

Storage Foundation for Sybase ASE CE 7.1 Configuration and Upgrade Guide - Solaris

Storage Foundation for Sybase ASE CE Configuration and Upgrade Guide

Last updated: 2016-04-28

Document version: 7.1 Rev 0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Configuring SF Sybase ASE CE	12
Chapter 1	Preparing to configure SF Sybase CE	13
	Supported Sybase ASE CE releases	13
	Supported SF Sybase CE configurations	13
	Coordinator disk requirements for I/O fencing	14
	Supported replication technologies for global clusters	14
	About planning to configure I/O fencing	14
	Typical SF Sybase CE cluster configuration with disk-based I/O fencing	16
	Planning for cluster management	16
Chapter 2	Configuring SF Sybase CE	18
	About configuring SF Sybase CE	18
	Configuring the SF Sybase CE components using the script-based installer	19
	Configuring the SF Sybase CE cluster	20
Chapter 3	Configuring SF Sybase CE clusters for data integrity	38
	Setting up disk-based I/O fencing using installer	38
	Initializing disks as VxVM disks	38
	Identifying disks to use as coordinator disks	39
	Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installer	39
	Checking shared disks for I/O fencing	41
	Configuring disk-based I/O fencing using installer	45
Chapter 4	Performing an automated SF Sybase CE configuration	48
	Configuring SF Sybase CE using response files	48
	Response file variables to configure SF Sybase CE	49
	Sample response files for configuring SF Sybase CE	59

Chapter 5	Performing an automated I/O fencing configuration using response files	62
	Configuring I/O fencing using response files	62
	Response file variables to configure disk-based I/O fencing	63
	Sample response file for configuring disk-based I/O fencing	65
Chapter 6	Configuring a Sybase cluster under VCS control using a response file	67
	Configuring a Sybase cluster under VCS control with a response file	67
	Response file variables to configure SF Sybase CE in VCS	68
Section 2	Post-installation and configuration tasks	71
Chapter 7	Verifying the installation	72
	Upgrading the disk group version	72
	Performing a postcheck on a node	73
	Verifying SF Sybase CE installation using VCS configuration file	73
	Verifying LLT, GAB, and cluster operation	74
	Verifying LLT	74
	Verifying GAB	76
	Verifying the cluster	78
	Verifying the cluster nodes	79
Chapter 8	Performing additional post-installation and configuration tasks	82
	Installing language packages	82
	About enabling LDAP authentication for clusters that run in secure mode	83
	Enabling LDAP authentication for clusters that run in secure mode	84
	Configuring Volume Replicator	88
	Running SORT Data Collector to collect configuration information	89

Section 3	Upgrade of SF Sybase CE	90
Chapter 9	Planning to upgrade SF Sybase CE	91
	About the upgrade	91
	Supported upgrade paths	92
	Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches	93
Chapter 10	Performing a full upgrade of SF Sybase CE using the product installer	96
	About full upgrades	96
	Preparing to perform a full upgrade to SF Sybase CE 7.1	97
	Upgrading to SF Sybase CE 7.1	99
	Upgrading SF Sybase CE using the Veritas installation program	100
Chapter 11	Performing an automated full upgrade of SF Sybase CE using response files	103
	Upgrading SF Sybase CE using a response file	103
	Response file variables to upgrade SF Sybase CE	104
	Sample response file for upgrading SF Sybase CE	107
Chapter 12	Performing a phased upgrade of SF Sybase CE	109
	About phased upgrade	109
	Performing a phased upgrade of SF Sybase CE from version 6.0 and later release	110
	Step 1: Performing pre-upgrade tasks on the first half of the cluster	111
	Step 2: Upgrading the first half of the cluster	114
	Step 3: Performing pre-upgrade tasks on the second half of the cluster	116
	Step 4: Performing post-upgrade tasks on the first half of the cluster	118
	Step 5: Upgrading the second half of the cluster	118
	Step 6: Performing post-upgrade tasks on the second half of the cluster	119

