:

```
# hagrps -offline grp_name -sys sys_name
```

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

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

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Veritas Storage Foundation for Sybase ASE CE to the latest release with minimal application downtime.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

To perform a rolling upgrade

- 1 Complete the preparatory steps on the first sub-cluster.
See [“Preparing to perform a rolling upgrade to SF Sybase CE 6.0.1”](#) on page 154.
- 2 Log in as superuser and mount the SF Sybase CE 6.0.1 installation media.
- 3 From root, start the installer.

```
# ./installer
```
- 4 From the menu, select `Upgrade` and from the sub menu, select `Rolling Upgrade`.
- 5 The installer suggests system names for the upgrade. Enter `Yes` to upgrade the suggested systems, or enter `No`, and then enter the name of any one system in the cluster on which you want to perform a rolling upgrade.
- 6 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type `y` to continue.

- 7 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 8 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 9 Review the end-user license agreement, and type **y** if you agree to its terms.
- 10 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you to do one of the following:

- Manually switch service groups
- Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

- 11 The installer prompts you to stop the applicable processes. Type **y** to continue. The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.
- 12 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs. When prompted, enable replication or global cluster capabilities, if required, and register the software.

The installer performs the upgrade configuration and re-starts processes.

If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.

Note: The SybaseCE service group is offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

- 13 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 14 Bring the database service group online on the subcluster that is upgraded.

```
# hagr -online sybase_group -sys sys1
# hagr -online sybase_group -sys sys2
```

- 15** Complete the preparatory steps on the nodes that you have not yet upgraded.

See [“Preparing to perform a rolling upgrade to SF Sybase CE 6.0.1”](#) on page 154.

- 16** The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

This completes phase 1 of the upgrade.

- 17** If the Sybase database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the service group online when VCS starts:

```
# haconf -makerw
# hagrps -modify sybasece AutoStart 1
# haconf -dump -makero
```

- 18** Phase 2 of the rolling upgrade begins here. This phase includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

- 19** The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.

- 20** The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.

- 21** Type **y** or **n** to help Symantec improve the automated installation.

- 22** If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

- 23** A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.

- 24** Upgrade Sybase ASE CE to the supported version.

See [“Supported database software”](#) on page 33.

For instructions, see the chapter *Upgrading Sybase ASE CE* in this document.

Performing post-upgrade tasks

This chapter includes the following topics:

- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)
- [Setting or changing the product license level](#)
- [Upgrading disk layout versions](#)
- [Upgrading CVM protocol version and VxVM disk group version](#)
- [Verifying the cluster](#)

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

See [“Performing a rolling upgrade using the installer”](#) on page 156.

To re-join the backup boot disk group

- ◆ Re-join the *backup_bootdg* disk group to the boot disk group.

```
# /etc/vx/bin/vxrootadm -Y join backup_bootdg
```

where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

See “[Performing a rolling upgrade using the installer](#)” on page 156.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vxdg` command to find the boot disk group where you are currently booted.

```
# vxdg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Setting or changing the product license level

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

After you upgrade, perform one of the following steps:

- Obtain a valid license key and run the `vxlicinst` command to add it to your system.
- Use the `vxkeyless` command to update the license keys to the keyless license model.

For more information and instructions, see the chapter *Licensing SF Sybase CE*.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 7, 8, and 9. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

See the `vxupgrade(1M)` manual page.

Support for disk layout Version 4 has been removed. You must upgrade any existing file systems with disk layout Version 4 to disk layout Version 7 or later using the `vxfsconvert` command.

See the `vxfsconvert(1M)` manual page.

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/vol1 | grep -i version
```

For more information about disk layout versions, see the *Veritas Storage Foundation Administrator's Guide*.

Upgrading CVM protocol version and VxVM disk group version

The default Cluster Volume Manager protocol version is 120.

Run the following command to verify the CVM protocol version:

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

If the protocol version is not 120, run the following command to upgrade the version:

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

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and

perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 180.

Check the existing disk group version:

```
# vxdg list dg_name | grep -i version
```

If the disk group version is not 180, run the following command on the master node to upgrade the version:

```
# vxdg -T 180 upgrade dg_name
```

Verifying the cluster

After completing the upgrade procedure, you must perform the following checks on each node of the cluster.

To verify the cluster

1 Verify that all ports are up on the cluster.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen      8ea203 membership 0123
Port b gen      8ea206 membership 0123
Port f gen      8ea21f membership 0123
Port h gen      8ea216 membership 0123
Port u gen      8ea21d membership 0123
Port v gen      8ea219 membership 0123
Port w gen      8ea21b membership 0123
Port y gen      8ea218 membership 0123
```

2 Verify all service groups and resources are online.

```
# hagr -state

hagr -state
#Group      Attribute      System      Value
binmnt      State          vcssx005   |ONLINE|
binmnt      State          vcssx012   |ONLINE|
binmnt      State          vcssx013   |ONLINE|
binmnt      State          vcssx014   |ONLINE|
cvm         State          vcssx005   |ONLINE|
cvm         State          vcssx012   |ONLINE|
cvm         State          vcssx013   |ONLINE|
cvm         State          vcssx014   |ONLINE|
sybasece   State          vcssx005   |ONLINE|
sybasece   State          vcssx012   |ONLINE|
sybasece   State          vcssx013   |ONLINE|
sybasece   State          vcssx014   |ONLINE|
```


Installation and upgrade of Sybase ASE CE

- [Chapter 17. Installing, configuring, and upgrading Sybase ASE CE](#)

Installing, configuring, and upgrading Sybase ASE CE

This chapter includes the following topics:

- [Before installing Sybase ASE CE](#)
- [Preparing for local mount point on VxFS for Sybase ASE CE binary installation](#)
- [Preparing for shared mount point on CFS for Sybase ASE CE binary installation](#)
- [Installing Sybase ASE CE software](#)
- [Preparing to create a Sybase ASE CE cluster](#)
- [Creating the Sybase ASE CE cluster](#)
- [Preparing to configure the Sybase instances under VCS control](#)
- [Configuring a Sybase ASE CE cluster under VCS control using the SF Sybase CE installer](#)
- [Upgrading Sybase ASE CE](#)

Before installing Sybase ASE CE

Before you install Sybase ASE CE, make sure that you perform the following tasks:

- Install SF Sybase CE
- Configure SF Sybase CE
- Set I/O fencing to Sybase mode

The high level flow for installing Sybase ASE CE in an SF Sybase CE environment:

- Create the Sybase user and groups. See Sybase ASE CE documentation.

- Create local or shared disk group, volume, and mount point for Sybase binary installation
- Install Sybase ASE CE
- Create a disk group, volume, and mount point for the Sybase quorum device
- Create a disk group, volume, and mount point for the Sybase datafiles
- Create the Sybase ASE CE cluster
- Configure Sybase ASE CE instances under VCS control

Preparing for local mount point on VxFS for Sybase ASE CE binary installation

The following procedure provides instructions for setting up the disk groups, volume, and mount point for installing Sybase ASE CE binaries for local mount point on VxFS.

To create the disk group, volume and mount point for Sybase binaries

- 1 Initialize the disk.

For example:

```
# vxdisksetup -i Disk_1 format=cdsdisk
```

- 2 Create a diskgroup.

For example:

```
# vxdg init sybbindg_101 Disk_1 Disk_2
```

- 3 Create a mirrored volume in the group:

```
# vxassist -g sybbindg_101 make sybbinvol  
12G layout=mirrored nmirrors=2
```

- 4 Create a VxFS file system on which to install the Sybase binaries:

```
# mkfs -t vxfs /dev/vx/rdisk/sybbindg_101/sybbinvol
```

For a binary installation on a local file system, run the command on each node.

- 5 Create the sybase home (\$SYBASE) directory on the node:

```
# mkdir /sybase
```


- 6 Mount the directory:

```
# mount -t vxfs /dev/vx/dsk/sybbindg_101/sybbinvol /sybase
```

- 7 Repeat the above steps on all other cluster nodes.
- 8 On each system, change permission of the directory to sybase.

```
# chown -R sybase:sybase /sybase
```

Preparing for shared mount point on CFS for Sybase ASE CE binary installation

The following procedure provides instructions for setting up the disk groups, volume, and mount point for installing Sybase ASE CE binaries for shared mount point on CFS.

To create the disk group, volume and mount point for Sybase binaries

- 1 Initialize the disk.

For example:

```
# vxdisksetup -i Disk_1 format=cdsdisk
```

- 2 Create a CVM diskgroup.

For example:

```
# vxdg -s init sybbindg_101 Disk_1 Disk_2
```

- 3 Create a mirrored volume in the group:

```
# vxassist -g sybbindg_101 make sybbinvol  
12G layout=mirrored nmirrors=2
```

- 4 Create a VxFS file system on which to install the Sybase binaries:

```
# mkfs -t vxfs -o largefiles /dev/vx/rdisk/sybbindg_101/sybbinvol
```

For a binary installation on a shared file system, you may run the command on any one node.

- 5 Create a Sybase ASE CHome directory (\$SYBASE) on all nodes:

```
# mkdir /sybase
```

- 6 Mount the directory:

```
# mount -t -o cluster vxfs /dev/vx/dsk/sybbindg_101/sybbinvol /sybase
```

- 7 On each system, change permission of the directory to sybase.

```
# chown -R sybase:sybase /sybase
```

Installing Sybase ASE CE software

For information on installing Sybase ASE CE software, see the Sybase ASE CE product documentation.

Requirements for the Sybase ASE CE configuration:

- Use the CFS mount points you created in the previous section for installing the binaries
See [“To create the disk group, volume and mount point for Sybase binaries”](#) on page 169.

Preparing to create a Sybase ASE CE cluster

The following procedure provides instructions for creating a file system for the quorum device.

To create the disk group, volume and mount point for a quorum device

- 1 Initialize the disk.

For example:

```
# vxdisksetup -i Disk_3 format=cdsdisk  
# vxdisksetup -i Disk_4 format=cdsdisk
```

- 2 As root user, from the CVM master, create a shared VxVM diskgroup for the quorum device.

```
# vxdg -s init quorum_101 Disk_3 Disk_4
```

- 3 As root user, from the CVM master, create a mirrored volume, *quorumvol*:

```
# vxassist -g quorum_101 make quorumvol  
1G layout=mirrored \  
nmirrors=2
```

- 4 As root user, from the CVM master, create a filesystem with the volume, *quorumvol*.

```
# mkfs -t vxfs /dev/vx/rdisk/quorum_101/quorumvol
```

- 5 On each system, create a directory, */quorum*:

```
# mkdir /quorum
```

- 6 On each system, mount */quorum*

```
# mount -t vxfs -o cluster /dev/vx/dsk/quorum_101/quorumvol  
/quorum
```

- 7 As root user, from any system, change permissions on */quorum*

```
# chown -R sybase:sybase /quorum
```

To create the disk group, volume and mount point for the datafiles

- 1 Initialize the disk.

For example:

```
# vxdisksetup -i Disk_5 format=cddisk  
# vxdisksetup -i Disk_6 format=cddisk
```

- 2 As root user, create a shared VxVM diskgroup for the datafiles.

```
# vxdg -s init sybdata_101 Disk_5 Disk_6
```

- 3 As root user, create a mirrored volume, *sybvol*:

```
# vxassist -g sybdata_101 make sybvol 1G layout=mirrored \  
nmirrors=2
```

- 4 As root user, create a filesystem with the volume, *sybvol*.

```
# mkfs -t vxfs /dev/vx/rdisk/sybdata_101/sybvol
```

- 5 On each system, create a directory, */sybdata*:

```
# mkdir /sybdata
```

- 6 On each system, mount */sybdata*

```
# mount -t vxfs -o cluster /dev/vx/dsk/sybdata_101/sybvol  
/sybdata
```
- 7 As root user, from any system, change permissions on */sybdata*

```
# chown -R sybase:sybase /sybdata
```

Creating the Sybase ASE CE cluster

For information on creating a Sybase ASE CE cluster, see the Sybase ASE CE product documentation. Follow the normal process.

Requirements for the Sybase ASE CE configuration:

- When you choose the private interconnect, set them on LLT links
- SF Sybase CE supports only one instance per node
- You can create a VCS cluster in local mode. Ignore the message "If you want to create a VCS cluster, specify "Shared" mode.", if it appears.
- Put the quorum device on the mount point created for the quorum device. See [“To create the disk group, volume and mount point for a quorum device”](#) on page 170.
- Put the datafiles on the mount point created in for the datafiles. See [“To create the disk group, volume and mount point for the datafiles”](#) on page 171.

Preparing to configure the Sybase instances under VCS control

Before putting the Sybase instances under VCS control, you may need to perform the following tasks:

- [Language settings for the Sybase agent](#)
- [Configuring Sybase for detail monitoring](#)
- [Encrypting passwords for Sybase](#)
- [About setting up detail monitoring for the agentfor Sybase](#)

Language settings for the Sybase agent

For the Veritas agent for Sybase to function with the desired locale, make sure that the Sybase installation has the correct localization files. For example, if the Sybase server requires 'LANG=en_US.UTF-8' environment variable, verify that the localization files corresponding to language 'en_US.UTF-8' are installed with Sybase.

Also, edit the file `$VCS_HOME/bin/vcsenv` to contain the following:

```
LANG=en_US.UTF-8;export LANG
```

This change affects all the agents that are configured on the nodes.

Configuring Sybase for detail monitoring

This section describes the tasks to be performed to configure a Sybase server for detail monitoring.

See [“About setting up detail monitoring for the agent for Sybase”](#) on page 175.

Note: The steps that are described here are specific to the sample script, `SqlTest.pl`, provided with the agent. If you use a custom script for detail monitoring, you must configure the Sybase database accordingly.

Perform these steps only once in a Sybase cluster.

To configure Sybase for detail monitoring

- 1 Source the `SYBASE.sh` file or `SYBASE.csh` file (depending on the user shell) to set the `$SYBASE` and `$SYBASE_ASE` environment variables.

- 2 Start the Sybase server.

```
# startserver -f ./SYBASE/$SYBASE_ASE/install/RUN_server_name
```

- 3 Start the Sybase client on any cluster node.

```
# isql -Usa -SSYBASE_SERVER_NAME
```

Enter the administrator password when prompted to do so.

- 4 Connect to the master database.

```
# use master
# go
```

5 Create a Sybase user account.

```
# sp_addlogin user_name, password  
# go
```

The detail monitor script should use this account to make transactions on the database.

6 Create a database.

```
# create database database_name  
# go
```

The detail monitor script should make transactions on this database.

7 If required, restrict the size of the log file for the database.

```
# sp_dboption database_name, "trunc log on chkpt", true  
# go
```

8 Connect to the database that is created in step 6.

```
# use database_name  
# go
```

9 Associate the user created in step 5 with the database created in step 6.

```
# sp_adduser user_name  
# go
```

10 Change the user to the one created in step 5.

```
# setuser user_name  
# go
```

11 Create a table in the database.

```
# create table table_name (lastupd datetime)  
# go
```

The detail monitor script should make transactions on this table.

If you use the SqlTest.pl for detail monitoring, make sure you create a table with a lastupd field of type datetime.

- 12 Verify the configuration by adding an initial value to the table.

```
# insert into table_name (lastupd) values (getdate())  
# go
```

- 13 Exit the database.

```
# exit
```

Encrypting passwords for Sybase

VCS provides a `vcscrypt` utility to encrypt user passwords. Encrypt passwords before specifying them for Sybase and SybaseBk resource type definition.

The `vcscrypt` utility also allows you to encrypt the agent passwords using a security key. The security key supports AES (Advanced Encryption Standard) encryption which creates a more secure password for the agent. See the *Veritas Cluster Server Administrator's Guide* for more information.

To encrypt passwords

- 1 From the path `$VCS_HOME/bin/`, run the `vcscrypt` utility.
- 2 Type the following command.

```
# vcscrypt -agent
```

The utility prompts you to enter the password twice. Enter the password and press Return.

```
Enter Password:  
Enter Again:
```

- 3 The utility encrypts the password and displays the encrypted password.
- 4 Enter this encrypted password as the value for the attribute.
Copy the encrypted password for future reference.

About setting up detail monitoring for the agent for Sybase

The Veritas agent for Sybase provides two levels of application monitoring: basic and detail. In basic monitoring, Sybase resource monitors the Sybase daemon processes to verify that they are continuously active.

In detail monitoring, the Sybase resource performs transactions on a table (provided by the user) in the database to ensure that the Sybase server functions properly. The agent uses this table for internal purposes. Symantec recommends

that you do not perform any other transaction on this table. The agent uses the script that is defined in the attribute `Monscript` of the Sybase resource. During detail monitoring, the agent executes the specified script. If the script successfully executes, the agent considers the database available. You can customize the default script according to your configurations.

To activate detail monitoring, the `LevelTwoMonitorFreq` attribute must be set to a positive integer and `User`, `UPword`, `Db`, and `Table` attributes must not be empty (""). The attribute `Monscript`, which contains the path of the detail monitor script, must also exist and must have execute permissions for the root.

Enabling detail monitoring for the agent for Sybase

Perform the following steps to enable detail monitoring on a database.

To enable detail monitoring

- 1 Make sure the Sybase server is configured for detail monitoring.
See [“Configuring Sybase for detail monitoring”](#) on page 173.
- 2 Make the VCS configuration writable.

```
# haconf -makerw
```


3 Enable detail monitoring for Sybase.