Chapter 13	Performing a rolling upgrade of SF Sybase CE	121
	About rolling upgrade	121
	Supported rolling upgrade paths	123
	Preparing to perform a rolling upgrade to SF Sybase CE 7.1	123
Chapter 14	Upgrading SF Sybase CE using Live Upgrade	125
	About Live Upgrade	125
	Supported upgrade paths for Live Upgrade and Boot Environment upgrade	126
	Before you upgrade SF Sybase CE using Solaris Live Upgrade	126
	Performing Live Upgrade on Solaris 10 systems	127
	Creating a new Solaris 10 boot environment on the alternate boot disk	128
	Upgrading SF Sybase CE using the installer for Solaris 10 Live Upgrade	133
	Completing the Solaris 10 Live Upgrade	134
	Verifying the Solaris 10 Live Upgrade of SF Sybase CE	135
	Administering boot environments in Solaris 10 Live Upgrade	136
	Reverting to the primary boot environment on a Solaris 10 system	136
	Switching the boot environment for Solaris 10 SPARC	137
	Upgrading the operating system and SF Sybase CE using Live Upgrade	138
	Upgrading SF Sybase CE only using Live Upgrade	139
	Upgrading Solaris only using Live Upgrade	139
	Completing the Live Upgrade	140
Chapter 15	Performing post-upgrade tasks	144
	Re-joining the backup boot disk group into the current disk group	144
	Reverting to the backup boot disk group after an unsuccessful upgrade	145
	Setting or changing the product license level	145
	Upgrading disk layout versions	146
	Upgrading CVM protocol version and VxVM disk group version	147

Section 4	Installation and upgrade of Sybase ASE CE	148
Chapter 16	Installing, configuring, and upgrading Sybase ASE CE	149
	Before installing Sybase ASE CE	149
	Preparing for local mount point on VxFS for Sybase ASE CE binary installation	150
	Preparing for shared mount point on CFS for Sybase ASE CE binary installation	151
	Installing Sybase ASE CE software	152
	Preparing to create a Sybase ASE CE cluster	152
	Creating the Sybase ASE CE cluster	154
	Preparing to configure the Sybase instances under VCS control	154
	Language settings for the Sybase agent	155
	Configuring Sybase for detail monitoring	155
	Encrypting passwords for Sybase	157
	About setting up detail monitoring for the agent for Sybase	158
	Configuring a Sybase ASE CE cluster under VCS control using the SF Sybase CE installer	160
	Upgrading Sybase ASE CE	168
Section 5	Adding and removing nodes	169
Chapter 17	Adding a node to SF Sybase CE clusters	170
	About adding a node to a cluster	170
	Before adding a node to a cluster	171
	Adding the node to a cluster manually	174
	Starting Veritas Volume Manager (VxVM) on the new node	174
	Configuring cluster processes on the new node	175
	Setting up the node to run in secure mode	177
	Starting fencing on the new node	180
	Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node	180
	After adding the new node	182
	Configuring the ClusterService group for the new node	182
	Adding a node to a cluster using the Veritas InfoScale installer	183
	Adding the new instance to the Sybase ASE CE cluster	186
	Creating Sybase user and groups	186

	Preparing the mount point for Sybase resources on the new node	186
	Adding a new Sybase ASE CE instance to the Sybase ASE CE cluster	187
	Bringing the new Sybase ASE CE instance under VCS control	188
Chapter 18	Removing a node from SF Sybase CE clusters	190
	About removing a node from a cluster	190
	Removing a node from a cluster	191
	Modifying the VCS configuration files on existing nodes	192
	Modifying the Cluster Volume Manager (CVM) configuration on the existing nodes to remove references to the deleted node	195
	Removing security credentials from the leaving node	195
Section 6	Configuration of disaster recovery environments	196
Chapter 19	Configuring disaster recovery environments	197
	Disaster recovery options for SF Sybase CE	197
	About setting up a global cluster environment for SF Sybase CE	198
	About configuring a parallel global cluster using Volume Replicator (VVR) for replication	198
Section 7	Installation reference	200
Appendix A	Installation scripts	201
	Installation script options	201
	About using the postcheck option	207
Appendix B	Sample installation and configuration values	210
	SF Sybase CE installation and configuration information	210
	SF Sybase CE worksheet	210

Appendix C	Tunable files for installation	215
	About setting tunable parameters using the installer or a response file	215
	Setting tunables for an installation, configuration, or upgrade	216
	Setting tunables with no other installer-related operations	217
	Setting tunables with an un-integrated response file	218
	Preparing the tunables file	219
	Setting parameters for the tunables file	220
	Tunables value parameter definitions	220
Appendix D	Configuration files	228
	About sample main.cf files	228
	Sample main.cf files for Sybase ASE CE configurations	228
	Sample main.cf for a basic Sybase ASE CE cluster configuration under VCS control with shared mount point on CFS for Sybase binary installation	229
	Sample main.cf for a basic Sybase ASE CE cluster configuration with local mount point on VxFS for Sybase binary installation	235
	Sample main.cf for a primary CVM VVR site	239
	Sample main.cf for a secondary CVM VVR site	246
Appendix E	Configuring the secure shell or the remote shell for communications	253
	About configuring secure shell or remote shell communication modes before installing products	253
	Manually configuring passwordless ssh	254
	Setting up ssh and rsh connection using the installer -comsetup command	258
	Setting up ssh and rsh connection using the pwdutil.pl utility	259
	Restarting the ssh session	262
	Enabling and disabling rsh for Solaris	263
Appendix F	High availability agent information	265
	About agents	265
	VCS agents included within SF Sybase CE	266
	VCS agent for Sybase included within SF Sybase CE	266
	CVMCluster agent	267
	Entry points for CVMCluster agent	267
	Attribute definition for CVMCluster agent	267

CVMCluster agent type definition	268
CVMCluster agent sample configuration	269
CVMVxconfigd agent	269
Entry points for CVMVxconfigd agent	270
Attribute definition for CVMVxconfigd agent	270
CVMVxconfigd agent type definition	271
CVMVxconfigd agent sample configuration	272
CVMVolDg agent	272
Entry points for CVMVolDg agent	272
Attribute definition for CVMVolDg agent	273
CVMVolDg agent type definition	274
CVMVolDg agent sample configuration	275
CFSMount agent	275
Entry points for CFSMount agent	276
Attribute definition for CFSMount agent	276
CFSMount agent type definition	278
CFSMount agent sample configuration	279
Process agent	279
Agent functions	279
State definitions	280
Attributes	280
Resource type definition	281
Sample configurations	281
Monitoring options for the Sybase agent	281
Sybase resource type	282
Type definition for the Sybase agent	282
Attribute definitions for the Sybase agent	283
Index	290

Configuring SF Sybase ASE CE

- Chapter 1. Preparing to configure SF Sybase CE
- Chapter 2. Configuring SF Sybase CE
- Chapter 3. Configuring SF Sybase CE clusters for data integrity
- Chapter 4. Performing an automated SF Sybase CE configuration
- Chapter 5. Performing an automated I/O fencing configuration using response files
- Chapter 6. Configuring a Sybase cluster under VCS control using a response file

Preparing to configure SF Sybase CE

This chapter includes the following topics:

- [Supported Sybase ASE CE releases](#)
- [Supported SF Sybase CE configurations](#)
- [Coordinator disk requirements for I/O fencing](#)
- [Supported replication technologies for global clusters](#)
- [About planning to configure I/O fencing](#)
- [Planning for cluster management](#)

Supported Sybase ASE CE releases

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

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

https://www.veritas.com/support/en_US/article.DOC8892

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.

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 must be DMP devices.
- Each of the coordinator disks must use a physically separate disk or LUN. Veritas 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.
- Coordinator devices can be attached over iSCSI protocol but they must be DMP devices and must support SCSI-3 persistent reservations.
- Veritas 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.
- The coordinator disk size must be at least 128 MB.

Supported replication technologies for global clusters

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


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.

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

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

Figure 1-1 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 installer | See “Setting up disk-based I/O fencing using installer” on page 38. |
| Using response files | See “Response file variables to configure disk-based I/O fencing” on page 63.
See “Configuring I/O fencing using response files” on page 62. |
| Manually editing configuration files | |

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

Figure 1-2 displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.

Figure 1-2 Typical SF Sybase CE cluster configuration with disk-based I/O fencing



Planning for cluster management

Table 1-1 lists the various agents supported in SF Sybase CE installations for effective cluster management.

Table 1-1 List of agents

Agent	Description
VCS agent for Sybase	Sybase database management 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.
VCS agents for CVM	Volume management An SF Sybase CE configuration automatically configures the CVMCluster resource and the CVMVxconfigd resource. You must configure the CVMVolDg agent for each shared disk group.

Table 1-1 List of agents *(continued)*

Agent	Description
VCS agents for CFS	File system management If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.
VCS process agent for vxfsend	vxfsend process/daemon management

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 Veritas InfoScale Enterprise 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

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

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

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

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

Note: If you want to reconfigure SF 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.

Common product installer

Run the program:

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

Choose **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 Cluster Server component to set up the SF Sybase CE cluster.