```
# hatype -modify Sybase LevelTwoMonitorFreq <value>
# hares -modify Sybase_resource User user_name
# hares -modify Sybase_resource UPword encrypted-password
# hares -modify Sybase_resource Db database_name
# hares -modify Sybase_resource Table table_name
# hares -modify Sybase_resource Monscript
"/opt/VRTSagents/ha/bin/Sybase/SqlTest.pl"
```

Note: To enable detail monitoring, the `LevelTwoMonitorFreq` attribute must be set to a positive value. You can also override the value of this attribute at the resource level.

4 Save the configuration.

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

Note: If detail monitoring is configured and the database is full, the SQL queries take considerable time to commit the results. In such a case, the monitor routine for the agent fails and attempts to fail over the service group. This issue is not encountered if detail monitoring is not configured.

Disabling detail monitoring for the agent for Sybase

1 Make the VCS configuration writable with:

```
# haconf -makerw
```

2 To disable detail monitoring for Sybase run the following command:

```
# hatype -modify Sybase LevelTwoMonitorFreq 0
```

3 Save the configuration with:

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

Configuring a Sybase ASE CE cluster under VCS control using the SF Sybase CE installer

A VCS service group is a collection of resources working together to provide application services to clients. A VCS service group typically includes multiple resources that are both hardware and software based. For example, a resource may be a physical component such as a disk or network interface card, or a software component such as Sybase or a Web server, or a configuration component such as an IP address or mounted file system.

For an example configuration file:

See [“Sample main.cf for a basic Sybase ASE CE cluster configuration under VCS control with shared mount point on CFS for Sybase binary installation”](#) on page 306.

The SF Sybase CE installer enables you to configure VCS service groups for putting a basic Sybase ASE CE cluster under VCS control. For examples of the VCS service group dependencies for SF Sybase CE see the following diagrams.

[Figure 17-1](#) displays the service group dependencies for an SF Sybase CE configuration on local disk group with VxFS.

Figure 17-1 Service group dependencies for an SF Sybase CE configuration on local disk group with VxFS


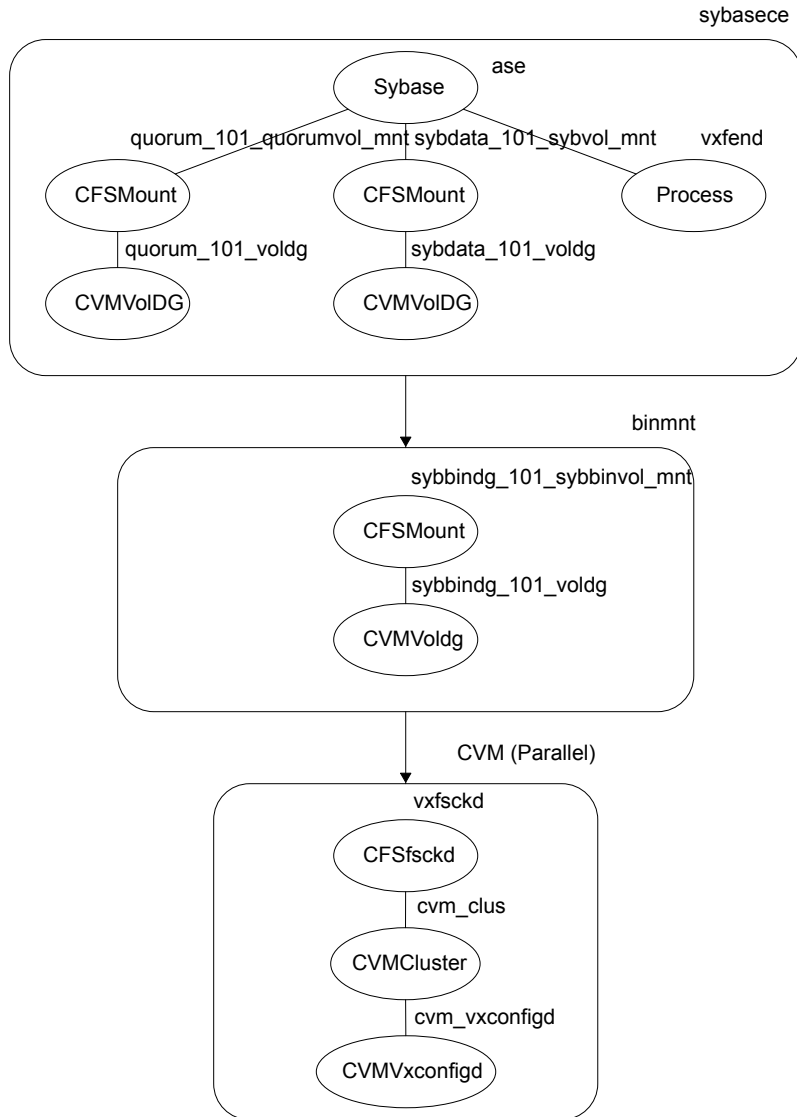


Figure 17-2 displays the service group dependencies for an SF Sybase CE configuration on shared disk group with CFS.

Figure 17-2 Service group dependencies for an SF Sybase CE configuration on shared disk group with CFS



Requirements for configuring the SF Sybase CE cluster under VCS control:

- Install SF Sybase CE.
- Configure SF Sybase CE.

- Configure I/O fencing in Sybase mode.
- Create Sybase user and group.
See Sybase documentation.
- Create a local or shared disk group, volume, and mount point for Sybase binary installation.
- Install the Sybase ASE CE software
- Create a shared disk group, volume and mount point for the Sybase ASE CE quorum device
- Create a shared disk group, volume and mount point for the Sybase ASE CE datafiles
- Create the Sybase ASE CE cluster

To put the Sybase ASE CE cluster and its resources under VCS control, the installer's configuration process will add the required resources to appropriate VCS service groups.

Table 17-1 lists the required resources for configuring Sybase ASE CE under VCS control.

Table 17-1 Required resources for configuring Sybase ASE CE under VCS control

Required resources	Example values	
Resources for the Sybase ASE CE binary installation:	Example values for shared mount point:	Example values for local mount point:
<ul style="list-style-type: none"> ■ Disk group ■ Mount point ■ Volume 	<ul style="list-style-type: none"> ■ <i>sybindg_101</i> ■ <i>/sybase</i> ■ <i>sybinvol</i> 	<ul style="list-style-type: none"> ■ <i>sybindg_101_voldg</i> ■ <i>/sybase</i> ■ <i>sybindg_101_vol</i>
Resources for the Sybase ASE CE quorum device:	Example values for shared mount point:	
<ul style="list-style-type: none"> ■ Disk group ■ Mount point ■ Volume 	<ul style="list-style-type: none"> ■ <i>quorum_101</i> ■ <i>/quorum</i> ■ <i>quorumvol</i> 	
Resources for the Sybase ASE CE datafiles:	Example values for shared mount point:	
<ul style="list-style-type: none"> ■ Disk group ■ Mount point ■ Volume 	<ul style="list-style-type: none"> ■ <i>sybdata_101</i> ■ <i>/sybdata</i> ■ <i>sybvol</i> 	

Table 17-1 Required resources for configuring Sybase ASE CE under VCS control
(continued)

Required resources	Example values
Any other CFS disk group, mount point, and volume used for Sybase ASE CE resources that are required by the Sybase ASE CE cluster	As needed
The quorum device name	/quorum/quorum.dat

Warning: You will not be able to proceed using the installer to configure the Sybase ASE CE cluster under VCS control without the items listed in [Table 17-1](#)

To configure VCS service groups for Sybase ASE CE

- 1 Log in to the installer if you are not currently logged in.
 See [“Configuring the SF Sybase CE components using the script-based installer”](#) on page 70.
- 2 When prompted to select an option from the main menu, choose the option: **Configure Sybase ASE CE Instance in VCS.**
 The installer will not be able to proceed any further unless you have the required resources available.
 See [Table 17-1](#) on page 181.
- 3 To select the type of file system where Sybase ASE CE binaries reside, choose one of the options.
 Symantec recommends CFS.
- 4 Configure the Sybase ASE CE binary installation resources under VCS control. These are the resources which were created while preparing to install Sybase ASE CE.
 See [“Preparing for shared mount point on CFS for Sybase ASE CE binary installation”](#) on page 169. for shared mount point.
 See [“Preparing for local mount point on VxFS for Sybase ASE CE binary installation”](#) on page 168. for local mount point.
 To configure the Sybase resources under VCS control:
 - To select a disk group used for Sybase ASE CE installation, choose one of the options.

Note: If you use Sybase ASE CE installation binaries on the local VxFS mount, you must specify the disk group for each node.

- To select the volume used for Sybase ASE CE installation, choose one of the options.
 - Enter the mount point for the selected volume.
- 5** The quorum device resources must be added into the resource group if it is under a different CFS than the Sybase database installation. These resources were created while preparing for a Sybase ASE CE cluster.

See [“Preparing to create a Sybase ASE CE cluster”](#) on page 170.

To configure the quorum device under VCS control:

- Enter **y** if the quorum device is under a different CFS than the Sybase database resources you have configured in the previous step, otherwise enter **n**.
 - If you entered **y**, select a disk group for the quorum device.
 - Select a volume for the quorum device.
 - Enter **y** if there is a CFS on the volume you selected, otherwise enter **n**. The quorum device can use either a volume which you have selected directly or a file under CFS created on the selected volume.
 - Enter the mount point for the volume.
- 6** If there are any other disk groups, volumes, or mount points used for the Sybase ASE CE cluster, such as other database files, for instance master, system, etc., which are using a different CFS, they must also be put under VCS control.

To add other disk groups, volumes, and mount points to the resource group, enter **y** when prompted, otherwise enter **n**.

- 7** Verify the disk groups, volumes and mount points information when prompted.
- 8** To configure the Sybase ASE CE resources:
- Enter the Sybase instance on ASE1 and ASE2 when prompted.
 - Enter the Sybase UNIX user name.
 - Enter Sybase home directory, where the Sybase binaries reside.
 - Enter Sybase version.

- If required, enter the username and password for the Admin user. The default username is 'sa', password is".
- Enter the Sybase quorum device information.
During configuration of Sybase instance under VCS control, if the quorum device is on CFS and is not mounted, the following warning message appears on the installer screen:

```
Error: CPI WARNING V-9-0-0 The quorum file /quorum/quorum.dat cannot be .  
This may be due to a file system not being mounted.
```

This message may be safely ignored. The resource will be onlined and available when the service is completed.

- Verify the Sybase configuration information by entering **y**, otherwise enter **n**. For example:

```
Sybase configuration information verification:
```

```
Sybase Server on sys1: ASE1  
Sybase Server on sys2: ASE2  
Sybase UNIX user name: sybase  
Sybase home directory where sybase binaries reside: /sybase  
Sybase version: 15  
Sybase sa: sa  
Passwords are not displayed  
Sybase quorum: /quorum/quorum.dat
```

Once you confirm the information is correct, the installer configures and onlines the VCS service groups for Sybase ASE CE. This completes the configuration of Sybase ASE CE under VCS control.

- Note the location of the configuration log files for future reference.
- 9 To verify the service groups have been created and are available online, enter:

```
# hagr -state
```

```
hagr -state  
#Group      Attribute      System      Value  
binmnt      State          system1     |ONLINE|  
binmnt      State          system2     |ONLINE|  
cvm         State          system1     |ONLINE|  
cvm         State          system2     |ONLINE|  
sybasece    State          system1     |ONLINE|  
sybasece    State          system2     |ONLINE|
```


Upgrading Sybase ASE CE

SF Sybase CE supports Sybase ASE CE 15.5 only at the time of publication.

For information on upgrading Sybase ASE CE software, see the Sybase ASE CE product documentation:

See infocenter.sybase.com.

Automated installation using response files

- [Chapter 18. About response files](#)
- [Chapter 19. Installing and configuring SF Sybase CE using a response file](#)
- [Chapter 20. Performing an automated I/O fencing configuration using response files](#)
- [Chapter 21. Configuring a cluster under VCS control using a response file](#)

About response files

This chapter includes the following topics:

- [About response files](#)
- [Response file syntax](#)
- [Guidelines for creating the SF Sybase CE response file](#)
- [Installation scenarios for response files](#)

About response files

Use response files to standardize and automate installations on multiple clusters. You can perform the following installation activities using a response file:

- Installing and configuring SF Sybase CE
- Uninstalling SF Sybase CE

[Table 18-1](#) lists the various options available for creating or obtaining a response file.

Table 18-1 Options for obtaining a response file

Option	Description
Create a response file	Create a response file based on the sample response file. See “Sample response files for installing and configuring SF Sybase CE” on page 207.

Table 18-1 Options for obtaining a response file (*continued*)

Option	Description
Reuse or customize the response files generated by an installation	<p>The response file generated by the installer is located in the following directory:</p> <pre data-bbox="556 406 1198 458">/opt/VRTS/install/logs/installsfbasece<version>-\ installernumber/installsfbasece<version>-installernumber.res</pre> <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 SF Sybase CE configuration option.</p>

At the end of the SF Sybase CE installation, the following files are created:

- A log file that contains executed system commands and output.
- A summary file that contains the output of the installation scripts.
- Response files to be used with the `-responsefile` option of the installer.

Note: The SF Sybase CE 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*.

Response file syntax

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

For example,

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

or, in the case of an integer value:

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

or, in the case of a list:

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

Guidelines for creating the SF Sybase CE response file

This section provides guidelines for creating the SF Sybase CE response file.

1. Create a response file using one of the available options.
 For various options on creating or obtaining an SF Sybase CE response file:
 See [“About response files”](#) on page 189.
2. Set the following master values to 1 to enable SF Sybase CE installation and configuration.

Note: The master settings must be set to 1 to enable the installer to read dependent variable definitions. For example, if the value `$CFG{opt}{install}` is not set to 1, the other dependent installation values in the response file will be disregarded. This is true for any master setting.

The following is the list of master values that must be set for installing and configuring SF Sybase CE.

```
Installing SF Sybase CE  $CFG{opt}{install}=1;
                        $CFG{opt}{installallpkgs}=1;
```

```
Configuring SF Sybase CE $CFG{opt}{configure}=1;
```

3. Now, set the appropriate value in the dependent variable definitions for installing and configuring SF Sybase CE.

The set of minimum definitions for a successful installation and configuration is as follows:

```
$CFG{accepteula}=1;
$CFG{config_cfs}=1;
$CFG{lltoverudp}="0";
$CFG{opt}{configure}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SFSYBASECE60";
$CFG{sfsybasece}{menu}=1;
$CFG{systems}=[ qw(sys1 sys2) ];
```

```
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=2455;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_lltlink1}{sys1}="eth1";
$CFG{vcs_lltlink1}{sys2}="eth1";
$CFG{vcs_lltlink2}{sys1}="eth2";
$CFG{vcs_lltlink2}{sys2}="eth2";
$CFG{vcs_userenpw}=[ qw(gmnEmhMjnInnLvnHmk) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];

1;
```

You can add more variable definitions, as required.

Installation scenarios for response files

The chapters in this section cover the following installation scenarios using response files:

- Installing and configuring SF Sybase CE
See [“Installing and configuring SF Sybase CE”](#) on page 193.
- Configuring a SF Sybase CE instance in VCS
See [“Configuring a Sybase cluster under VCS control with a response file ”](#) on page 215.
- Configuring I/O fencing for SF Sybase CE with a response file
See [“Configuring I/O fencing using response files”](#) on page 209.

Installing and configuring SF Sybase CE using a response file

This chapter includes the following topics:

- [Installing and configuring SF Sybase CE](#)
- [Response file variables to install Veritas Storage Foundation for Sybase ASE CE](#)
- [Response file variables to configure Veritas Storage Foundation for Sybase ASE CE](#)
- [Sample response files for installing and configuring SF Sybase CE](#)

Installing and configuring SF Sybase CE

You can create a single response file or separate response files for installing and configuring SF Sybase CE.

The installer performs the following tasks:

- Installs SF Sybase CE.
- Configures SF Sybase CE.

The following sample procedure uses a single response file for installing and configuring SF Sybase CE.

To install and configure SF Sybase CE using response files

- 1 Make sure that the systems meet the installation requirements.
See [“Hardware requirements”](#) on page 32.
- 2 Complete the preparatory steps before starting the installation.
For instructions, see the chapter "Preparing to install and configure SF Sybase CE" 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 189.

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

For guidelines on creating a response file:

See [“Guidelines for creating the SF Sybase CE response file”](#) on page 191.

For a sample response file:

See [“Sample response files for installing and configuring SF Sybase CE”](#) on page 207.

See [“Sample response files for installing and configuring SF Sybase CE”](#) on page 207.

- 4 Mount the product disc and navigate to the product directory that contains the installation program.
- 5 Start the installation and configuration:

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

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

6 Configure I/O fencing.

Note: Before you configure I/O fencing, make sure that you complete the required pre-configuration tasks.

For instructions on configuring I/O fencing using a response file, see the chapter *Configuring I/O fencing using a response file* in this document.

7 Complete the SF Sybase CE post-installation tasks.

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

Response file variables to install Veritas Storage Foundation for Sybase ASE CE

[Table 19-1](#) lists the response file variables that you can define to install SF Sybase CE.

Table 19-1 Response file variables for installing SF Sybase CE

Variable	Description
CFG{opt}{install}	<p>Installs SF Sybase CE RPMs. Configuration can be performed at a later time using the <code>-configure</code> option.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	<p>Instructs the installer to install SF Sybase CE RPMs based on the variable that has the value set to 1:</p> <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> <p>List or scalar: scalar</p> <p>Optional or required: required</p>

Table 19-1 Response file variables for installing SF Sybase CE (*continued*)

Variable	Description
CFG{accepteula}	<p>Specifies whether you agree with the EULA.pdf file on the media.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>
CFG{opt}{vxkeyless}	<p>Installs the product with keyless license.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{license}	<p>Installs the product with permanent license.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{keys}{hostname}	<p>List of keys to be registered on the system if the variable \$CFG{opt}{vxkeyless} is set to 0 or if the variable \$CFG{opt}{licence} is set to 1.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
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{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>

Table 19-1 Response file variables for installing SF Sybase CE (*continued*)

Variable	Description
CFG{opt}{pkgpath}	<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>
CFG{opt}{tmppath}	<p>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.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{rsh}	<p>Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{logpath}	<p>Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{prodmode}	<p>List of modes for product</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>

Response file variables to configure Veritas Storage Foundation for Sybase ASE CE

[Table 19-2](#) lists the response file variables that you can define to configure SF Sybase CE.

Table 19-2 Response file variables specific to configuring Veritas Storage Foundation for Sybase ASE CE

Variable	List or Scalar	Description
\$CFG{config_cfs}	Scalar	Performs the Cluster File System configuration for SF Sybase CE (Required) Set the value to 1 to configure Cluster File System for SF Sybase CE
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (Required) Set the value to 1 to configure SF Sybase CE.
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. The value is SFSYBASECE60 for SFSYBASECE (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)

Table 19-2 Response file variables specific to configuring Veritas Storage Foundation for Sybase ASE CE (*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 a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtpsev), the SNMP trap notification (snmpport, snmpcons, and snmpsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

[Table 19-3](#) lists the response file variables that specify the required information to configure a basic SF Sybase CE cluster.

Table 19-3 Response file variables specific to configuring a basic SF Sybase CE cluster

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

Table 19-3 Response file variables specific to configuring a basic SF Sybase CE 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 19-4 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 19-4 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. You must enclose the system name within double quotes. (Required)

Table 19-4 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 19-5 lists the response file variables that specify the required information to configure LLT over UDP.

Table 19-5 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	<p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p>
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 19-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 19-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 19-6 lists the response file variables that specify the required information to configure virtual IP for SF Sybase CE cluster.

Table 19-6 Response file variables specific to configuring virtual IP for SF Sybase CE 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 19-7 lists the response file variables that specify the required information to configure the SF Sybase CE cluster in secure mode.

Table 19-7 Response file variables specific to configuring SF Sybase CE cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonnode}	Scalar	Specifies that the securityonnode option is being used.
CFG{securityonnode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. <ul style="list-style-type: none"> ■ 1—Configure the first node ■ 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.
CFG{opt}{fips}	Scalar	Specifies that the FIPS option is being used.
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

Table 19-8 lists the response file variables that specify the required information to configure VCS users.

Table 19-8 Response file variables specific to configuring VCS users

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

Table 19-8 Response file variables specific to configuring VCS users (*continued*)

Variable	List or Scalar	Description
CFG{vcs_username}	List	List of names of VCS users (Optional)
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 19-9](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 19-9 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table 19-10 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 19-10 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

Table 19-11 lists the response file variables that specify the required information to configure SF Sybase CE global clusters.

Table 19-11 Response file variables specific to configuring SF Sybase CE 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)

Table 19-11 Response file variables specific to configuring SF Sybase CE global clusters (*continued*)

Variable	List or Scalar	Description
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response files for installing and configuring SF Sybase CE

The following sample response file installs and configures SF Sybase CE on two nodes, sys1 and sys2.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{config_cfs}=1;
$CFG{fencingenabled}=0;
$CFG{lltoverudp}=0;
$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vr}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SFSYBASECE60";
$CFG{sfsybasece}{menu}=1;
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=24731;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_lltlink1}{sys1}="eth1";
$CFG{vcs_lltlink1}{sys2}="eth1";
$CFG{vcs_lltlink2}{sys1}="eth2";
$CFG{vcs_lltlink2}{sys2}="eth2";
$CFG{vcs_userenpw}=[ qw(JqrJqlQnrMrrPzrLqo) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];
```

```
1;
```