See [“Configuring the SF Sybase CE cluster”](#) on page 20.

6 Add VCS users.

See [“Adding VCS users”](#) on page 32.

7 Configure SMTP email notification.

See [“Configuring SMTP email notification”](#) on page 33.

8 Configure SNMP trap notification.

See [“Configuring SNMP trap notification”](#) on page 35.

Configuring the SF Sybase CE cluster

Configure the systems on which you installed Veritas InfoScale Enterprise 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 19.
- 2** Select the **Configure Cluster File System** option from the main menu.
Press Enter to continue.
If there are any SF Sybase CE processes running, these processes are stopped.
Press Enter 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 21.
 - Configure private heartbeat links.
See [“Configuring private heartbeat links”](#) on page 21.

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 LLT over UDP (User Datagram Protocol). Veritas 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.

You must not configure LLT heartbeat using the links that are part of aggregated links. For example, link1, link2 can be aggregated to create an aggregated link, aggr1. You can use aggr1 as a heartbeat link, but you must not use either link1 or link2 as heartbeat links.

The following procedure helps you configure LLT heartbeat links.

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 LLT over UDP.
 - Option 1: Configure the heartbeat links using 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: Configure the heartbeat links using 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.
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.

Answer the installer prompts. The following example shows different NICs based on architecture:

- For Solaris SPARC:
You must not enter the network interface card that is used for the public network (typically bge0.)

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] bge1
```

```
Would you like to configure a second private heartbeat link?
```

```
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on sys1:
[b,q,?] bge2
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
```

- For Solaris x64:

You must not enter the network interface card that is used for the public network (typically e1000g0.)

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] e1000g1
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on sys1:
[b,q,?] e1000g2
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 or step 5 for option 3.

- 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 InfoScale Operations Manager, or to specify in the RemoteGroup resource.

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

See the *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: bge0
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[b,q,?] (bge0)
```

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

Do one of the following:

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

```
Is bge0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

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

See [“Setting up trust relationships for your SF Sybase CE cluster”](#) on page 27.

See [“Configuring a secure cluster node by node”](#) on page 28.

Configuring SF Sybase 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.

To configure SF Sybase CE in secure mode

1 To install and configure SF Sybase CE in secure mode, run the command:

```
# ./installer -security
```

2 The installer displays the following question before the installer stops the product processes:

- Do you want to grant read access to everyone? [y,n,q,?]
 - To grant read access to all authenticated users, type *y*.
 - To grant usergroup specific permissions, type *n*.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
 - To specify usergroups and grant them read access, type *y*
 - To grant read access only to root users, type *n*. The installer grants read access read access to the root users.

- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]
- 3** 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.

Setting up trust relationships for your SF Sybase CE cluster

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

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

To set up a trust relationship

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

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

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

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

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

Input the broker port: (14545)

Specify a port number on which broker is running or press Enter to accept the default port.

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

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

4 The installer performs one of the following actions:

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

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

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

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

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

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

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

Configuring a secure cluster node by node

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

[Table 2-1](#) lists the tasks that you must perform to configure a secure cluster.

Table 2-1 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 28.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 29.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 30.

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/installer -securityonnode
```

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

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

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

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

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

Warning: All 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/installer -securityonnode
```

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

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

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

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

```
Select the option you would like to perform [1-2,q.?] 2
Enter the security conf file directory: [b]
```

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** To grant access to all users, add or modify `SecureClus=1` and `DefaultGuestAccess=1` in the cluster definition.

For example:

To grant read access to everyone:

```
Cluster clus1 (
  SecureClus=1
  DefaultGuestAccess=1
)
```

Or

To grant access to only root:

```
Cluster clus1 (
  SecureClus=1
)
```

Or

To grant read access to specific user groups, add or modify `SecureClus=1` and `GuestGroups={}` to the cluster definition.

For example:

```
cluster clus1 (
  SecureClus=1
  GuestGroups={staff, guest}
```

- 5 Modify `/etc/VRTSvcs/conf/config/main.cf` file on the first node, and add `-secure` to the WAC application definition if GCO is configured.

For example:

```
Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = {" /opt/VRTSvcs/bin/wac -secure"}
    RestartLimit = 3
)
```

- 6 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

- 7 On the first node, start VCS. Then start VCS on the remaining nodes.

```
# /opt/VRTSvcs/bin/hastart
```

- 8 On all nodes, start CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer
```

- 9 On the first node, unfreeze the service groups.

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

```
# /opt/VRTSvcs/bin/hagrp -list Frozen=1
```

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

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

Adding VCS users

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

To add VCS users

- 1 Review the required information to add VCS users.

- 2 Reset the password for the Admin user, if necessary.

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

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

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

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

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

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

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

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

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

Configuring SMTP email notification

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

Refer to the *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.

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 35.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server's host name.

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

```
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 *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.

See [“Configuring global clusters”](#) on page 36.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the SNMP trap daemon port.

```
Enter the SNMP trap daemon port: [b,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.

```
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 *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.

If you skip this option, the installer proceeds to configure VCS based on the configuration details you provided.

3 Provide information to configure this cluster as global cluster.

The installer prompts you for a NIC, a virtual IP address, and value for the netmask.

4 Verify and confirm the configuration of the global cluster.

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

```
NIC: bge0
```

```
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 InfoScale™ 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 installer](#)

Setting up disk-based I/O fencing using installer

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

On Solaris 10 x64, disk-based I/O fencing is to be used when Storage Foundation and InfoScale Availability products co-exist.

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:

```
# vxdisk list
```

- 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 *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 c2t13d0
```

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

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 38.

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 41.

Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installer

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for disk-based I/O fencing using the installer

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

```
# /opt/VRTS/install/installer -fencing
```

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

Note down the location of log files that you can access if there is a 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 the remote nodes and checks whether SF Sybase CE 7.1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-6,q]
```

- 4 Ensure that the disk group constitution that is used by the fencing module contains the same disks that are currently used as coordination disks.

- 5 Verify the coordination points.

```
For example,  
Disk Group: fendg  
Fencing disk policy: dmp  
Fencing disks:  
    emc_clariion0_62  
    emc_clariion0_65  
    emc_clariion0_66
```

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

```
Successfully completed the vxfsenwap operation
```

The keys on the coordination disks are refreshed.

- 6 Do you want to send the information about this installation to us to help improve installation in the future? [y,n,q,?] **(y)**.
- 7 Do you want to view the summary file? [y,n,q] **(n)**.

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 `vxfcntlshdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

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

The `vxfcntlshdw` 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 *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 41.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 42.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfcntlshdw utility”](#) on page 43.

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 Veritas technical support.

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

The following output is a sample:

```
# vxddladm listsupport all
```

LIBNAME	VID	PID
libvx3par.so	3PARdata	VV
libvxCLARiiON.so	DGC	All
libvxFJTSYe6k.so	FUJITSU	E6000
libvxFJTSYe8k.so	FUJITSU	All
libvxap.so	Oracle	All
libvxatf.so	VERITAS	ATFNODES
libvxcompellent.so	COMPELNT	Compellent Vol
libvxcopan.so	COPANSYS	8814, 8818

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the vxfsthdw 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 Veritas InfoScale Enterprise.
- 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/rdisk/c1t1d0s2` path on node A and the `/dev/rdisk/c2t1d0s2` path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdisk/c1t1d0s2
```

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

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/rdisk/c2t1d0s2` path.

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

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

Testing the disks using `vxfcntlsthaw` utility

This procedure uses the `/dev/rdisk/c1t1d0s2` disk in the steps.

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

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

```
The disk /dev/rdisk/c1t1d0s2 is ready to be configured for I/O Fencing on
node sys1
```

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

To test the disks using vxfststhdw utility

- 1 Make sure system-to-system communication functions properly.
See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 253.
- 2 From one node, start the utility.
- 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 Review the output as the utility performs the checks and reports its activities.
- 5 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/rdisk/ctl0s2 have PASSED  
The disk is now ready to be configured for I/O fencing on node  
sys1
```

- 6 Run the vxfststhdw utility for each disk you intend to verify.

Note: Only dmp disk devices can be used as coordinator disks.

Configuring disk-based I/O fencing using installer

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 installer

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

```
# /opt/VRTS/install/installer -fencing
```

The installer 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 Enter the host name of one of the systems in the cluster.

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

The program checks that the local node running the script can communicate with remote nodes and checks whether SF Sybase CE 7.1 is configured properly.