The following sample response file only installs SF Sybase CE on two nodes, sys1 and sys2.

```
our %CFG;  
  
$CFG{accepteula}=1;  
$CFG{opt}{gco}=1;  
$CFG{opt}{install}=1;  
$CFG{opt}{installrecpkgs}=1;  
$CFG{opt}{vr}=1;  
$CFG{opt}{vxkeyless}=1;  
$CFG{prod}="SFSYBASECE60";  
$CFG{systems}=[ qw(sys1 sys2) ];
```

```
1;
```

The following sample response file only configures CFS on two nodes, sys1 and sys2.

```
our %CFG;  
  
$CFG{config_cfs}=1;  
$CFG{fencingenabled}=0;  
$CFG{lltoverudp}=0;  
$CFG{opt}{configure}=1;  
$CFG{prod}="SFSYBASECE60";  
$CFG{sfsybasece}{menu}=1;  
$CFG{systems}=[ qw(sys1 sys2) ];  
$CFG{vcs_allowcomms}=1;  
$CFG{vcs_clusterid}=60037;  
$CFG{vcs_clustername}="clus1";  
$CFG{vcs_lltlink1}{sys1}="eth1";  
$CFG{vcs_lltlink1}{sys2}="eth1";  
$CFG{vcs_lltlink2}{sys1}="eth2";  
$CFG{vcs_lltlink2}{sys2}="eth2";  
$CFG{vcs_userenpw}=[ qw(bMNFmHmJNiNNlVNhMK) ];  
$CFG{vcs_username}=[ qw(admin) ];  
$CFG{vcs_userpriv}=[ qw(Administrators) ];
```

```
1;
```


Performing an automated I/O fencing configuration using response files

This chapter includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for SF Sybase CE.

To configure I/O fencing using response files

- 1 Make sure that SF Sybase CE is configured.
- 2 Make sure you have completed the preparatory tasks.
See [“About planning to configure I/O fencing”](#) on page 38.
- 3 Copy the response file to one of the cluster systems where you want to configure I/O fencing.
See [“Sample response file for configuring disk-based I/O fencing”](#) on page 213.
- 4 Edit the values of the response file variables as necessary.
See [“Response file variables to configure disk-based I/O fencing”](#) on page 210.

- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfbasece<version>
-responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

See “[About the Veritas installer](#)” on page 59.

Response file variables to configure disk-based I/O fencing

Table 20-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for SF Sybase CE.

Table 20-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. <ul style="list-style-type: none"> ■ 2—Sybase Mode fencing ■ 3—Fencing migration when the cluster is online (Required)
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the I/O fencing mechanism. This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the <code>fencing_scsi3_disk_policy</code> variable and either the <code>fencing_dgname</code> variable or the <code>fencing_newdgs_disks</code> variable. (Optional)

Table 20-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{fencing_dgname}	Scalar	<p>Specifies the disk group for I/O fencing. (Optional)</p> <p>Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.</p>
CFG{fencing_newdg_disks}	List	<p>Specifies the disks to use to create a new disk group for I/O fencing. (Optional)</p> <p>Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.</p>

Table 20-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{fencing_cpagent_monitor_freq}	Scalar	<p>Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution.</p> <p>Note: Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidentally deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If LevelTwoMonitorFreq attribute is not set, the agent will not monitor any changes to the Coordinator Disk Group. 0 means not to monitor the Coordinator Disk Group constitution.</p>
CFG {fencing_config_cpagent}	Scalar	<p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p>
CFG {fencing_cpagentgrp}	Scalar	<p>Name of the service group which will have the Coordination Point agent resource as part of it.</p> <p>Note: This field is obsolete if the fencing_config_cpagent field is given a value of '0'.</p>

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 210.

```
#
# Configuration Values:
#
our %CFG;

$CFG{fencing_config_cpagent}="0";
$CFG{fencing_dgname}="fendg";
$CFG{fencing_newdg_disks}=[ qw(emc_clariion0_236
    emc_clariion0_237 emc_clariion0_238) ];
$CFG{fencing_option}=2;
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{opt}{configure}=1;
$CFG{prod}="SFSYBASECE601";
$CFG{sfsybasece}{menu}=2;
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=8950;
$CFG{vcs_clustname}="clus1";

1;
```

214 | Performing an automated I/O fencing configuration using response files
Sample response file for configuring disk-based I/O fencing

Configuring a cluster under VCS control using a response file

This chapter includes the following topics:

- [Configuring a Sybase cluster under VCS control with a response file](#)
- [Response file variables to configure SF Sybase CE in VCS](#)

Configuring a Sybase cluster under VCS control with a response file

Observe the following prerequisites prior to configuring a Sybase cluster under VCS with a response file:

- SF Sybase CE must be installed and configured on the system.
- Sybase must be installed.
- The Sybase cluster must already be created.

To configure a Sybase cluster under VCS using a response file

- ◆ Use the configuration response file to configure the product:

```
# installsfbasece -responsefile /opt/VRTS/install/logs/\
installsfbasece-installernumber/installsfbasece-installer\
number.response
```

The following sample response file configures SF Sybase CE under VCS control.

```

our %CFG;

$CFG{opt}{configure}=1;
$CFG{prod}="SFSYBASECE60";
$CFG{sfsybasece}{ase_home}="/sybase_home";
$CFG{sfsybasece}{ase_owner}="sybase";
$CFG{sfsybasece}{ase_quorum}="/qrmnt/newqrm1";
$CFG{sfsybasece}{ase_sa}="sa";
$CFG{sfsybasece}{ase_server}{vcslx003}{SERVER}="inst1";
$CFG{sfsybasece}{ase_server}{vcslx004}{SERVER}="inst2";
$CFG{sfsybasece}{ase_server}{vcslx005}{SERVER}="inst3";
$CFG{sfsybasece}{ase_server}{vcslx006}{SERVER}="inst4";
$CFG{sfsybasece}{ase_version}=15;
$CFG{sfsybasece}{menu}=3;
$CFG{sfsybasece}{storage_resource}{qrmhg304}{vol1}{mount}="/qrmnt";
$CFG{sfsybasece}{storage_resource}{qrmhg304}{vol1}{usage}="quorum device";
$CFG{sfsybasece}{storage_resource}{sybdatdg304}{vol1}{mount}="/datamnt";
$CFG{sfsybasece}{storage_resource}{sybdatdg304}{vol1}{usage}="database devices";
$CFG{sfsybasece}{storage_resource}{sybhomedg304}{vol1}{mount}="/sybase_home";
$CFG{sfsybasece}{storage_resource}{sybhomedg304}{vol1}{usage}="sybase installation";
$CFG{sybase_location}=1;
$CFG{systems}=[ qw(vcslx003 vcslx004 vcslx005 vcslx006) ];

1;

```

Response file variables to configure SF Sybase CE in VCS

Table 21-1 lists the response file variables that you can define to configure SF Sybase CE in VCS.

Table 21-1 Response file variables specific to configuring SF Sybase CE in VCS

Variable	List or Scalar	Description
CFG{sfsybasece}{ase_home}	Scalar	Defines the SF Sybase CE home directory.
CFG{sfsybasece}{ase_owner}	Scalar	Defines the SF Sybase CE owner name.
CFG{sfsybasece}{ase_quorum}	Scalar	Defines the SF Sybase CE quorum device.

Table 21-1 Response file variables specific to configuring SF Sybase CE in VCS
(continued)

Variable	List or Scalar	Description
CFG{sfsybasece}{ase_sa}	Scalar	Defines the SF Sybase CE administrator name.
CFG{sfsybasece}{ase_server} {system}{SERVER}	Scalar	Defines the SF Sybase CE instance name on sles95243.
CFG{sfsybasece}{ase_version}	Scalar	Defines the SF Sybase CE version.
CFG{sfsybasece}{menu}=3	Scalar	Option for configuring SF Sybase CE under VCS.
CFG{sfsybasece}{storage_resource} {database_dontuse}{dbname}	Scalar	Lists the SF Sybase CE database devices that reside on the disgroupname diskgroup and the volumename volume.
CFG{sfsybasece}{storage_resource} {master_dontuse}{mastervol} {usage}	Scalar	Lists the SF Sybase CE database devices that reside on the master_dontuse diskgroup and the mastervol volume.
CFG{sfsybasece}{storage_resource} {proc_dontuse}{proc01vol} {mount}	Scalar	Specifies the mount point for the proc_dontuse database device.
CFG{sfsybasece}{storage_resource} {proc_dontuse}{proc01vol} {usage}	Scalar	Lists the SF Sybase CE database devices that reside on the database_dontuse diskgroup and the database01vol volume.
CFG{sfsybasece}{storage_resource} {quorum_dontuse}{quorumvol} {usage}	Scalar	Lists the SF Sybase CE quorum devices that reside on the quorum_dontuse diskgroup and the quorumvol volume.
CFG{sfsybasece}{storage_resource} {sybase1_dontuse}{sybasevol} {mount}="/opt/sybase"	Scalar	Specifies the SF Sybase CE installation location under /opt/sybase.

Table 21-1 Response file variables specific to configuring SF Sybase CE in VCS
(continued)

Variable	List or Scalar	Description
CFG{sfSybasece}{storage_resource} {sybase1_dontuse}{sybasevol} {usage}="sybase installation"	Scalar	Specifies the SF Sybase CE installation location that resides on the sybase1_dontuse diskgroup and the sybasevol volume.
CFG{sybase_location}	Scalar	Specifies the SF Sybase CE location type. <ul style="list-style-type: none"> ■ 1—Location on CFS ■ 2—Location on a local VxFS file system.

Adding and removing nodes

- [Chapter 22. Adding a node to SF Sybase CE clusters](#)
- [Chapter 23. Removing a node from SF Sybase CE clusters](#)

Adding a node to SF Sybase CE clusters

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Adding the node to a cluster manually](#)
- [Adding a node to a cluster using the SF Sybase CE installer](#)
- [Adding the new instance to the Sybase ASE CE cluster](#)

About adding a node to a cluster

After you install SF Sybase CE and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 4 nodes.

You can add a node:

- Using the product installer
- Manually

The following table provides a summary of the tasks required to add a node to an existing SF Sybase CE cluster.

Table 22-1 Tasks for adding a node to a cluster

Step	Description
Complete the prerequisites and preparatory tasks before adding a node to the cluster.	See “Before adding a node to a cluster” on page 222.
Add a new node to the cluster.	See “Adding a node to a cluster using the SF Sybase CE installer” on page 234. See “Adding the node to a cluster manually” on page 225.
Complete the configuration of the new node after adding it to the cluster.	See “Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node” on page 231.
Prepare the new node for installing Sybase.	See “Adding the new instance to the Sybase ASE CE cluster” on page 237.
Add the node to Sybase.	See “Adding the new instance to the Sybase ASE CE cluster” on page 237.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing SF Sybase CE cluster, perform the required preparations.

- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

- 1 Review hardware and software requirements for SF Sybase CE.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 Verify the existing cluster is an SF Sybase CE cluster and that SF Sybase CE is running on the cluster.

- 4 If the cluster is upgraded from the previous version, you must check the cluster protocol version to make sure it has the same version as the node to be added. If there is a protocol mismatch, the node is unable to join the existing cluster.

Check the cluster protocol version using:


```
# vxctl protocolversion
Cluster running at protocol 120
```

- 5 If the cluster protocol on the master node is below 120, upgrade it using:

```
# vxctl upgrade [version]
```

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 22-1](#).

Figure 22-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SF Sybase CE 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 two-node 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 22-1 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.

Complete the following preparatory steps on the new node before you add it to an existing SF Sybase CE cluster.

To prepare the new node

1 Verify that the new node meets installation requirements.

```
# ./installsfsybasece -precheck
```

2 Install SF Sybase CE on the new system. Make sure all the VRTS RPMs available on the existing nodes are also available on the new node.

```
# cd /opt/VRTS/install  
  
# ./installsfsybasece<version>
```

Where *<version>* is the specific release version.

Do not configure SF Sybase CE when prompted.

3 You can restart the new node after installation is complete. Configure the new node using the configuration from the existing cluster nodes.

Adding the node to a cluster manually

Perform this procedure after you install SF Sybase CE only if you plan to add the node to the cluster manually.

Table 22-2 Procedures for adding a node to a cluster manually

Step	Description
Start the Veritas Volume Manager (VxVM) on the new node.	See “Starting Veritas Volume Manager (VxVM) on the new node” on page 225.
Configure the cluster processes on the new node.	See “Configuring cluster processes on the new node” on page 226.
Configure fencing for the new node to match the fencing configuration on the existing cluster.	See “Starting fencing on the new node” on page 231.
Start VCS.	See “To start VCS on the new node” on page 233.
Configure CVM and CFS.	See “Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node” on page 231.
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 233.

Starting Veritas Volume Manager (VxVM) on the new node

Veritas Volume Manager (VxVM) 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 `installsfbasece` program.

To start VxVM on the new node

- 1 To start VxVM on the new node, use the vxinstall utility:

```
# vxinstall
```

- 2 Enter **n** when prompted to set up a system wide disk group for the system. The installation completes.

- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

Configuring cluster processes on the new node

Perform the steps in the following procedure to configure cluster processes on the new node.

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using `vi` or another text editor, add the line for the new node to the file. The file resembles:

```
0 sys1
1 sys2
2 sys5
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.
- 3 Create an `/etc/llttab` file on the new system. For example:

```
set-node system3
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.

- 4 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/llt
/etc/sysconfig/gab
/etc/sysconfig/vcs
```

- 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:

```
system3
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -rsh -clus -copy \
-from_sys sys1 -to_sys system3
```

- 9 Start the LLT and GAB drivers on the new node:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vxgln start
```

- 10 On the new node, verify GAB port membership:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen df204 membership 01
```

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 22-3](#) uses the following information for the following command examples.

Table 22-3 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
sys5	sys5.nodes.example.com	The new node that you are adding to the cluster.

Configuring the authentication broker on node sys5

To configure the authentication broker on node sys5

- 1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvcs/vcsauth/bkup  
# cd /tmp; gunzip -c /opt/VRTSvcs/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvcs/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

- 3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \  
./broker_setup.sh/tmp/eat_setup  
  
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f  
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \  
/VRTSatlocal.conf -b 'Security\Authentication \  
\Authentication Broker' -k UpdatedDebugLogFileName \  
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

- 4 Copy the broker credentials from one node in the cluster to sys5 by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/  
  
# ls  
  
CMDSERVER  CPSADM  HAD  VCS_SERVICES  WAC
```

- 5 Import the `VCS_SERVICES` domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \  
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \  
/VCS_SERVICES -p password
```

- 6 Import the credentials for `HAD`, `CMDSERVER`, `CPSADM`, and `WAC`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \  
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \  
/HAD -p password
```

- 7 Start the `vcsauthserver` process on `sys5`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvcs/vcsauth/data/CLIENT
# mkdir /var/VRTSvcs/vcsauth/data/TRUST
# export EAT_DATA_DIR='/var/VRTSvcs/vcsauth/data/TRUST'
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \
localhost:14149 -s high
```

9 Create the /etc/VRTSvcs/conf/config/.secure file:

```
# touch /etc/VRTSvcs/conf/config/.secure
```

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

1 For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/vxfen
/etc/vxfendg
/etc/vxfenmode
```

2 Start fencing on the new node:

3 On the new node, verify that the GAB port memberships are a and b:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 57c004 membership 012
Port b gen 57c019 membership 012
```

Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node

Modify the existing cluster configuration to configure Cluster Volume Manager (CVM) and Cluster File System (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 system3
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagrpl -modify cvm SystemList -add system3 2
# hagrpl -modify cvm AutoStartList -add system3
# hares -modify cvm_clus CVMNodeId -add system3 2
# haconf -dump -makero
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```


- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \  
system3:/etc/VRTSvcs/conf/config/main.cf  
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \  
system3:/etc/VRTSvcs/conf/config/CFSTypes.cf  
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \  
system3:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

- 7 The `/etc/vx/tunefstab` file sets non-default tunables for local-mounted and cluster-mounted file systems.

If you have configured a `/etc/vx/tunefstab` file to tune cluster-mounted file systems on any of the existing cluster nodes, you may want the new node to adopt some or all of the same tunables.

To adopt some or all tunables, review the contents of the file, and copy either the file, or the portions desired, into the `/etc/vx/tunefstab` file on the new cluster node.

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:

```
# hstart
```

VCS brings the CVM and CFS groups online.

- 2 Verify that the CVM and CFS groups are online:

```
# hagrps -state
```

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

- 1 On an existing node, for example sys1, write-enable the configuration:

```
# haconf -makerw
```

- 2 Add the node system3 to the existing ClusterService group.