- 4 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
```

- 5 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.

- 6 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. Veritas 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.
- 7 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 `vxfcntlshdw` utility and then return to this configuration program. See [“Checking shared disks for I/O fencing”](#) on page 41.
- 8 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 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/vxfermode` file to a date and time suffixed file `/etc/vxfermode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Updates the I/O fencing configuration file `/etc/vxfermode`.
 - 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.

- 16** Enable auto refresh of coordination points.

Do you want to enable auto refresh of coordination points
if registration keys are missing
on any of them? [y,n,q,b,?] **(n)**

Performing an automated SF Sybase CE configuration

This chapter includes the following topics:

- [Configuring SF Sybase CE using response files](#)
- [Response file variables to configure SF Sybase CE](#)
- [Sample response files for configuring SF Sybase CE](#)

Configuring SF Sybase CE using response files

To configure SF Sybase CE using response files

- 1 Make sure the Veritas InfoScale Availability or Enterprise packages are installed on the systems where you want to configure SF Sybase CE.
- 2 Copy the response file to one of the cluster systems where you want to configure SF Sybase CE.
- 3 Edit the values of the response file variables as necessary.

To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.

See [“Response file variables to configure SF Sybase CE”](#) on page 49.

- Start the configuration from the system to which you copied the response file.
For example:

```
# /opt/VRTS/install/installer -responsefile
/tmp/response_file
```

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

- 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.

Response file variables to configure SF Sybase CE

[Table 4-1](#) lists the response file variables that you can define to configure SF Sybase CE.

Table 4-1 Response file variables specific to configuring SF Sybase 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 packages 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)

Table 4-1 Response file variables specific to configuring SF Sybase CE
(continued)

Variable	List or Scalar	Description
CFG{activecomponent}	List	Defines the component to be configured. The value is SFSYBASECE71 for SFSYBASECE (Required)
CFG{keys}{keyless} CFG{keys}{license}	List	CFG{keys}{keyless} gives a list of keyless keys to be registered on the system. CFG{keys}{license} gives a list of user defined keys to be registered on the system. (Optional)
CFG{systems}	List	List of systems on which the product is to be configured. (Required)
CFG{prod}	Scalar	Defines the product for operations. The value is ENTERPRISE71 for Veritas InfoScale Enterprise. (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 rsh must be used instead of ssh as the communication method between systems. (Optional)

Table 4-1 Response file variables specific to configuring SF Sybase 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 Veritas website. The value 0 indicates that the installation logs are not uploaded to the Veritas website. (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 smtpsevr), the SNMP trap notification (snmppport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

[Table 4-2](#) lists the response file variables that specify the required information to configure a basic SF Sybase CE cluster.

Table 4-2 Response file variables specific to configuring a basic SF Sybase CE cluster

Variable	List or Scalar	Description
CFG{donotreconfigurevcs}	Scalar	Defines if you need to re-configure VCS. (Optional)

Table 4-2 Response file variables specific to configuring a basic SF Sybase CE cluster *(continued)*

Variable	List or Scalar	Description
CFG{donotreconfigurefencing}	Scalar	Defines if you need to re-configure fencing. (Optional)
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)

[Table 4-3](#) lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 4-3 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. At least 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 4-3 Response file variables specific to configuring private LLT over Ethernet (*continued*)

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

[Table 4-4](#) lists the response file variables that specify the required information to configure LLT over UDP.

Table 4-4 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)
CFG{vcs_udplink<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 4-4 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_address} {<sys1>}	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} {<sys1>}	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} {<sys1>}	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} {<sys1>}	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 4-4 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_netmask} {<sys1>}	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 4-5 lists the response file variables that specify the required information to configure virtual IP for SF Sybase CE cluster.

Table 4-5 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 4-6 lists the response file variables that specify the required information to configure the SF Sybase CE cluster in secure mode.

Table 4-6 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. ■ 1—Configure the first node ■ 2—Configure the other node
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. List or scalar: list Optional or required: optional
CFG{rootsecusrgrps}	Scalar	Defines the read access to the cluster only for root and other users or user groups which are granted explicit privileges in VCS objects. (Optional)
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{defaultaccess}	Scalar	Defines if the user chooses to grant read access to everyone. Optional or required: optional
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

[Table 4-7](#) lists the response file variables that specify the required information to configure VCS users.

Table 4-7 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)
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 4-8](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 4-8 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.example.com) of the SMTP server to be used for web notification. (Optional)
CFG{vcs_smtprecip}	List	List of full email addresses (example: user@example.com) of SMTP recipients. (Optional)

Table 4-8 Response file variables specific to configuring VCS notifications using SMTP (*continued*)

Variable	List or Scalar	Description
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, 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 4-9](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 4-9 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 4-10](#) lists the response file variables that specify the required information to configure SF Sybase CE global clusters.

Table 4-10 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)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response files for 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{keys}{keyless}=[qw(ENTERPRISE)];
$CFG{prod}="ENTERPRISE71";
$CFG{activecomponent}=[ qw(SFSYBASECE71) ];
$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}="bge1";
$CFG{vcs_lltlink1}{sys2}="bge1";
$CFG{vcs_lltlink2}{sys1}="bge2";
$CFG{vcs_lltlink2}{sys2}="bge2";
$CFG{vcs_userenpw}=[ qw(JqrJqlQnrMrrPzrLqo) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];

1;

```

The following sample response file only configures CFS on two nodes, sys1 and sys2.

```

our %CFG;

$CFG{activecomponent}=[ qw(SFSYBASECE71) ];
$CFG{config_cfs}=1;
$CFG{defaultaccess}=1;
$CFG{fencingenabled}=0;
$CFG{lltoverudp}=0;
$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{noipc}=1;
$CFG{opt}{vr}=1;
$CFG{prod}="ENTERPRISE71";
$CFG{reconfigurevcs}=1;
$CFG{secusrgrps}=[ qw("root") ];
$CFG{sfsybasece}{menu}=1;
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=10056;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_eat_security}=1;
$CFG{vcs_lltlink1}{sys1}="bge1";
$CFG{vcs_lltlink1}{sys2}="bge1";
$CFG{vcs_lltlink2}{sys1}="bge2";
$CFG{vcs_lltlink2}{sys2}="bge2";
$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 14.
- 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 65.
- 4 Edit the values of the response file variables as necessary.
See “[Response file variables to configure disk-based I/O fencing](#)” on page 63.

- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installer  
-responsefile /tmp/response_file
```

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

Response file variables to configure disk-based I/O fencing

Table 5-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for SF Sybase CE.

Table 5-1 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. ■ 2—Sybase Mode fencing ■ 3—Replace/Add/Remove coordination points ■ 4—Refresh keys/registrations on the existing coordination points (Required)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.

Table 5-1 Response file variables specific to configuring disk-based I/O fencing (*continued*)

Variable	List or Scalar	Description
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.
CFG{fencing_cpagent_monitor_freq}	Scalar	Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution. 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 <code>LevelTwoMonitorFreq</code> attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If <code>LevelTwoMonitorFreq</code> 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.
CFG {fencing_config_cpagent}	Scalar	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.

Table 5-1 Response file variables specific to configuring disk-based I/O fencing (*continued*)

Variable	List or Scalar	Description
CFG {fencing_cpagentgrp}	Scalar	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the fencing_config_cpagent field is given a value of '0'.
CFG(fencing_auto_refresh_reg)	Scalar	Enable the auto refresh of coordination points variable in case registration keys are missing on any of CP servers.

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 63.

```
# 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{activecomponent}=[ qw(SFSYBASECE71) ];
$CFG{fencing_option}=2;
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{opt}{noipc}=1;
$CFG{prod}="ENTERPRISE71";
$CFG{sfsybasece}{menu}=2;
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=10056;
$CFG{vcs_clustername}="clus1";
```

1;

Configuring a Sybase 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:

```
# installer -responsefile /opt/VRTS/install/logs/\
installer-installernumber/installer-installer\
number.response
```