```
# hagrps -modify ClusterService SystemList -add system3 2
```

```
# hagrps -modify ClusterService AutoStartList -add system3
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device eth0 -sys system3
```

```
# hares -modify gconic Device eth0 -sys system3
```

- 4 Save the configuration by running the following command from any node.

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

Adding a node to a cluster using the SF Sybase CE installer

You can add a node to a cluster using the `-addnode` option with the SF Sybase CE installer.

The SF Sybase CE installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and RPMs installed but not configured on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.

- Creates the following files on the new node:

```
/etc/llttab
```

```
/etc/VRTSvcs/conf/sysname
```

- Copies the following files on the new node:

```
/etc/llthosts
```

```
/etc/gabtab
```

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

- Copies the following files from the existing cluster to the new node:

```
/etc/vxfenmode
/etc/vxfendg
/etc/vx/.uuids/clusuuid
/etc/sysconfig/llt
/etc/sysconfig/gab
/etc/sysconfig/vxfen
```

- Configures fencing.
- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

- Starts SF Sybase CE processes and configures CVM and CFS on the new node.

At the end of the process, the new node joins the SF Sybase CE cluster.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SF Sybase CE installer with the -addnode option.

```
# cd /opt/VRTS/install
# ./installsfbasece<version> -addnode
```

Where *<version>* is the specific release version.

See [“About the Veritas installer”](#) on page 59.

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing SF Sybase CE cluster.
The installer uses the node information to identify the existing cluster.

```
Enter one node of the SFSYBASECE cluster to which
you would like to add one or more new nodes: sys1
```

- 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: system3
```

Confirm if the installer prompts if you want to add the node to the cluster.

The installer checks the installed products and RPMs 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 system3: [b,q,?] eth1
```

```
Enter the NIC for the second private heartbeat  
link on system3: [b,q,?] eth2
```

Note: At least two private heartbeat links must be configured for high availability of the cluster.

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 8 Review and confirm the information.

- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

Enter the NIC for VCS to use on system3: **eth3**

The installer starts the SF Sybase CE processes and configures CVM and CFS on the new node. The new node is now part of the cluster.

SF Sybase CE is configured on the cluster. Do you want to configure it on the new node(s)? [y,n,q] (y) n

To add the new node into the Sybase ASE CE cluster and database:

See [“Adding the new instance to the Sybase ASE CE cluster”](#) on page 237.

- 10 Configure the following service groups manually to include the new node in the VCS configuration:
- The *binmnt* group which contains the Sybase binaries
 - The *Sybase* group which contains:
 - The new instance on the added node
 - The database mounts where the database resides
 - The quorum mounts where the quorum device resides.
 - See [“Adding the new instance to the Sybase ASE CE cluster”](#) on page 237.
- 11 Confirm that the new node has joined the SF Sybase CE cluster using `lltstat -n` and `gabconfig -a` commands.

If the new node has not joined the cluster, verify if it has been added to the SystemList.

Adding the new instance to the Sybase ASE CE cluster

To add a new Sybase ASE CE instance to the cluster you must complete the following tasks:

- [Creating Sybase user and groups](#)
- [Preparing the mount point for Sybase resources on the new node](#)
- [Adding a new Sybase ASE CE instance to the Sybase ASE CE cluster](#)
- [Bringing the new Sybase ASE CE instance under VCS control](#)

Creating Sybase user and groups

To prepare the new node for a Sybase ASE CE instance, create the Sybase user and groups.

See your Sybase ASE CE documentation.

Preparing the mount point for Sybase resources on the new node

To prepare the new node for installing Sybase, you must prepare mount points on the new node for Sybase binaries, quorum device, and datafiles.

See [“Preparing for shared mount point on CFS for Sybase ASE CE binary installation”](#) on page 169.

See [“Preparing to create a Sybase ASE CE cluster”](#) on page 170.

Create the mount point for the file system with the Sybase binary files.

For example:

```
# mkdir -p /sybase  
# chown -R sybase:sybase /sybase
```

Create the mount point for the file system with the Sybase quorum device.

For example:

```
# mkdir -p /quorum  
# chown -R sybase:sybase /quorum
```

Create the mount point for the file system with the Sybase datafiles.

For example:

```
# mkdir -p /sybdata  
# chown -R sybase:sybase /sybdata
```

Adding a new Sybase ASE CE instance to the Sybase ASE CE cluster

For a CFS shared installation of Sybase ASE CE binaries, the new Sybase ASE CE instance on the new node can share the existing cluster's Sybase ASE CE binaries.

For a local VxFS installation of Sybase ASE CE binaries, you need to create diskgroups for binaries and install Sybase ASE CE binaries on the new node.

To configure the new node

- 1 From an existing node in the cluster, write enable the cluster configuration:

```
# haconf -makerw
```

- 2 In case of Sybase binaries on CFS, add the new node to the VCS service group for the Sybase binaries:

```
# hagrps -modify binmnt SystemList -add sys5 2
```

```
# hagrps -modify binmnt AutoStartList -add sys5
```

- 3 In case of Sybase binaries on local VxFS, add the name of the DiskGroup for the new node.

```
# hares -modify sybase_install_dg DiskGroup
sybase_new_diskgroup -sys sys5
```

```
# hares -modify sybase_install_mnt BlockDevice
/dev/vx/dsk/sybase_new_diskgroup/sybase_new_volume -sys sys5
```

```
# hares -modify sybase_install_vol DiskGroup
sybase_new_diskgroup -sys sys5
```

```
# hares -modify sybase_install_vol Volume
sybase_new_volume -sys sys5
```

- 4 Save the configuration changes.

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

- 5 Bring the VCS group for Sybase binaries group online on the new node:

```
# hagrps -online binmnt -sys sys5
```

To add the new node to the Sybase ASE CE cluster

- ◆ Follow the procedures in your Sybase ASE CE documentation.

Bringing the new Sybase ASE CE instance under VCS control

After adding a new instance to the Sybase ASE CE cluster you must bring it under VCS control.

To configure the new instance under VCS control

- 1 From an existing node in the cluster, write enable the cluster configuration:

```
# haconf -makerw
```

- 2 Add the node to the VCS service group for managing Sybase resources:

```
# hagrps -modify sybasece SystemList -add sys5 2
```

```
# hagrps -modify sybasece AutoStartList -add sys5
```

- 3 Add the new instance to the VCS resource used to manage Sybase instances:

```
# hares -modify ase Server ase3 -sys sys5
```

- 4 Save the configuration changes.

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

- 5 Bring the Sybase service group online on the new node:

```
# hagrps -online sybasece -sys sys5
```

Note: Before you bring the Sybase service group online, make sure you have manually created the Run file for the added instance on the added node, with appropriate instance information.

This completes the addition of the new node to the cluster. You now have a three node cluster.

Removing a node from SF Sybase CE clusters

This chapter includes the following topics:

- [About removing a node from a cluster](#)
- [Removing a node from a cluster](#)
- [Modifying the VCS configuration files on existing nodes](#)
- [Removing security credentials from the leaving node](#)

About removing a node from a cluster

You can remove one or more nodes from an SF Sybase CE cluster. The following table provides a summary of the tasks required to add a node to an existing SF Sybase CE cluster.

Table 23-1 Tasks for removing a node from a cluster

Step	Description
Prepare to remove the node: <ul style="list-style-type: none">■ Back up the configuration file.■ Check the status of the nodes and the service groups.■ Take the service groups offline and removing the database instances.	See “Removing a node from a cluster” on page 242.
Remove the node from the cluster.	See “Removing a node from a cluster” on page 242.

Table 23-1 Tasks for removing a node from a cluster (*continued*)

Step	Description
Modify the cluster configuration on remaining nodes. <ul style="list-style-type: none"> ■ Edit the /etc/llthosts file. ■ Edit the /etc/gabtab file. ■ Modify the VCS configuration to remove the node. 	See “Modifying the VCS configuration files on existing nodes” on page 243.
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See “Removing security credentials from the leaving node ” on page 246.

The Veritas product installer does not support removing a node. You must remove a node manually. The example procedures describe how to remove a node from a cluster with three nodes.

Removing a node from a cluster

Perform the following steps to remove a node from a cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

To prepare to remove a node from a cluster

- 1 Take the Sybase ASE CE service group offline (if under VCS control) on the node you want to remove.

```
# hagrps -offline sybase_group -sys system3
```

- 2 Remove the Sybase ASE CE database instance from the node.

For instructions, see the Sybase ASE CE documentation.

- 3 Take the *binmnt* service group offline (if under VCS control) on the node you want to remove.

```
# hagrps -offline binmnt_group -sys system3
```

- 4 Stop the applications that use Veritas File System (VxFS) or Cluster Files System (CFS) mount points and are not configured under VCS. Use native application commands to stop the applications.

- 5 Uninstall Sybase ASE CE from the node.

For instructions, see the Sybase ASE CE documentation.

To remove a node from a cluster

- 1 Unmount the VxFS/CFS file systems that are not configured under VCS.

```
# umount mount_point
```

- 2 Stop VCS on the node:

```
# hastop -local
```

- 3 Uninstall SF Sybase CE from the node using the SF Sybase CE installer.

```
# cd /opt/VRTS/install
```

```
# ./uninstallsfsybasece system3
```

The installer stops all SF Sybase CE processes and uninstalls the SF Sybase CE RPMs.

- 4 Modify the VCS configuration files on the existing nodes to remove references to the deleted node.

Modifying the VCS configuration files on existing nodes

Modify the configuration files on the remaining nodes of the cluster to remove references to the deleted nodes.

Tasks for modifying the cluster configuration files:

- Edit the `/etc/llhosts` file
- Edit the `/etc/gabtab` file
- Modify the VCS configuration to remove the node

For an example `main.cf`:

To edit the `/etc/llthosts` file

- ◆ On each of the existing nodes, edit the `/etc/llthosts` file to remove lines that contain references to the removed nodes.

For example, if `system3` is the node removed from the cluster, remove the line "2 system3" from the file:

```
0 sys1
1 sys2
2 system3
```

Change to:

```
0 sys1
1 sys2
```

To edit the `/etc/gabtab` file

- ◆ Modify the following command in the `/etc/gabtab` file to reflect the number of systems after the node is removed:

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

where `N` is the number of remaining nodes in the cluster.

For example, with two nodes remaining, the file resembles:

```
/sbin/gabconfig -c -n2
```

Modify the VCS configuration file `main.cf` to remove all references to the deleted node.

Use one of the following methods to modify the configuration:

- Edit the `/etc/VRTSvcs/conf/config/main.cf` file
This method requires application down time.
- Use the command line interface
This method allows the applications to remain online on all remaining nodes.

The following procedure uses the command line interface and modifies the sample VCS configuration to remove references to the deleted node. Run the steps in the procedure from one of the existing nodes in the cluster. The procedure allows you to change the VCS configuration while applications remain online on the remaining nodes.

To modify the cluster configuration using the command line interface (CLI)

- 1 Back up the `/etc/VRTSvcs/conf/config/main.cf` file.

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.3node.bak
```

- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Remove the node from the `AutoStartList` attribute of the service group by specifying the remaining nodes in the desired order:

```
# hagr -modify cvm AutoStartList sys1 sys2
```

- 4 Remove the deleted node from the system list of any other parent service groups to CVM that exist on the cluster before removing CVM. For example, to delete the node `system3`:

```
# hagr -modify syb_grp SystemList -delete system3
# hagr -modify Sybase SystemList -delete system3
# hagr -modify cvm SystemList -delete system3
# hares -modify cvm_clus CVMNodeId -delete system3
```

- 5 If you have a local VxFS configuration, will also need to remove the diskgroup of node to be removed from `binmnt`.

```
# hares -modify sybase_install_dg DiskGroup -delete \
sybase_new_diskgroup
```

- 6 Remove the node from the `SystemList` attribute of the service group:

If the system is part of the `SystemList` of a parent group, it must be deleted from the parent group first.

- 7 Remove the node from the `CVMNodeId` attribute of the service group:

```
# hares -modify cvm_clus CVMNodeId -delete system3
```

- 8 Remove the deleted node from the `NodeList` attribute of all CFS mount resources:

```
# hares -modify CFSMount NodeList -delete system3
```

- 9 Remove the deleted node from the cluster system list:

```
# hasys -delete system3
```

- 10 Save the new configuration to disk:

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

- 11 Verify that the node is removed from the VCS configuration.

```
# grep -i system3 /etc/VRTSvcs/conf/config/main.cf
```

If the node is not removed, use the VCS commands as described in this procedure to remove the node.

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node sys5. Perform the following steps.

To remove the security credentials

- 1 Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \  
stop
```

- 2 Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Configuration of disaster recovery environments

- [Chapter 24. Configuring disaster recovery environments](#)

Configuring disaster recovery environments

This chapter includes the following topics:

- [Disaster recovery options for SF Sybase CE](#)
- [About setting up a global cluster environment for SF Sybase CE](#)
- [About configuring a parallel global cluster using Veritas Volume Replicator \(VVR\) for replication](#)

Disaster recovery options for SF Sybase CE

Storage Foundation for Sybase CE supports configuring a disaster recovery environment using global clustering (GCO) using Veritas Volume Replicator (VVR) for replication.

For more about planning for disaster recovery environments:

See [“About global clusters”](#) on page 22.

See [“About Veritas Volume Replicator”](#) on page 23.

See [“Supported replication technologies for global clusters”](#) on page 33.

You can install and configure clusters for your disaster recovery environment as you would for any cluster using the procedures in this installation guide.

For a high level description of the tasks for implementing disaster recovery environments:

See [“About setting up a global cluster environment for SF Sybase CE”](#) on page 250.

See [“About configuring a parallel global cluster using Veritas Volume Replicator \(VVR\) for replication”](#) on page 250.

For complete details for configuring your disaster recovery environment once clusters are installed and configured:

See the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*.

About setting up a global cluster environment for SF Sybase CE

Configuring a global cluster for environment with parallel clusters requires the coordination of many component setup tasks. The procedures provided here are guidelines. You will need this guide to install and configure SF Sybase CE on each cluster. Refer to the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Guide* to configure a global cluster environment and replication between the two clusters.

- Configure a SF Sybase CE cluster at the primary site
- Configure an SF Sybase CE cluster at the secondary site
- Configure a global cluster environment
- Test the HA/DR configuration

Upon successful testing, you can bring the environment into production

For global cluster configuration details:

See the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Guide*.

About configuring a parallel global cluster using Veritas Volume Replicator (VVR) for replication

Configuring a global cluster for environment with SF Sybase CE and Veritas Volume Replicator requires the coordination of many component setup tasks. The tasks listed below are guidelines.

Before configuring two clusters for global clustering, you must verify that:

- You have the correct installation options enabled for SF Sybase CE, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster and VVR enabled.
Review SF Sybase CE requirements and licensing information.
- Both clusters have SF Sybase CE software installed and configured.

About configuring a parallel global cluster using Veritas Volume Replicator (VVR) for replication

Note: You can install and configure both clusters at the same time, or you can configure the second cluster at a later time than the first.

You can use this guide to install and configure SF Sybase CE on each cluster. For details for configuring a global cluster environment and replication between the the clusters using VVR:

See the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*.

See the *Veritas Storage Foundation for Sybase ASE CE Installation and Configuration Guide*.

With two clusters installed and configured , you are ready to configure a global cluster environment using VVR. You must perform the following tasks to modify both cluster configurations to support replication in the global cluster environment.

Table 24-1 Tasks for configuring a parallel global cluster with VVR

Task	Description
Setting up replication on the primary site	<ul style="list-style-type: none"> ■ Create the Storage Replicator Log (SRL) in the disk group for the database. ■ Create the Replicated Volume Group (RVG) on the primary site.
Setting up replication on the secondary site	<ul style="list-style-type: none"> ■ Create a disk group to hold the data volume, SRL, and RVG on the storage on the secondary site. You must match the names and sizes of these volumes with the names and sizes of the volumes on the primary site. ■ Edit the <code>/etc/vx/vras/.rdg</code> file on the secondary site. ■ Use resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites. ■ Create the replication objects on the secondary site.
Starting replication of the database.	<p>You can use either of the following methods to start replication:</p> <ul style="list-style-type: none"> ■ Automatic synchronization ■ Full synchronization with Storage Checkpoint

Table 24-1 Tasks for configuring a parallel global cluster with VVR (*continued*)

Task	Description
Configuring VCS for replication on clusters at both sites.	Configure Veritas Cluster Server (VCS) to provide high availability for the database: <ul style="list-style-type: none">■ Modify the VCS configuration on the primary site■ Modify the VCS configuration on the secondary site

Once the global clusters and replication with VVR are configured, the following replication use cases are supported for it:

- Migration of the role of the primary site to the remote site
- Takeover of the primary site role by the secondary site
- Migrate the role of primary site to the secondary site
- Migrate the role of new primary site back to the original primary site
- Take over after an outage
- Resynchronize after an outage
- Update the rlink to reflect changes

For details on the replication use cases:

See the *Veritas Storage Foundation and High Availability Solutions Disaster Recovery Implementation Guide*.

Uninstallation of SF Sybase CE

- [Chapter 25. Preparing to uninstall SF Sybase CE from a cluster](#)
- [Chapter 26. Uninstalling SF Sybase CE using the product installer](#)
- [Chapter 27. Performing an automated uninstallation of SF Sybase CE using response files](#)

Preparing to uninstall SF Sybase CE from a cluster

This chapter includes the following topics:

- [About uninstalling SF Sybase CE from a cluster](#)
- [Options for uninstalling SF Sybase CE](#)
- [Preparing to uninstall SF Sybase CE from a cluster](#)

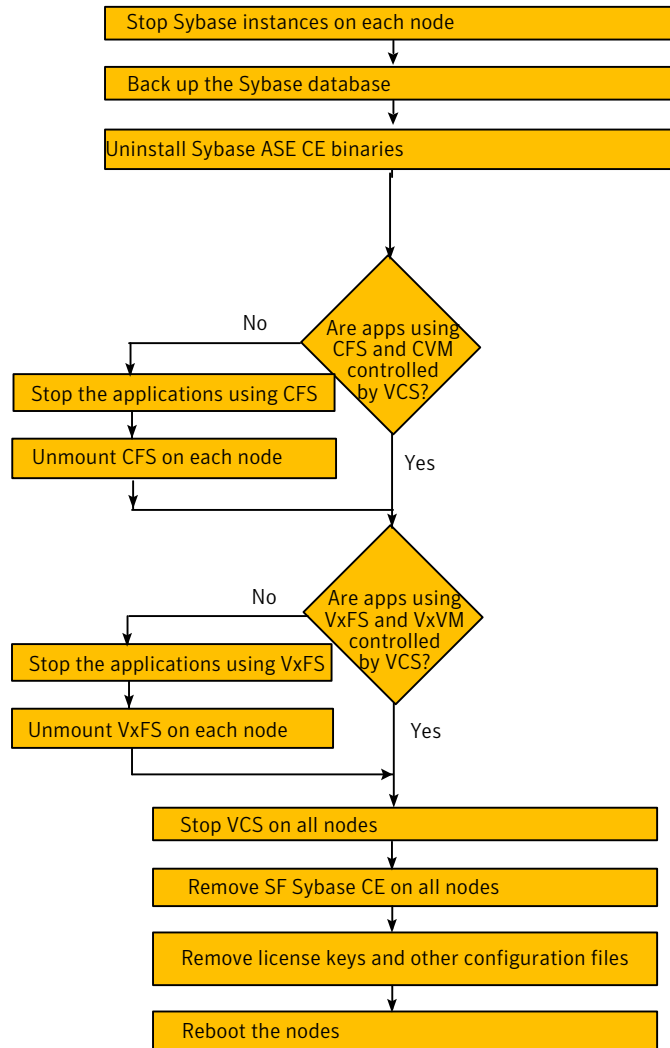
About uninstalling SF Sybase CE from a cluster

You can uninstall SF Sybase CE using the `uninstallsfsybasece`.

Note: After you uninstall SF Sybase CE, you cannot access the Sybase database as Veritas Volume Manager and Veritas File System are uninstalled from the cluster. Make sure that you back up the Sybase database before you uninstall SF Sybase CE.

[Figure 25-1](#) illustrates the steps that are required to uninstall SF Sybase CE from a cluster.

Figure 25-1 SF Sybase CE uninstallation



Options for uninstalling SF Sybase CE

Table 25-1 lists the available options for uninstalling SF Sybase CE:

Table 25-1 Options for uninstalling SF Sybase CE

Options	Description
SF Sybase CE uninstallation program	Use the <code>uninstallsybasece</code> program to uninstall SF Sybase CE.
Response file	Use a response file to automate or perform an unattended uninstallation of SF Sybase CE. See “Uninstalling SF Sybase CE using a response file” on page 269.

Preparing to uninstall SF Sybase CE from a cluster

Perform the steps in the following procedure before you uninstall SF Sybase CE from a cluster.

To prepare to uninstall SF Sybase CE from a cluster

- 1 Stop applications that use the Sybase ASE CE database.
See [“Stopping applications that use the Sybase database”](#) on page 258.
- 2 Stop Sybase instances.
See [“Stopping Sybase instances”](#) on page 258.
- 3 Back up the Sybase database.
See [“Backing up the Sybase database”](#) on page 259.
- 4 Uninstalling Sybase ASE CE (optional)
See [“Uninstalling Sybase ASE CE \(optional\)”](#) on page 259.
- 5 Remove root disk encapsulation.
See [“Removing root disk encapsulation”](#) on page 260.
- 6 Stop the applications that use CFS (outside of VCS control).
See [“Stopping the applications that use CVM or CFS \(outside of VCS control\)”](#) on page 261.
- 7 Unmount CFS file systems (outside of VCS control).
See [“Unmounting CFS file systems \(outside of VCS control\)”](#) on page 261.
- 8 Stop VCS.
See [“Stopping VCS”](#) on page 262.

- 9 Stop the applications that use VxFS (outside of VCS control).
See “[Stopping the applications that use VxVM or VxFS \(outside of VCS control\)](#)” on page 262.
- 10 Unmount VxFS file systems (outside of VCS control).
See “[Unmounting VxFS file systems \(outside of VCS control\)](#)” on page 263.

Stopping applications that use the Sybase database

Stop the applications that are dependent on service groups that contain Sybase resources.

To stop applications that use the Sybase database

- 1 Review the dependencies between service groups:

```
# hagrps -dep
```

- 2 Stop the service groups on each node:

```
# hagrps -offline app_group -sys node_name
```

Stopping Sybase instances

You need to stop Sybase CE and the Sybase instances on the cluster nodes where you want to uninstall SF Sybase CE. Before you stop the Sybase instances, stop the applications that are dependent on the service groups that contain Sybase.

The procedure in this section provides instructions to stop the instances on a two-node cluster; the nodes are sys1 and sys2. Depending on the VCS configuration, the procedure to stop Sybase instances may vary.

To stop Sybase instances

- 1 Log in as the superuser on one of the nodes in the cluster.
- 2 On each node, take the Sybase resources in the VCS configuration file (main.cf) offline.

```
# hagr -offline Sybase_group -sys node_name
```

For example:

```
# /opt/VRTSvcs/bin/hagr -offline sybasece -sys sys1
```

```
# /opt/VRTSvcs/bin/hagr -offline sybasece -sys sys2
```

These commands stop the Sybase resources under VCS control.

- 3 Verify that the state of the Sybase and CVM service groups are offline and online respectively.

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

Group	Attribute	System	Value
binmnt	State	sys1	ONLINE
binmnt	State	sys2	ONLINE
cvm	State	sys1	ONLINE
cvm	State	sys2	ONLINE
sybasece	State	sys1	OFFLINE
sybasece	State	sys2	OFFLINE

Backing up the Sybase database

If you plan to retain the Sybase database, you must back up the Sybase database.

For instructions on backing up the Sybase database, see the Sybase documentation.

Uninstalling Sybase ASE CE (optional)

Uninstall Sybase ASE CE before uninstalling SF Sybase CE. For information about the Sybase ASE CE uninstall utility, see the Sybase ASE CE product documentation.

To uninstall Sybase ASE CE

- 1 Log in as the Sybase user.

Note: In case of CFS binary installation, log in to any node. In case of Sybase ASE CE binary installation on local VxFS, you must uninstall from each node in cluster.

- 2 Set the DISPLAY variable. Depending on the shell you use, run the following command:

```
Bourne Shell (sh or ksh)  $ DISPLAY=host:0.0;export DISPLAY
```

```
C Shell (csh or tcsh)    $ setenv DISPLAY host:0.0
```

- 3 Run the uninstall utility.

```
# /cd $SYBASE_HOME/sybuninstallASESuite
```

- 4 Run `uninstall`.

```
# ./uninstall
```

Removing root disk encapsulation

Perform this step only if you plan to remove the VxVM and VVR RPMs.

If you have VxVM and VVR installed, you need to indicate to the installer whether or not you want to remove the VxVM RPMs from all nodes in the cluster. If you want to remove these RPMs, you need to ensure that the root disk is not encapsulated. The uninstallation fails if you choose to remove these RPMs while the root disk is encapsulated.

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, `unmirror` and `unencapsulate` the root disk as described in the following procedure.

To remove root disk encapsulation

- 1 Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- 2 Convert all the encapsulated volumes in the root disk to make them accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

- 3 To check if the root disk is unencapsulated:

```
# df -v /
```

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

Stopping the applications that use CVM or CFS (outside of VCS control)

You need to stop the applications that use CVM volumes or CFS mount points not controlled by VCS.

To stop the applications that use CVM or CFS (outside of VCS control)

- 1 Stop the applications that use a CFS mount point. The procedure varies for different applications. Use the procedure appropriate for the application.
- 2 Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

Unmounting CFS file systems (outside of VCS control)

You need to unmount CFS file systems that are not under VCS control on all nodes.

To unmount CFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the mount command. The command lists all the mounted clustered file systems. Consult the main.cf file for identifying the files that are under VCS control.

```
# mount | grep vxfs | grep cluster
```

- 2 Unmount each file system that is not controlled by VCS:

```
# umount mount_point
```

Stopping VCS

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

To stop VCS

- 1 Log in as the superuser on one of the cluster nodes.
- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped.

In this command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This output indicates that VCS has been stopped.

```
# /sbin/gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

```
Port a gen 5c3d0b membership 01
```

```
Port b gen 5c3d10 membership 01
```

Stopping the applications that use VxVM or VxFS (outside of VCS control)

You need to stop all applications that use VxVM volumes or VxFS mount points not under VCS control.

To stop the applications that use VxVM or VxFS (outside of VCS control)

- 1 Stop the applications that use a VxFS mount point. The procedure varies for different applications. Use the procedure that is appropriate for your application.
- 2 Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

Unmounting VxFS file systems (outside of VCS control)

You need to unmount VxFS file systems that are not under VCS control on all nodes.

Note: To avoid issues on rebooting, you must remove all entries of VxFS from the `/etc/fstab` file.

To unmount VxFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the `mount` command. The command lists all the mounted file systems.

```
# mount | grep vxfs
```

- 2 Unmount each file system that is not under VCS control:

```
# umount mount_point
```


Uninstalling SF Sybase CE using the product installer

This chapter includes the following topics:

- [Uninstalling SF Sybase CE with the script-based installer](#)

Uninstalling SF Sybase CE with the script-based installer

Perform the steps in the following procedure to remove SF Sybase CE from a cluster.

To remove SF Sybase CE from a cluster

- 1 Remove the SF Sybase CE RPMs. You can remove the RPMs using the uninstallation program or using the response file.

Using the uninstallation program:

See [“Removing the SF Sybase CE RPMs”](#) on page 265.

Using the response file:

See [“Uninstalling SF Sybase CE using a response file”](#) on page 269.

- 2 Remove other configuration files (optional).

See [“Removing other configuration files \(optional\)”](#) on page 267.

Removing the SF Sybase CE RPMs

The `uninstallsfsybasece` can remove these RPMs only if there are no open volumes.

The installer performs the following tasks:

- Removes the SF Sybase CE RPMs.
- Removes the language RPMs, if installed.

Note: The following directories remain after uninstallation: `/opt/VRTS`, `/opt/VRTSperl`, `/etc/VRTSvcs`, `/var/VRTSvcs`. They contain logs and configuration information for future reference. You may or may not remove them.

To remove the SF Sybase CE RPMs

- 1 Log in as the superuser on any node in the cluster.
- 2 Navigate to the directory that contains the `uninstallsfsybasece`:

```
# cd /opt/VRTS/install
```

- 3 Start the `uninstallsfsybasece`:

```
# ./uninstallsfsybasece<version> [-rsh]
```

Where `<version>` is the specific release version.

See “[About the Veritas installer](#)” on page 59.

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

- 4 If you have VxVM and VVR installed, indicate whether or not you want to remove VxVM and VVR RPMs from all nodes in the cluster. Enter `y` only if the root disk is outside of VxVM control.
- 5 To check if the root disk is under VxVM control:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Convert all the encapsulated volumes in the root disk to make them accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

- 6 If you invoked the `uninstallsfbasece` from a remote system in the same subnet, enter the name of the systems from which you want to uninstall SF Sybase CE.

If you invoked the `uninstallsfbasece` from a node in the SF Sybase CE cluster, review the cluster information and confirm to uninstall SF Sybase CE.

The `uninstallsfbasece` performs the following task:

- Checks the operating system on each node
 - Verifies the system-to-system communication
 - Verifies the licenses
 - Checks for the SF Sybase CE RPMs installed on the nodes. This process involves identifying system uninstallation requirements and dependencies between RPMs to determine the safety and order of uninstalling RPMs.
- 7 Confirm to uninstall SF Sybase CE.

The program performs the following tasks:

- Stops the agents and performs verifications on each node to proceed with uninstallation
- Stops the SF Sybase CE processes and uninstalls the SF Sybase CE RPMs
- Displays the location of the uninstallation summary, response file, and log files for reference.

Removing other configuration files (optional)

You can remove the Veritas configuration files and the RPMs that are left after running the `uninstallsfbasece`.

To remove residual Veritas configuration files (optional)

- 1 List all VRTS RPMs that can be removed.

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

- 2 Run the `rpm -e rpm_name` command to remove the remaining VRTS RPMs.
- 3 Move the residual Veritas configuration files to the `vrts.bkp` directory:

```
# cd /var
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# mv vx vrts.bkp
# cd /var/opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /etc
# mkdir vrts.bkp
# mv vx *llt* *fen* *gab* *vcs* vrts.bkp
```

You can remove the `vrts.bkp` directories at a later time.

Performing an automated uninstalation of SF Sybase CE using response files

This chapter includes the following topics:

- [Uninstalling SF Sybase CE using a response file](#)
- [Response file variables to uninstall Veritas Storage Foundation for Sybase ASE CE](#)
- [Sample response file for uninstalling SF Sybase CE](#)

Uninstalling SF Sybase CE using a response file

Perform the steps in the following procedure to uninstall SF Sybase CE using a response file.

To uninstall SF Sybase CE 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 189.

Note: You must replace the host names in the response file with that of the systems from which you want to uninstall SF Sybase CE.

For a sample response file:

See [“Sample response file for uninstalling SF Sybase CE”](#) on page 271.

- 3 Navigate to the directory containing the SF Sybase CE uninstallation program:

```
# cd /opt/VRTS/install
```

- 4 Start the uninstallation:

```
# ./uninstallsfsybasece<version> -responsefile /tmp/response_file
```

Where *<version>* is the specific release version.

See [“About the Veritas installer”](#) on page 59.

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

- 5 Optionally, remove residual configuration files, if any.

See [“Removing other configuration files \(optional\)”](#) on page 267.

Response file variables to uninstall Veritas Storage Foundation for Sybase ASE CE

[Table 27-1](#) lists the response file variables that you can define to configure SF Sybase CE.

Table 27-1 Response file variables for uninstalling SF Sybase CE

Variable	Description
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required

Table 27-1 Response file variables for uninstalling SF Sybase CE (*continued*)

Variable	Description
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
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}{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}{uninstall}	Uninstalls SF Sybase CE RPMs. List or scalar: scalar Optional or required: optional

Sample response file for uninstalling SF Sybase CE

The following sample response file uninstalls SF Sybase CE from nodes, sys1 and sys2.

```
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{prod}="SFSYBASECE60";
$CFG{systems}=[ qw(sys1 sys2) ];

1;
```


Installation reference

- [Appendix A. SF Sybase CE installation RPMs](#)
- [Appendix B. Installation scripts](#)
- [Appendix C. Sample installation and configuration values](#)
- [Appendix D. Tunable files for installation](#)
- [Appendix E. Configuration files](#)
- [Appendix F. High availability agent information](#)
- [Appendix G. Compatibility issues when installing Storage Foundation for Sybase ASE CE with other products](#)

SF Sybase CE installation RPMs

This appendix includes the following topics:

- [SF Sybase CE installation RPMs](#)

SF Sybase CE installation RPMs

[Table A-1](#) lists the RPM name and contents for each SF Sybase CE RPM.

Table A-1 List of SF Sybase CE RPMs

RPM	Content	Configuration
VRTSgab	Depends on VRTSllt. Contains the binaries for Veritas Cluster Server group membership and atomic broadcast services.	Minimum
VRTSllt	Contains the binaries for Veritas Cluster Server low-latency transport.	Minimum
VRTSamf	Contains the binaries for the Veritas Asynchronous Monitoring Framework kernel driver functionality for the process and mount based agents.	Minimum
VRTSperl	Contains Perl for Veritas.	Minimum
VRTSspt	Contains the binaries for Veritas Software Support Tools.	Recommended

Table A-1 List of SF Sybase CE RPMs (*continued*)

RPM	Content	Configuration
VRTSvcs	<p>Depends on VRTSvxfen, VRTSgab, and VRTSslt.</p> <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. 	Minimum
VRTSvcsag	<p>Depends on VRTSvcs.</p> <p>Contains the binaries for Veritas Cluster Server bundled agents.</p>	Minimum
VRTSvcssea	<p>Required for VCS with the high availability agent for Sybase.</p> <p>VRTSvcssea contains the binaries for Veritas high availability agents for DB2, Sybase, and Oracle.</p>	Recommended
VRTSvlic	<p>Contains the binaries for Symantec License Utilities.</p>	Minimum
VRTSvxfen	<p>Depends on VRTSgab.</p> <p>Contains the binaries for Veritas I/O fencing.</p>	Minimum
VRTScavf	<p>Veritas Cluster Server Agents for Storage Foundation Cluster File System</p>	Minimum
VRTSfssdk	<p>Veritas File System Software Developer Kit</p> <p>For VxFS APIs, the RPM 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	<p>Veritas Group Lock Manager for Storage Foundation Cluster File System</p>	Minimum
VRTSob	<p>Veritas Enterprise Administrator</p>	Recommended
VRTSvxfs	<p>Veritas File System binaries</p>	Minimum

Table A-1 List of SF Sybase CE RPMs (*continued*)

RPM	Content	Configuration
VRTSvxvm	Veritas Volume Manager binaries	Minimum
VRTSaslapm	Volume Manager ASL/APM	Minimum
VRTSsfcp601	Veritas Storage Foundation Common Product Installer The Storage Foundation Common Product installer RPM contains the scripts that perform the following functions: installation, configuration, upgrade, uninstallation, adding nodes, and removing nodes. You can use this script to simplify the native operating system installations, configurations, and upgrades.	Minimum
VRTSsfmh	Veritas Storage Foundation Managed Host	Recommended
VRTSfsadv	Veritas File System Advanced Features by Symantec	Minimum
VRTSvcsdr	Veritas Cluster Server Disk Reservation Modules and Utilities by Symantec	Recommended
VRTSlvconv	Veritas Volume Manager by Symantec, LVM configuration converter. Tool for conversion of LVM configuration to Veritas Volume Manager.	All
VRTSvbs	Veritas Virtual Business Service.	Recommended

Installation scripts

This appendix includes the following topics:

- [Installation script options](#)
- [About using the postcheck option](#)

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 the Veritas installer”](#) on page 59.

Table B-1 Available command line options

Commandline Option	Function
-addnode	Adds a node to a high availability cluster.
-allpkgs	Displays all RPMs required for the specified product. The RPMs are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-configure	Configures the product after installation.
-fencing	Configures I/O fencing in a running cluster.
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-installallpkgs	The <code>-installallpkgs</code> option is used to select all RPMs.
-installrecpkgs	The <code>-installrecpkgs</code> option is used to select the recommended RPMs set.
-installminpkgs	The <code>-installminpkgs</code> option is used to select the minimum RPMs set.
-ignorepatchreqs	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
-kickstart <i>dir_path</i>	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec RPMs in the correct order for installing, in a format that can be used for Kickstart installations. The <i>dir_path</i> indicates the path to the directory in which to create the file.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	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.
-makeresponsefile	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-minpkgs	Displays the minimal RPMs required for the specified product. The RPMs 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.
-nolic	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkginfo	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.
-pkgpath <i>package_path</i>	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.
-pkgset	Discovers and displays the RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
-pkgtable	Displays product's RPMs in correct installation order by group.
-postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-recpkgs	Displays the recommended RPMs required for the specified product. The RPMs 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.
-redirect	Displays progress details without showing the progress bar.
-requirements	The <code>-requirements</code> option displays required OS version, required RPMs and patches, file system space, and other system requirements in order to install the product.
-responsefile <i>response_file</i>	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.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version.
-rollingupgrade_phase2	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The <code>-timeout</code> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 prevents the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
-tmppath <i>tmp_path</i>	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.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file <i>tunables_file</i>	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-upgrade_kernelpkgs	The <code>-upgrade_kernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase1</code> .
-upgrade_nonkernelpkgs	The <code>-upgrade_nonkernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase2</code> .
-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. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
-yumgroupxml	The <code>-yumgroupxml</code> option is used to generate a yum group definition XML file. The <code>createrepo</code> command can use the file on Redhat Linux to create a yum group for automated installation of all RPMs for a product. An available location to store the XML file should be specified as a complete path. The <code>-yumgroupxml</code> option is supported on Redhat Linux only.

About using the postcheck option

You can use the installer's `post-check` to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists (if LLT is not running at the check time).

- The `VRTSllt` pkg version is not consistent on the nodes.
- The `llt-linkinstall` value is incorrect.
- The `llthosts(4)` or `llttab(4)` configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB `linkinstall` value exists.
- The `VRTSgab` pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The `VRTSvcs` pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg` file is invalid.
- The `vxfen` link-install value is incorrect.
- The `VRTSvxfen` pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the `Autostartlist` value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because `Vxfen` is not started.
- Cluster Volume Manager cannot start because `gab` is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

You can use the installer's post-check option to perform the following checks:

General checks for all products:

- All the required RPMs are installed.
- The versions of the required RPMs are correct.
- There are no verification issues for the required RPMs.

Checks for Volume Manager (VM):

- Lists the daemons which are not running (`vxattachd`, `vxconfigbackupd`, `vxesd`, `vxrelocd` ...).
- Lists the disks which are not in 'online' or 'online shared' state (`vxdisk list`).
- Lists the diskgroups which are not in 'enabled' state (`vx dg list`).
- Lists the volumes which are not in 'enabled' state (`vxprint -g <dgname>`).
- Lists the volumes which are in 'Unstartable' state (`vxinfo -g <dgname>`).
- Lists the volumes which are not configured in (AIX) `/etc/filesystems`, (Linux/HP-UX) `/etc/fstab`, or (SunOS) `/etc/vfstab`.

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`).
- Whether all VxFS file systems present in (AIX) `/etc/filesystems`, (Linux/HP-UX) `/etc/fstab`, or (SunOS) `/etc/vfstab` file are mounted.
- Whether all VxFS file systems present in (AIX) `/etc/filesystems`, (Linux/HP-UX) `/etc/fstab`, or (SunOS) `/etc/vfstab` are in disk layout 6 or higher.
- Whether all mounted VxFS file systems are in disk layout 6 or higher.

Checks for Cluster File System:

- Whether FS and ODM are running at the latest protocol level.
- Whether all mounted CFS file systems are managed by VCS.
- Whether `cvm` service group is online.

See [“Performing a postcheck on a node”](#) on page 101.

Sample installation and configuration values

This appendix includes the following topics:

- [SF Sybase CE installation and configuration information](#)
- [SF Sybase CE worksheet](#)

SF Sybase CE installation and configuration information

The SF Sybase CE installation and configuration program prompts you for information about SF Sybase CE. It also provides default values for some information which you can choose to use. The worksheets provide sample values that you can use as examples of the information required for an SF Sybase CE installation and configuration.

Symantec recommends using the worksheets provided to record values for your systems before you begin the installation and configuration process.

SF Sybase CE worksheet

[Table C-1](#) contains the sample values that may be used when you install and configure SF Sybase CE. Enter the SF Sybase CE values for your systems in the following table:

Table C-1 SF Sybase CE worksheet

Installation information	Sample value	Assigned value
Number of nodes in the cluster	2	
Host names for Primary cluster	sys1 and sys2	
Host names for added or removed node	sys5	
SF Sybase CE License key	License keys are in the format: XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX	
Required SF Sybase CE RPMs vs. all SF Sybase CE	Install only the required RPMs if you do not want to configure any optional components or features. Default option is to install all RPMs.	
Primary cluster name	clus1	
Primary cluster ID number	101	
<p>Private network links</p> <p>You can choose a network interface card that is not part of any aggregated interface, or you can choose an aggregated interface.</p> <p>The interface names that are associated with each NIC for each network link must be the same on all nodes.</p> <p>Do not use the network interface card that is used for the public network, which is typically eth0.</p>	eth1,eth2	
Cluster Manager NIC (Primary NIC)	eth0	
Cluster Manager IP	10.10.12.1, 10.10.12.2	
Netmask for the virtual IP address	255.255.240.0	

Table C-1 SF Sybase CE worksheet (*continued*)

Installation information	Sample value	Assigned value
<p>Mode for Authentication Service:</p> <ul style="list-style-type: none"> ■ Automatic mode ■ Semiautomatic mode using encrypted files ■ Semiautomatic mode without using encrypted files <p>Default option is automatic mode.</p>	Automatic mode	
<p>User name</p> <p>Adding users is required if when using secure cluster mode. Otherwise it is optional.</p>	smith	
<p>User password</p>	password	
<p>User privilege</p> <p>VCS privilege levels include:</p> <ul style="list-style-type: none"> ■ Administrators— Can perform all operations, including configuration options on the cluster, service groups, systems, resources, and users. ■ Operators—Can perform specific operations on a cluster or a service group. ■ Guests—Can view specified objects. 	admin	
<p>Domain-based address of the SMTP server</p> <p>The SMTP server sends notification email about the events within the cluster.</p>	smtp.symantecexample.com	
<p>Email address of each SMTP recipient to be notified</p>	john@symantecexample.com	

Table C-1 SF Sybase CE worksheet (*continued*)

Installation information	Sample value	Assigned value
<p>Minimum severity of events for SMTP email notification</p> <p>The severity levels are defined as follows:</p> <ul style="list-style-type: none"> ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption 	E	
<p>Email address of SMTP notification recipients</p>	admin@symantecexample.com	
<p>SNMP trap daemon port number the console</p>	162	
<p>System name for the SNMP console</p>	system2	
<p>Minimum severity level of events for SMTP notification</p> <p>The severity levels are defined as follows:</p> <ul style="list-style-type: none"> ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption 	i	
<p>CVM enclosure-based naming</p> <p>Requires Dynamic Multi-pathing (DMP).</p>	yes	

Table C-1 SF Sybase CE worksheet (*continued*)

Installation information	Sample value	Assigned value
Default disk group You can select the name of a default disk group of a system for running Veritas Volume Manager commands which require a disk group to be specified.	vxfencoordg	
The name of three disks that form the coordinator disk group.	■ sdd ■ sde ■ sdf	
Vxfen disk group	vxfencoordg	

Tunable files for installation

This appendix includes the following topics:

- [About setting tunable parameters using the installer or a response file](#)
- [Setting tunables for an installation, configuration, or upgrade](#)
- [Setting tunables with no other installer-related operations](#)
- [Setting tunables with an un-integrated response file](#)
- [Preparing the tunables file](#)
- [Setting parameters for the tunables file](#)
- [Tunables value parameter definitions](#)

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

```
# ./installer -tunablesfile tunables_file_name
```

See “[Setting tunables for an installation, configuration, or upgrade](#)” on page 294.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
system1 system2 ...]
```

See “[Setting tunables with no other installer-related operations](#)” on page 295.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name
```

See “[Setting tunables with an un-integrated response file](#)” on page 296.

For more information on response files, see the *chapter: About response files*.

You must select the tunables that you want to use from this guide.

See “[Tunables value parameter definitions](#)” on page 298.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See “[Tunables value parameter definitions](#)” on page 298.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

- 1 Prepare the tunables file.
See “[Preparing the tunables file](#)” on page 297.
- 2 Make sure the systems where you want to install SF Sybase CE meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters. Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 298.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 297.
- 2 Make sure the systems where you want to install SF Sybase CE meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [
sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters. Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 298.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

- 1 Make sure the systems where you want to install SF Sybase CE meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 297.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-responsefile` and `-tunablesfile` options.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

- 7 Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

For more information on response files, see the *chapter: About response files*.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

You can use the installer to create a tunables file template, or manually format tunables files you create.

To create a tunables file template

- ◆ Start the installer with the `-tunables` option. Enter the following:

```
# ./installer -tunables
```

You see a list of all supported tunables, and the location of the tunables file template.

To manually format tunables files

- ◆ Format the tunable parameter as follows:

```
$TUN{"tunable_name"}{"system_name"|"*" }=value_of_tunable;
```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#  
# Tunable Parameter Values:  
#  
our %TUN;  
  
$TUN{"tunable1" }{"*" }=1024;  
$TUN{"tunable3" }{"sys123" }="SHA256";  
  
1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See [“Tunables value parameter definitions”](#) on page 298.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the dmp_daemon_count value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table D-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table D-1 Supported tunable parameters

Tunable	Description
dmp_cache_open	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_daemon_count	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_delayq_interval	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_fast_recovery	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_health_time	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_log_level	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer sets only the system default value of read_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer sets only the system default value of read_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
vol_checkpoint_default	(Veritas File System) Size of VxVM checkpoints (sectors). This tunable requires system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.

Table D-1 Supported tunable parameters (*continued*)

Tunable	Description
voliot_jobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Configuration files

This appendix includes the following topics:

- [About sample main.cf files](#)
- [Sample main.cf files for Sybase ASE CE configurations](#)

About sample main.cf files

You can examine the VCS configuration file, main.cf, to verify the SF Sybase CE installation and configuration.

- The main.cf file is located in the folder /etc/VRTSvcs/conf/config.
- After an SF Sybase CE installation, several sample main.cf file types can be viewed in the following directory: /etc/VRTSagents/ha/conf/Sybase
- All sample configurations assume that the Veritas High Availability Agent for Sybase binaries are installed on local disks and that they are managed by the operating system. These file systems must be specified in the file /etc/fstab
- For the following configuration samples, please note the "cluster" definition in all of the configurations should specify UseFence=SCSI3.

Sample main.cf files for Sybase ASE CE configurations

Sample main.cf file examples are provided for the following Sybase ASE CE configurations:

- Basic cluster configuration
 - With shared mount point on CFS for Sybase binary installation
 - With local mount point on VxFS for Sybase binary installation
- Replicating data between two clusters

- For a primary site in a CVM VVR configuration
- For a secondary site in a CVM VVR configuration

Sample main.cf for a basic Sybase ASE CE cluster configuration under VCS control with shared mount point on CFS for Sybase binary installation

This sample main.cf is for a single site with a basic cluster configuration with shared mount point on CFS for Sybase binary installation.

The following are the configuration details for this Sybase ASE CE configuration sample main.cf:

- File name: sybasece_cfs_main.cf
- File location: /etc/VRTSagents/ha/conf/Sybase/

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "SybaseTypes.cf"

cluster clus1 (
    UserNames = { admin = HopHojOlpKppNxpJom }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
    UseFence=SCSI3
)

system system1 (
)

system system2 (
)

// binmounts group for configuring CFS mounts for Sybase binaries.

group binmnt (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)
```

```

CFSMount sybbindg_101_sybbinvol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbindg_101/sybbin_vol"
)

CVMVolDg sybbindg_101_voldg (
    CVMDiskGroup = sybbindg_101
    CVMVolume = { sybbin_vol }
    CVMActivation = sw
)

requires group cvm online local firm
sybbindg_101_sybbinvol_mnt requires sybbindg_101_voldg

// resource dependency tree
//
// group binmnt
// {
//   CFSMount sybbindg_101_sybbinvol_mnt
//   {
//     CVMVolDg sybbindg_101_voldg
//   }
// }

// cvm group for CVM and CFS specific agents.

group cvm (
    SystemList = { system1 = 0, system2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = sfsyb_90
    CVMNodeId = { system1 = 0, system2 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

```

```
    )

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

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
// group cvm
// {
//     CFSfsckd vxfsckd
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
// }

// sybasece group for:
// 1. CVM volumes for Sybase database and quorum device
// 2. CFS mount for Sybase database and quorum device
// 3. Process agent for vxfsckd process.
// 4. Sybase database instance.

group sybasece (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount quorum_101_quorumvol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_101/quorumvol"
)
```

```

CFSMount sybdata_101_sybvool_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/sybdata_101/sybvool"
)

CVMVolDg quorum_101_voldg (
    CVMDiskGroup = quorum_101
    CVMVolume = { quorumvol }
    CVMActivation = sw
)

CVMVolDg sybdata_101_voldg (
    CVMDiskGroup = sybdata_101
    CVMVolume = { sybvool }
    CVMActivation = sw
)

Process vx fend (
    PathName = "/sbin/vxfend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

Sybase ase (
    Server @system1 = ase1
    Server @system2 = ase2
    Owner = sybase
    Home = "/sybase"
    Version = 15
    SA = sa
    Quorum_dev = "/quorum/q.dat"
)

requires group binmnt online local firm
ase requires quorum_101_quorumvol_mnt
ase requires sybdata_101_sybvool_mnt
ase requires vx fend
quorum_101_quorumvol_mnt requires quorum_101_voldg
sybdata_101_sybvool_mnt requires sybdata_101_voldg

// resource dependency tree
//
// group sybasece

```

```

// {
// Sybase ase
//   {
//     CFSMount quorum_101_quorumvol_mnt
//       {
//         CVMVolDg quorum_101_voldg
//       }
//     CFSMount sybdata_101_sybvol_mnt
//       {
//         CVMVolDg sybdata_101_voldg
//       }
//     Process vxfsend
//   }
// }

```

Sample main.cf for a basic Sybase ASE CE cluster configuration with local mount point on VxFS for Sybase binary installation

This sample main.cf is for a single site with a basic cluster configuration with local mount point on VxFS for Sybase binary installation.

The following are the configuration details for this Sybase ASE CE configuration sample main.cf:

- File name: sybasece_vxfs_main.cf
- File location: /etc/VRTSagents/ha/conf/Sybase/

```

include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "SybaseTypes.cf"

cluster clus1 (
    UserNames = { admin = HopHojOlpKppNxpJom }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
    UseFence=SCSI3
)

system system1 (
)

```

```

system system2 (
)

// binmounts group for configuring VxFS mounts for Sybase binaries.

group binlocalmnt (
    SystemList = { system1 = 0, system2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

DiskGroup sybbindg_101_voldg (
    DiskGroup = sybbindg
)

Mount sybbindg_101_sybbinvol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbindg_101/sybbin_vol"
    FSType = vxfs
    FsckOpt = "-y"
)

Volume sybbindg_101_vol (
    DiskGroup = sybbindg
    Volume = sybbinvol
)

requires group cvm online local firm
sybbindg_101_sybbinvol_mnt requires sybbindg_101_vol
sybbindg_101_vol requires sybbindg_101_voldgdg

// resource dependency tree
//
// group binlocalmnt
// {
// Mount sybbindg_101_sybbinvol_mnt
//     {
//     Volume sybbindg_vol
//         {

```

```
//          DiskGroup sybbindg_101_voldg
//          }
//      }
// }

// cvm group for CVM and CFS specific agents.

group cvm (
    SystemList = { system1 = 0, system2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = clus1
    CVMNodeId = { system1 = 0, system2 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

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

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
// group cvm
// {
//     CFSfsckd vxfsckd
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
// }
```



```

//      }
// }

// sybasece group for:
// 1. CVM volumes for Sybase database and quorum device
// 2. CFS mount for Sybase database and quorum device
// 3. Process agent for vxfeed process.
// 4. Sybase database instance.

group sybasece (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount quorum_101_quorumvol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_101/quorumvol"
)

CFSMount sybdata_101_sybvoldg_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/sybdata_101/sybvoldg"
)

CVMVolDg quorum_101_voldg (
    CVMDiskGroup = quorum_101
    CVMVolume = { quorumvol }
    CVMActivation = sw
)

CVMVolDg sybdata_101_voldg (
    CVMDiskGroup = sybdata_101
    CVMVolume = { sybvoldg }
    CVMActivation = sw
)

Process vxfeed (
    PathName = "/sbin/vxfeed"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

```

```

Sybase ase (
    Server @system1 = ase1
    Server @system2 = ase2
    Owner = sybase
    Home = "/sybase"
    Version = 15
    SA = sa
    Quorum_dev = "/quorum/q.dat"
)

requires group binlocalmnt online local firm
ase requires quorum_101_quorumvol_mnt
ase requires sybdata_101_sybvol_mnt
ase requires vxfsend
quorum_101_quorumvol_mnt requires quorum_101_voldg
sybdata_101_sybvol_mnt requires sybdata_101_voldg

// resource dependency tree
//
// group sybasece
// {
// Sybase ase
// {
//     CFSMount quorum_101_quorumvol_mnt
//     {
//         CVMVolDg quorum_101_voldg
//     }
//     CFSMount sybdata_101_sybvol_mnt
//     {
//         CVMVolDg sybdata_101_voldg
//     }
//     Process vxfsend
// }
// }

```

Sample main.cf for a primary CVM VVR site

This sample main.cf is for a primary site in a CVM VVR configuration. It is one of two sample main.cfs for replicating data between two clusters.

The following are the configuration details for this Sybase ASE CE configuration sample main.cf:

- File name: sybasece_cvmvvr_primary_main.cf
- File location: /etc/VRTSagents/ha/conf/Sybase

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "SybaseTypes.cf"

cluster clus1 (
    UserNames = { admin = HopHojOlpKppNxpJom }
    ClusterAddress = "10.180.88.188"
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
    UseFence=SCSI3
)

remoteclass clus2 (
    ClusterAddress = "10.190.99.199"
)

heartbeat Icmp (
    ClusterList = { clus2 }
    Arguments @clus2 = { "10.190.99.199" }

)

system system1 (
)

system system2 (
)

group ClusterService (
    SystemList = { system1 = 0, system2 = 1 }
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
```

```
        StartProgram = "/opt/VRTSvcs/bin/wacstart"
        StopProgram = "/opt/VRTSvcs/bin/wacstop"
        MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
        RestartLimit = 3
    )

IP gcoip (
    Device = eth0
    Address = "10.180.88.188"
    NetMask = "255.255.255.0"
)

NIC csgnic (
    Device = eth0
)

gcoip requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//     Application wac
//     {
//     IP gcoip
//     {
//     NIC csgnic
//     }
//     }
//     }
//     }

group RVGgroup (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
)

CVMVolDg sybdata_voldg (
    CVMDiskGroup = sybdata_101
    CVMActivation = sw
```

```

    )

    RVGShared sybdata_rvg (
        RVG = syb_rvg
        DiskGroup = sybdata_101
    )

    requires group binmnt online local firm
    sybdata_rvg requires sybdata_voldg

group binmnt (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount sybbindg_101_sybbinvol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbindg_101/sybbin_vol"
)

CVMVolDg sybbindg_101_voldg (
    CVMDiskGroup = sybbindg_101
    CVMVolume = { sybbin_vol }
    CVMActivation = sw
)

requires group cvm online local firm
sybbindg_101_sybbinvol_mnt requires sybbindg_101_voldg

group cvm (
    SystemList = { system1 = 0, system2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (

```

```
        CVMClustName = clus1
        CVMNodeId = { system1 = 0, system2 = 1 }
        CVMTransport = gab
        CVMTimeout = 200
    )

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

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
//     group cvm
//     {
//     CFSfsckd vxfsckd
//     {
//         CVMcluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
//     }
// }

group logowner (
    SystemList = { system1 = 0, system2 = 1 }
    AutoStartList = { system1, system2 }
)

IP logowner_ip (
    Device = eth0
    Address = "10.10.9.101"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = eth0
)
```

```

RVGLogowner rvg_logowner (
    RVG = syb_rvg
    DiskGroup = sybdata_101
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic

// resource dependency tree
//
//     group logowner
//     {
//     RVGLogowner rvg_logowner
//         {
//             IP logowner_ip
//                 {
//                     NIC nic
//                 }
//             }
//         }
//     }

group sybasece (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    ClusterList = { clus1 = 0, clus2 = 1 }
    AutoStartList = { system1, system2 }
    ClusterFailOverPolicy = Manual
    Authority = 1
    OnlineRetryLimit = 3
    TriggerResStateChange = 1
    OnlineRetryInterval = 120
)

CFSMount quorum_101_quorumvol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_101/quorumvol"
)

CFSMount sybdata_101_sybvol_mnt (
    MountPoint = "/sybdata"

```

Sample main.cf files for Sybase ASE CE configurations

```

        BlockDevice = "/dev/vx/dsk/sybdata_101/sybv1"
    )

    CVMVolDg quorum_101_voldg (
        CVMDiskGroup = quorum_101
        CVMVolume = { quorumvol }
        CVMActivation = sw
    )

    Process vxfend (
        PathName = "/sbin/vxfend"
        Arguments = "-m sybase -k /tmp/vcmp_socket"
    )

    RVGSharedPri syb_vvr_shpri (
        RvgResourceName = sybdata_rvg
        OnlineRetryLimit = 0
    )

    Sybase ase (
        Server @system1 = ase1
        Server @system2 = ase2
        Owner = sybase
        Home = "/sybase"
        Version = 15
        SA = sa
        Quorum_dev = "/quorum/q.dat"
    )

    requires group RVGgroup online local firm
    sybdata_101_sybv1_mnt requires syb_vvr_shpri
    ase requires vxfend
    ase requires sybdata_101_sybv1_mnt
    ase requires quorum_101_quorumvol_mnt
    quorum_101_quorumvol_mnt requires quorum_101_voldg

    // resource dependency tree
    //
    //     group sybasece
    //     {
    //     Sybase ase
    //     {
    //         CFSMount sybdata_101_sybv1_mnt

```



```

//          {
//          RVGSharedPri syb_vvr_shpri
//          }
//          Process vx fend
//          CFMount quorum_101_quorumvol_mnt
//          {
//          CVMVolDg quorum_101_voldg
//          }
//          }
//          }

```

Sample main.cf for a secondary CVM VVR site

This sample main.cf is for a secondary site in a CVM VVR configuration. It is the second of two sample main.cfs for replicating data between two clusters.

The following are the configuration details for this Sybase ASE CE configuration sample main.cf:

- File name: sybasece_cvmvvr_secondary_main.cf
- File location: /etc/VRTSagents/ha/conf/Sybase

This is main.cf for CVM VVR configuration on Secondary site.

```

-----

include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "SybaseTypes.cf"

cluster clus2 (
    UserNames = { admin = HopHojOlpKppNxpJom }
    ClusterAddress = "10.190.99.199"
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
    UseFence=SCSI3
)

remoteclass clus1 (
    ClusterAddress = "10.180.88.188"
)

heartbeat Icmp (
    ClusterList = { clus1 }
)

```

```
        Arguments @clus1 = { "10.180.88.188" }
    )

system system3 (
)

system system4 (
)

group ClusterService (
    SystemList = { system3 = 0, system4 = 1 }
    AutoStartList = { system3, system4 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)

IP gcoip (
    Device = eth0
    Address = "10.190.99.199"
    NetMask = "255.255.255.0"
)

NIC csgnic (
    Device = eth0
)

gcoip requires csgnic
wac requires gcoip

// resource dependency tree
//
// group ClusterService
// {
//   Application wac
//   {
//     IP gcoip
```

```

//      {
//      NIC csgnic
//      }
//  }
// }

group RVGgroup (
    SystemList = { system3 = 0, system4 = 1 }
    Parallel = 1
    AutoStartList = { system3, system4 }
)

CVMVolDg sybdata_voldg (
    CVMDiskGroup = sybdata_101
    CVMActivation = sw
)

RVGShared sybdata_rvg (
    RVG = syb_rvg
    DiskGroup = sybdata_101
)

requires group binmnt online local firm
sybdata_rvg requires sybdata_voldg

group binmnt (
    SystemList = { system3 = 0, system4 = 1 }
    Parallel = 1
    AutoStartList = { system3, system4 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount sybbindg_101_sybbinvol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbindg_101/sybbin_vol"
)

CVMVolDg sybbindg_101_voldg (
    CVMDiskGroup = sybbindg_101
    CVMVolume = { sybbin_vol }
    CVMActivation = sw
)

```

```

    )
    requires group cvm online local firm
    sybbindg_101_sybbinvol_mnt requires sybbindg_101_voldg

group cvm (
    SystemList = { system3 = 0, system4 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system3, system4 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = clus2
    CVMNodeId = { system3 = 0, system4 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

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

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
//     group cvm
//     {
//     CFSfsckd vxfsckd
//         {
//             CVMCluster cvm_clus
//                 {
//                     CVMVxconfigd cvm_vxconfigd
//                 }
//             }
//         }
//     }

```

```

group logowner (
    SystemList = { system3 = 0, system4 = 1 }
    AutoStartList = { system3, system4 }
)

IP logowner_ip (
    Device = eth0
    Address = "10.11.9.102"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = eth0
)

RVGLogowner rvg_logowner (
    RVG = syb_rvg
    DiskGroup = sybdata_101
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic

// resource dependency tree
//
// group logowner
// {
//   RVGLogowner rvg_logowner
//   {
//     IP logowner_ip
//     {
//       NIC nic
//     }
//   }
// }

group sybasece (
    SystemList = { system3 = 0, system4 = 1 }
    Parallel = 1
    ClusterList = { clus2 = 0, clus1 = 1 }
    AutoStartList = { system3, system4 }
)

```

```
OnlineRetryLimit = 3
OnlineRetryInterval = 120
)

CFSMount quorum_101_quorumvol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_101/quorumvol"
)

CVMVolDg quorum_101_voldg (
    CVMDiskGroup = quorum_101
    CVMVolume = { quorumvol }
    CVMActivation = sw
)

CFSMount sybdata_101_sybvol_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/sybdata_101/sybvol"
)

Process vxfend (
    PathName = "/sbin/vxfend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

RVGSharedPri syb_vvr_shpri (
    RvgResourceName = sybdata_rvg
    OnlineRetryLimit = 0
)

Sybase ase (
    Server @system3 = ase1
    Server @system4 = ase2
    Owner = sybase
    Home = "/sybase"
    Version = 15
    SA = sa
    Quorum_dev = "/quorum/q.dat"
)

requires group RVGgroup online local firm
sybdata_101_sybvol_mnt requires syb_vvr_shpri
ase requires vxfend
```

```
ase requires sybdata_101_sybvol_mnt
ase requires quorum_101_quorumvol_mnt
quorum_101_quorumvol_mnt requires quorum_101_voldg
```


High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount agent](#)
- [Process agent](#)
- [Monitoring options for the Sybase agent](#)
- [Sybase resource type](#)

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 Sybase or a Web server, or a configuration component such as an IP address or mounted file system.

Attributes are either optional or required, although sometimes the attributes that are optional in one configuration may be required in other configurations. Many optional attributes have predefined or default values, which you should change as required. A variety of internal use only attributes also exist. Do not modify these attributes—modifying them can lead to significant problems for your clusters. Attributes have type and dimension. Some attribute values can accept numbers, others can accept alphanumeric values or groups of alphanumeric values, while others are simple boolean on/off values.

The entry points and attributes for each SF Sybase CE agent are described in this appendix.

VCS agents included within SF Sybase CE

SF Sybase CE includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDg agent
- CFSMount agent

An SF Sybase CE installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

You must configure the CVMVolDg agent for each disk group that is used by an agent for Sybase service group. Configure a disk group for only a single agent for Sybase service group. If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.

Use the information in this appendix about the entry points and attributes of the listed agents to make necessary configuration changes. For information on how to modify the VCS configuration:

See the *Veritas Cluster Server Administrator's Guide*

VCS agent for Sybase included within SF Sybase CE

SF Sybase CE includes an additional agent for Sybase.

See the *Veritas Cluster Server Agent for Sybase Installation and Configuration Guide* for more information on the Sybase agent.

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 F-1](#) describes the entry points used by the CVMCluster agent.

Table F-1 CVMCluster agent entry points

Entry Point	Description
Online	Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by 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 F-2](#) describes the user-modifiable attributes of the CVMCluster resource type.

Table F-2 CVMCluster agent attributes

Attribute	Description
CVMClustName	Name of the cluster. <ul style="list-style-type: none"> ■ Type and dimension: string-scalar
CVMNodeAddr	List of host names and IP addresses. <ul style="list-style-type: none"> ■ Type and dimension: string-association
CVMNodeId	Associative list. The first part names the system; the second part contains the LLT ID number for the system. <ul style="list-style-type: none"> ■ Type and dimension: string-association

Table F-2 CVMCluster agent attributes (*continued*)

Attribute	Description
CVMTransport	<p>Specifies the cluster messaging mechanism.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar ■ Default = gab <p>Note: Do not change this value.</p>
PortConfigd	<p>The port number that is used by CVM for vxconfigd-level communication.</p> <ul style="list-style-type: none"> ■ Type and dimension: integer-scalar
PortKmsgd	<p>The port number that is used by CVM for kernel-level communication.</p> <ul style="list-style-type: none"> ■ Type and dimension: integer-scalar
CVMTimeout	<p>Timeout in seconds used for CVM cluster reconfiguration.</p> <ul style="list-style-type: none"> ■ Type and dimension: integer-scalar ■ Default = 200

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```

type CVMCluster (
    static keylist RegList = { CVMNodePreference }
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName,
                            CVMNodeAddr, CVMNodeId, PortConfigd,
                            PortKmsgd, CVMTimeout }

    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
    str CVMNodePreference
    int PortConfigd
    int PortKmsgd
    int CVMTimeout
)

```

Note: The attributes `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` are not used in an SF Sybase CE 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 = clus1  
    CVMNodeId = { sys1 = 0, sys2 = 1 }  
    CVMTransport = gab  
    CVMTimeout = 200  
)
```

CVMVxconfigd agent

The CVMVxconfigd agent starts and monitors the vxconfigd daemon. The vxconfigd daemon maintains disk and disk group configurations, communicates configuration changes to the kernel, and modifies the configuration information that is stored on disks. CVMVxconfigd must be present in the CVM service group.

The CVMVxconfigd agent is an OnOnly agent; the agent starts the resource when the cluster starts up and VCS restarts the resource when necessary. The Operations attribute specifies these default aspects of startup.

Symantec recommends starting the vxconfigd daemon with the `syslog` option, which enables logging of debug messages. Note that the SF Sybase CE 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 F-3](#) describes the entry points for the CVMVxconfigd agent.

Table F-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
<code>imf_init</code>	Initializes the agent to interface with the AMF kernel module. This function runs when the agent starts up.
<code>imf_getnotification</code>	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.
<code>imf_register</code>	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 F-4](#) describes the modifiable attributes of the CVMVxconfigd resource type.

Table F-4 CVMVxconfigd agent attribute

Attribute	Description
<code>CVMVxconfigdArgs</code>	<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> <ul style="list-style-type: none"> ■ Type and dimension: keylist

Table F-4 CVMVxconfigd agent attribute (*continued*)

Attribute	Description
IMF	<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. Valid values are as follows: <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 Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: <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 <code>imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the <code>RegisterRetyLimit</code> key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the <code>Mode</code> key changes. Default: 3. ■ Type and dimension: integer-association <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 IMF{} = { Mode=2, MonitorFreq=1, RegisterRetryLimit=3 }
```

```
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 }
)
```

For a more extensive `main.cf` that includes the `CVMVxconfigd` resource:

See [“About sample main.cf files”](#) on page 305.

CVMVolDg agent

The `CVMVolDg` agent represents and controls CVM diskgroups and CVM volumes within the diskgroups. The global nature of CVM diskgroups and volumes requires importing them only once on the CVM master node.

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 Sybase service group. A disk group must be configured to only one Sybase 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 F-5 describes the entry points used by the CVMVolDg agent.

Table F-5 CVMVolDg agent entry points

Entry Point	Description
Online	<p>Starts all volumes in the shared disk group specified by the CVMVolume attribute.</p> <p>Imports the shared disk group from the CVM master node, if the disk group is not already imported.</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>Sets the activation mode of the shared disk group to “off.”</p> <p>If the CVMDeportOnOffline attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.</p>
Monitor	<p>Monitors specified critical volumes in the diskgroup. The CVMVolume attribute specifies these volumes. SF Sybase CE requires specifying at least one volume in a disk group.</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 CFMount 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 F-6 describes the user-modifiable attributes of the CVMVolDg resource type.

Table F-6 CVMVolDg agent attributes

Attribute	Description
CVMDiskGroup (required)	<p>Shared disk group name.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar
CVMVolume (required)	<p>Lists critical volumes in the disk group. SF Sybase CE requires specifying at least one volume in the disk group.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-keylist
CVMActivation (required)	<p>Activation mode for the disk group.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar ■ Default = sw (shared-write) <p>This is a localized attribute.</p>
CVMDeportOnOffline (optional)	<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> <ul style="list-style-type: none"> ■ Type and dimension: integer-scalar ■ Default = 0 <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 <code>CVMDeportOnOffline</code> attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.</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 }
    static str ArgList[] = { CVMDiskGroup, CVMVolume,
        CVMActivation }
```

```

    str CVMDiskGroup
    keylist CVMVolume[]
    str CVMActivation
    temp int voldg_stat
)

```

CVMVolDg agent sample configuration

Each Sybase service group requires a CVMVolDg resource type to be defined. The following is a sample configuration:

```

CVMVolDg cvmvoldg1 (
    Critical = 0
    CVMDiskgroup = testdg
    CVMVolume = { vol1, vol2, mvoll, mvoll2, snapvol, vset1 }
    CVMVolumeIoTest = { snapvol, vset1 }
    CVMActivation @system1 = sw
    CVMActivation @system2 = sw
    CVMDeportOnOffline = 1
)

CVMVolDg sybbindg_101_voldg (
    CVMDiskGroup = sybbindg_101
    CVMVolume = { sybbin_vol }
    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 F-7](#) provides the entry points for the CFSMount agent.

Table F-7 CFSMount agent entry points

Entry Point	Description
Online	Mounts a block device in cluster mode.
Offline	Unmounts the file system, forcing unmount if necessary, and sets primary to secondary if necessary.
Monitor	Determines if the file system is mounted. Checks mount status using the <code>fsclustadm</code> command.
Clean	Generates a null operation for a cluster file system mount.
<code>imf_init</code>	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.
<code>imf_getnotification</code>	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
<code>imf_register</code>	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSMount agent

[Table F-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

Table F-8 CFSMount Agent attributes

Attribute	Description
MountPoint	Directory for the mount point. <ul style="list-style-type: none"> ■ Type and dimension: string-scalar
BlockDevice	Block device for the mount point. <ul style="list-style-type: none"> ■ Type and dimension: string-scalar
NodeList	List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list. <ul style="list-style-type: none"> ■ Type and dimension: string-keylist

Table F-8 CFSMount Agent attributes (*continued*)

Attribute	Description
IMF	<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. <ul style="list-style-type: none"> Valid values are as follows: <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 Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. <ul style="list-style-type: none"> Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: <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 imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. <ul style="list-style-type: none"> Default: 3. ■ Type and dimension: integer-association

Table F-8 CFSMount Agent attributes (*continued*)

Attribute	Description
MountOpt (optional)	<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 style="margin-left: 40px;">ro ro,cluster blkclear,mincache=closesync</pre> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar
Policy (optional)	<p>List of nodes to assume the primaryship of the cluster file system if the primary fails. If set to NULL or if none of the hosts specified in the list is active when the primary fails, a node is randomly selected from the set of active nodes to assume primaryship.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar

CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition:

```
type CFSMount (
    static int IMF{} = { Mode=3, MonitorFreq=1, RegisterRetryLimit=3 }
    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 Sybase service group requires a CFSMount resource type to be defined:

```

CFSMount sybbindg_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbindg/sybbinvol"
    Primary = sys2;
)

```

To see CFSMount defined in a more extensive example:

See [“About sample main.cf files”](#) on page 305.

Process agent

The Process agent starts, stops, and monitors a process that you specify. You can use the agent to make a process highly available or to monitor it.

Agent functions

Online	Starts a process in the background with optional arguments and priority in the specified user context.
Offline	Terminates the process with a SIGTERM. If the process does not exit, a SIGKILL is sent.
Monitor	Checks to see if the process is running by scanning the process table for the name of the executable pathname and argument list.
Clean	Terminates all ongoing resource actions and takes the resource offline, forcibly when necessary.

State definitions

ONLINE	Indicates that the specified process is running in the specified user context.
--------	--

OFFLINE	Indicates that the specified process is not running in the specified user context.
FAULTED	Indicates that the process has terminated unexpectedly.
UNKNOWN	Indicates that the agent can not determine the state of the process.

Attributes

Table F-9 Required attribute

Required attribute	Description
PathName	Complete pathname to access an executable program. This path includes the program name. If a script controls the process, the PathName defines the complete path to the shell. Type and dimension: string-scalar

Table F-10 Optional attributes

Optional attribute	Description
Arguments	Passes arguments to the process. If a script controls the process, the script is passed as an argument. Separate multiple arguments with a single space. A string cannot accommodate more than one space between arguments, nor allow for leading or trailing whitespace characters. Type and dimension: string-scalar

Table F-10 Optional attributes (*continued*)

Optional attribute	Description
	<p>The file that contains the process ID for the monitoring process. Specify the <code>PidFile</code> attribute for the monitoring process to use the Pid. Otherwise, to complete the monitoring process the agent uses the <code>ps</code> output.</p> <p>Note that when you use scripts, or other indirect mechanisms, to start processes, you must set the <code>PidFile</code> attribute if the <code>ps</code> output is different from the configured values for the <code>PathName</code> or <code>Arguments</code> attributes.</p> <p>Type and dimension: string-scalar Example: <code>"/var/lock/sendmail.pid"</code></p>
Priority	<p>Priority that the process runs. Priority values range between -20 (highest) to +19 (lowest).</p> <p>Type and dimension: string-scalar Default: 10</p>
UserName	<p>This attribute is the owner of the process. The process runs with the user ID.</p> <p>Type and dimension: string-scalar Default: root</p>

Resource type definition

```

type Process (
  static keylist SupportedActions = { "program.vfd", getcksum }
  static str ArgList[] = { PathName, Arguments, UserName,
    Priority, PidFile }
  str PathName
  str Arguments
  str UserName = root
  str Priority = 10
  str PidFile
)

```

Sample configurations

```
Process vxfend (  
    PathName = "/sbin/vxfend"  
    Arguments = "-m sybase -k /tmp/vcmp_socket"  
)
```

Monitoring options for the Sybase agent

The Veritas agent for Sybase provides two levels of application monitoring: basic and detail.

In the basic monitoring mode, the agent for Sybase monitors the Sybase daemon processes to verify whether they are running.

For Sybase cluster edition, the agent uses `qrmutil` utility that Sybase provides to get the status of the Sybase instance. If the state returned by `qrmutil` utility is 'failure pending', the agent panics the node. When the Sybase agent detects that the configured Sybase server is not running on a system, based on the value of the `OnlineRetryLimit` attribute of the Sybase service group, the service group is restarted on the same system on which the group faulted.

For example:

```
# qrmutil --quorum_dev=/quorum/quorum.dat --monitor=ase1  
Executing 'monitor' command for instance 'ase1'  
Instance 'ase1' has a failure pending.  
# echo $?  
99
```

In this example instance 'ase1' has a failure pending state. The agent will panic the node running the instance 'ase1'. The node will automatically rejoin the cluster after reboot.

In the detail monitoring mode, the agent performs a transaction on a test table in the database to ensure that Sybase functions properly. The agent uses this test table for internal purposes. Symantec recommends that you do not perform any other transaction on the test table.

See [“About setting up detail monitoring for the agentfor Sybase”](#) on page 175.

Sybase resource type

The type definitions and attribute definitions for the Sybase resource type are described as follows.

Type definition for the Sybase agent

The resource type definition for the agent for Sybase is as follows.

```
type Sybase (
    static boolean AEPTIMEOUT = 1
    static keylist SupportedActions = { "checkpoint_all" }
    str Server
    str Owner
    str Home
    str Version
    str SA
    str SAPswd
    str Run_ServerFile
    int DetailMonitor = 0
    str User
    str UPword
    str Db
    str Table
    str Monscript = "/opt/VRTSagents/ha/bin/Sybase/SqlTest.pl"
    boolean WaitForRecovery = 0
    str Quorum_dev
    str interfaces_File
    int ShutdownWaitLimit = 60
    int DelayAfterOnline = 10
    int DelayAfterOffline = 2
    static int ToleranceLimit = 1
    static str ArgList[] = { Server, Owner, Home, Version, SA,
        SAPswd, User, UPword, Db, Table, Monscript, DetailMonitor,
        WaitForRecovery, Run_ServerFile, Quorum_dev, State,
        interfaces_File, ShutdownWaitLimit, DelayAfterOnline,
        DelayAfterOffline }
    static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }
    static str IMFRegList[] = { Server, Owner, Quorum_dev }
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Sybase"
)
```

Attribute definitions for the Sybase agent

Review the description of the Sybase agent attributes. The agent attributes are classified as required, optional, and internal.

[Table F-11](#) lists the required attributes.

Table F-11 Required attributes

Required Attributes	Definition
Server	<p>The \$DSQUERY ASE name. Only one server should be configured in a Sybase service group. The advantage of configuring Sybase resources in a separate service group is, each Sybase data server can failover independently.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>
Owner	<p>Sybase user as the defined owner of executables and database files in any of the sources (such as NIS+, /etc/hosts, and so on) specified in the /etc/nsswitch.conf file for passwd entry. The Sybase executables and database files are accessed in the context of this user.</p> <p>Type and dimension: string-scalar</p>
Home	<p>The \$SYBASE path to Sybase binaries and configuration files.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>
Version	<p>Version of Sybase ASE.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p> <p>Caution: Once the Sybase resource is online in VCS, you must not modify the Home and Version attributes. For the Sybase cluster edition, setting invalid values for Home and Version attributes when the resource is in Online state causes the node to panic.</p>
SA	<p>Sybase database administrator. This attribute is required to connect to the ASE for shutdown.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>

Table F-11 Required attributes (*continued*)

Required Attributes	Definition
SApswd	<p>Encrypted password for Sybase database administrator. This password is required to connect to the ASE for shutdown.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p> <p>See “Encrypting passwords for Sybase” on page 175.</p> <p>Note: You need not specify a value for this attribute if the SA user does not require a password.</p>
Quorum_dev	<p>The quorum device manages the cluster membership, stores cluster configuration data, and contains information shared among server instances and nodes. The quorum device is a disk that is accessible to all the nodes in the cluster. Specify a fully qualified quorum device name.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p> <p>Note: This attribute should be specified only for the cluster edition.</p> <p>Caution: If you do not set the Quorum_dev attribute for Sybase Cluster Edition, the Sybase agent does not perform the qrmutil-based checks. This error in configuration may lead to undesirable results. For example, if qrmutil returns failure pending, the agent does not panic the system. Thus, the Sybase agent does not perform qrmutil-based checks because the Quorum_dev attribute is not set.</p>

[Table F-12](#) lists the optional attributes.

Table F-12 Optional attributes

Optional Attributes	Definition
DetailMonitor	<p>Specifies whether the Sybase server is monitored in detail. A positive integer value indicates that the resource monitors the Sybase server in detail. Value 0 denotes it does not. Default is 0.</p> <p>Type and dimension: int-scalar</p> <p>Note: The DetailMonitor attribute is deprecated in SF Sybase CE 6.0.1. Instead, LevelTwoMonitorFreq attribute of Sybase agent may be used. The default value of LevelTwoMonitorFreq attribute is 0 (zero).</p>
User	<p>The database user, in the context of which, the transactions are performed on the database. You must specify a value for this attribute if LevelTwoMonitorFreq is set to a positive integer value.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>
UPword	<p>Encrypted password for the database user. You must specify a value for this attribute if LevelTwoMonitorFreq is set to a positive integer value. However, you need not specify a value for this attribute if the database user does not require a password.</p> <p>See “Encrypting passwords for Sybase” on page 175.</p> <p>intercType and dimension: string-scalar</p> <p>Default value: No default value</p>
Db	<p>Name of the database used for detailed monitoring. The table used by the detail monitor script resides in this database. You must specify a value for this attribute if LevelTwoMonitorFreq is set to a positive integer value.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>
Table	<p>Name of the table on which the detail monitoring script performs the transactions. You must specify a value for this attribute if LevelTwoMonitorFreq is set to a positive integer value.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
Monscript	<p>The path to the detail monitor script; the default value for this attribute is the path for the script, SqlTest.pl, provided with the agent. You must specify a value for this attribute if LevelTwoMonitorFreq is set to a positive integer value.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p> <p>Note: By default, SqlTest.pl script has the execute permission set. If you specify custom detail monitor script, ensure that custom detail monitor script also has the execute permissions set.</p>
Run_ServerFile	<p>Specifies the location of the RUN_SERVER file for the Sybase instance. The default location of this file is used if no value is specified for this attribute.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p>

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
IMF	

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
	<p>This resource-type level attribute determines whether the Sybase 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: 3</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: 5</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 <ul style="list-style-type: none"> ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>sybase_imf_register</code> agent function to register the resource with the AMFkernel 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. <p>Default: 3</p>

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
	Type and dimension: Integer-association.
interfaces_File	<p>Specifies the location of interfaces file, including the directory name and the file name for the Sybase instance. If this attribute is configured, [-I interfaces file] option is used when connecting to the isql session. If this attribute is not configured, the agent does not use the -I option.</p> <p>Type and dimension: string-scalar</p> <p>Default value: No default value</p> <p>For example: /sybase/my_interfaces_file</p> <p>Note: It is assumed that you have modified the RUN_ServerFile with the non-default interface file location if the interfaces_File attribute is configured.</p>
DelayAfterOnline	<p>Specifies the number of seconds that elapse after the Online entry point is complete and before the next monitor cycle is invoked.</p> <p>Type and dimension: integer-scalar</p> <p>Default value: 10</p>
DelayAfterOffline	<p>Specifies the number of seconds that elapse after the Offline entry point is complete and before the next monitor cycle is invoked.</p> <p>Type and dimension: integer-scalar</p> <p>Default value: 2</p>
ShutdownWaitLimit	<p>Maximum number of seconds for which the agent waits for the Sybase instance to stop after issuing the shutdown with wait command, and before attempting to issue the kill -15 <data server-pid> command, if required.</p> <p>Type and dimension: integer-scalar</p> <p>Default value: 60</p>

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
Quorum_dev	<p>The quorum device manages the cluster membership, stores cluster configuration data and contains information shared among server instances and nodes. It must be a disk accessible to all nodes in the cluster. Specify fully qualified quorum device name.</p> <p>Caution: If you do not set the Quorum_dev attribute for Sybase Cluster Edition, the Sybase agent does not perform the qrmutil-based checks. This error in configuration may lead to undesirable results. For example, if qrmutil returns failure pending, the agent does not panic the system. Thus, the Sybase agent does not perform qrmutil-based checks because the Quorum_dev attribute is not set.</p> <p>Type and dimension: String-scalar</p> <p>Default value: No default value</p>
Run_ServerFile	<p>Specifies the location of the RUN_SERVER file of the Sybase instance. The default location of the file is used if no value is specified for this attribute.</p> <p>Type and dimension: String-scalar</p> <p>Default value: No default value</p>

[Table F-13](#) lists the internal attribute for Sybase agent.

This attribute is for internal use only. Symantec recommends not to modify the value of this attribute.

Table F-13 Internal attribute

Internal attribute	Definition
AgentDirectory	<p>Specifies the location of the binaries, scripts, and other files related to the agent for Sybase.</p> <p>Type and dimension: static-string</p>

Compatibility issues when installing Storage Foundation for Sybase ASE CE with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present](#)

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

Installing Storage Foundation when other Veritas products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

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

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host RPMs as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

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

If you plan to install or upgrade Storage Foundation on systems where ApplicationHA has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not consider VCS as an installed product even though it uses the bundled VRTSvcS RPM.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not allow the installation or upgrade for products that use VCS. The following products cannot be installed or upgraded: VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SFCFSRAC or SFSYBASECE.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer allows the installation or upgrade of VM, FS, SF, or DMP.
- When you uninstall Storage Foundation products where ApplicationHA is present, the installer does not uninstall VRTSsfmh.

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

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx and VRTSisco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx, VRTSisco, and VRTSat.

Index

A

- adding
 - users 81
- agent for SQL server
 - attribute definitions 347
 - resource type 346
 - type definition 346
- agents
 - about 329
 - CFSMount 339
 - CVMCluster 331
 - CVMVolDg 336
 - CVMVxconfigd 333
 - of VCS 330
- attributes
 - about agent attributes 329
 - CFSMount agent 340
 - CVMCluster agent 331
 - CVMVolDg agent 331, 337
 - CVMVxconfigd agent 334

B

- backup boot disk group 159–160
 - rejoining 159

C

- cables
 - cross-over Ethernet 223
- CFS
 - stopping applications 261
 - unmounting file systems 261
- CFSMount agent 339
 - attributes 340
 - entry points 339
 - sample configuration 342–343
 - type definition 342
- CFSTypes.cf 342
- cluster
 - removing a node from 242
 - verifying operation 106

- Cluster Manager 24
- clusters
 - basic setup 26
 - four-node illustration 26
- commands
 - hastatus 106
 - hasys 107
 - lltstat 102
 - vxdisksetup (initializing disks) 89
- configuration
 - required information
 - for SF Sybase CE, 287
- configuration files
 - removing 267
 - See also* main.cf samples
- configuring
 - ssh 53
- configuring VCS
 - adding users 81
 - event notification 82, 84
 - global clusters 86
- CVM
 - CVMTypes.cf file 332
 - upgrading protocol version 161
- CVMCluster agent 331
 - attributes 331
 - entry points 331
 - sample configuration 333
 - type definition 332
- CVMTypes.cf
 - definition, CVMCluster agent 332
 - definition, CVMVolDg agent 338
 - definition, CVMVxconfigd agent 335
- CVMVolDg agent 336
 - attributes 337
 - entry points 337
 - sample configuration 339
 - type definition 338
- CVMVxconfigd agent 333
 - attributes 334
 - CVMTypes.cf 335
 - entry points 333

CVMVxconfigd agent (*continued*)
 sample configuration 336
 type definition 335

D

detail monitoring
 disabling 177
 enabling 176
 disks
 adding and initializing 89
 testing with vxfcntlshdw 90
 verifying node access 92

E

environment variables
 MANPATH 56
 Ethernet controllers 223

F

files. *See* configuration files

G

GAB
 port memberships 104
 GAB ports 105
 Global Cluster Option (GCO)
 overview 29
 global clusters
 about 22
 configuration 86

H

hastatus -summary command 106
 hasys -display command 107
 hubs
 independent 223

I

I/O fencing
 checking disks 90
 shared storage 90
 installation
 pre-installation tasks
 mounting product disc 55
 workflow 52
 preparation 287
 installation worksheets 287

installsfsybasece
 installing SF Sybase CE 65
 upgrading SF Sybase CE 130

J

Java Console 24

K

kernel.panic tunable
 setting 56

L

LLT
 interconnects 57
 verifying 102
 llstat command 102

M

MANPATH environment variable 56
 media speed 57
 optimizing 57
 monitoring
 basic 346
 detail 346

N

nodes
 adding Sybase ASE CE nodes
 configuring GAB 226
 configuring LLT 226
 configuring VXFEN 226
 starting Volume Manager 225
 preparing Sybase ASE CE nodes
 about 237
 configuring CVM 231
 creating Sybase user and groups 238
 preparing Sybase resource mount
 points 238
 removing a node from a cluster
 tasks 241
 removing nodes
 GAB configuration 244
 LLT configuration 244
 modifying VCS configuration 245

O

- optimizing
 - media speed 57
- options
 - SF Sybase CE configuration 69

P

- PATH variable
 - VCS commands 102
- ports
 - GAB 105

R

- rejoining
 - backup boot disk group 159
- removing a node from a cluster
 - editing VCS configuration files 243
 - procedure 242
 - tasks 241
- response file
 - about 189
 - syntax 190
- rolling upgrade 156
 - versions 151
- rsh 53

S

- sample configuration files. *See* main.cf samples
- setting
 - kernel.panic tunable 56
- SF Sybase CE
 - about 19
 - high-level view 26
- SF Sybase CE configuration
 - of components 70
- SF Sybase CE configuration
 - about 69
 - options 69
 - preparation
 - worksheets 287
 - required information 287
- SF Sybase CE installation
 - pre-installation tasks
 - setting MANPATH 56
 - setting up shared storage 55
 - synchronizing time settings 53
 - verifying systems 58
 - preinstallation information 31

SF Sybase CE installation (*continued*)

- preparation
 - worksheets 287
- requirements
 - hardware 32
- using installsfsybasece 65
- verifying
 - cluster operations 102
 - GAB operations 102
 - LLT operations 102
- SF Sybase CE uninstallation
 - preparation
 - stopping applications, CFS 261
 - stopping applications, VxFS 262
 - stopping Sybase instances 258
 - stopping VCS 262
 - uninstalling 259
 - unmounting CFS file systems 261
 - unmounting VxFS file systems 263
 - removing configuration files 267
 - removing RPMs 265
 - using uninstallsfsybasece 265
 - workflow 255
- SF Sybase CE upgrade
 - post-upgrade tasks
 - upgrading CVM protocol version 161
 - upgrading disk group version 161
 - preparation 126, 154
 - restoring configuration files 131
 - stopping cluster resources 126, 154
 - stopping Sybase ASE CE 126
 - using installsfsybasece 130
- shared storage
 - setting up 55
- SMTP email notification 82
- SNMP trap notification 84
- SQL server agent attributes
 - AgentDirectory 355
 - Db 350
 - DelayAfterOffline 354
 - DelayAfterOnline 354
 - DetailMonitor 350
 - home 348
 - monscript 351
 - owner 348
 - Quorum_dev 349
 - Run_ServerFile 351
 - SA 348
 - SAPswd 349

- SQL server agent attributes *(continued)*
 - server 348
 - ShutdownWaitLimit 354
 - table 350
 - UPword 350
 - user 350
 - version 348
- ssh 53
 - configuring 53
- Storage Foundation for Sybase ASE CE configuration
 - verifying 104
- Sybase
 - language settings 173
 - setting up for detail monitoring 173
 - stopping instances 258
- Sybase agent
 - detail monitoring 175
 - monitoring options 346
- Sybase ASE CE
 - pre-installation
 - setting up storage 37
- synchronizing time settings, before installing 53
- system communication using rsh
 - ssh 53
- system state attribute value 106

T

- tunables file
 - about setting parameters 293
 - parameter definitions 298
 - preparing 297
 - setting for configuration 294
 - setting for installation 294
 - setting for upgrade 294
 - setting parameters 297
 - setting with no other operations 295
 - setting with un-integrated response file 296
- type definition 347

U

- uninstallation
 - of SF Sybase CE 265
- uninstallsybasece
 - removing RPMs 265
- unsuccessful upgrade 160
- upgrade
 - stopping Sybase ASE CE 154

- upgrading
 - rolling 156

V

- VCS
 - command directory path variable 102
 - notifications 22
 - stopping 262
- VCS notifications
 - SMTP notification 22
 - SNMP notification 22
- Veritas File System
 - stopping applications 262
 - unmounting 263
- Veritas Operations Manager 24
- VVR
 - about 23
 - global cluster overview 250
- vxdisksetup command 89
- VxFS. *See* Veritas File System

W

- worksheets
 - for SF Sybase CE 287