The following sample response file configures SF Sybase CE under VCS control.

```
our %CFG;

$CFG{opt}{configure}=1;
$CFG{prod}="ENTERPRISE70";
$CFG{activecomponent}=[ qw(SFSYBASECE70) ];
$CFG{sfsybasece}{ase_home}="/sybase_home";
$CFG{sfsybasece}{ase_owner}="sybase";
$CFG{sfsybasece}{ase_quorum}="/qrmmt/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}{qrm dg304}{vol1}{mount}="/qrmmt";
$CFG{sfsybasece}{storage_resource}{qrm dg304}{vol1}{usage}="quorum device";
$CFG{sfsybasece}{storage_resource}{sybdatadg304}{vol1}{mount}="/datamnt";
$CFG{sfsybasece}{storage_resource}{sybdatadg304}{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 6-1 lists the response file variables that you can define to configure SF Sybase CE in VCS.

Table 6-1 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.

Table 6-1 Response file variables specific to configuring SF Sybase CE in VCS (*continued*)

Variable	List or Scalar	Description
CFG{sfsybasece}{ase_owner}	Scalar	Defines the SF Sybase CE owner name.
CFG{sfsybasece}{ase_quorum}	Scalar	Defines the SF Sybase CE quorum device.
CFG{sfsybasece}{ase_sa}	Scalar	Defines the SF Sybase CE administrator name.
CFG{sfsybasece}{ase_server} {sol58236}{SERVER}	Scalar	Defines the SF Sybase CE instance name on sol58236.
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} {<diskgroupname>} {<volumename>}{usage}	Scalar	Lists the SF Sybase CE database devices that reside on the diskgroupname 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}{quorum_vol} {usage}	Scalar	Lists the SF Sybase CE quorum devices that reside on the quorum_dontuse diskgroup and the quorum_vol volume.

Table 6-1 Response file variables specific to configuring SF Sybase CE in VCS (*continued*)

Variable	List or Scalar	Description
CFG{sfsybasece}{storage_resource} {sybase1_dontuse}{sybasevol} {mount}="/opt/sybase"	Scalar	Specifies the SF Sybase CE installation location under /opt/sybase.
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. ■ 1—Location on CFS ■ 2—Location on a local VxFS file system.

Post-installation and configuration tasks

- [Chapter 7. Verifying the installation](#)
- [Chapter 8. Performing additional post-installation and configuration tasks](#)

Verifying the installation

This chapter includes the following topics:

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

Upgrading the disk group version

After you upgrade from previous versions to 7.1, you have to upgrade the disk group version manually.

To upgrade disk group version, you have to first upgrade the cluster protocol version using the `vxddctl upgrade` command.

```
# vxddctl list
Volboot file
version: 3/1
seqno: 0.1
cluster protocol version: 140
hostid: sys1
hostguid: {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
#
# vxddctl upgrade
#

# vxddctl list

Volboot file
version: 3/1
```

```
seqno:    0.2
cluster protocol version: 160
hostid:   sys1
hostguid: {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
```

Verify if the cluster protocol version shows 140 and disk group version is upgraded to 200.

```
# vxctl list |grep version

version:    140
#
# vxdg upgrade dg_name
#
# vxdg list dg_name |grep version

version:    220
```

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 207.

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 74.
- 4 Verify GAB operation.
See [“Verifying GAB”](#) on page 76.
- 5 Verify the cluster operation.
See [“Verifying the cluster”](#) on page 78.

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
```

The output on sys1 resembles:

LLT node information:

Node	State	Links
*0 sys1	OPEN	2
1 sys2	OPEN	2

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

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

LLT node information:

Node	State	Links
* 0 sys1	OPEN	2
1 sys2	OPEN	2
2 sys5	OPEN	1

- 3** Log in as superuser on the node sys2.
- 4** Run the `lltstat` command on the node sys2 to view the status of LLT.

```
lltstat -n
```

The output on sys2 resembles:

LLT node information:

Node	State	Links
0 sys1	OPEN	2
*1 sys2	OPEN	2

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

For example, run the following command on the node sys1 in a two-node cluster:

```
lltstat -nvv active
```

The output on sys1 resembles the following:

- For Solaris SPARC:

Node	State	Link	Status	Address
*0 sys1	OPEN			
		bge1	UP	08:00:20:93:0E:34
		bge2	UP	08:00:20:93:0E:38
1 sys2	OPEN			

```
bge1 UP      08:00:20:8F:D1:F2
bge2 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 7-1](#) lists the different ports that the software configures for different functions.

Table 7-1 GAB port description

Port	Function
a	GAB
b	I/O fencing
f	Cluster File System (CFS)
h	Cluster Server (VCS: High Availability Daemon)
m	Cluster Volume Manager (CVM) CVM uses port m for SmartIO VxVM cache coherency using Group Lock Manager (GLM).
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 <code>>= 100</code> . 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 *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  fd6a01 membership 01
Port b gen  fd6a03 membership 01
Port f gen  fd6a12 membership 01
Port h gen  fd6a06 membership 01
Port m gen  fd6a0b membership 01
Port u gen  fd6a10 membership 01
Port v gen  fd6a09 membership 01
Port w gen  fd6a0d membership 01
Port y gen  fd6a08 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 *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 *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 in the following procedure is for SPARC and it 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	CPU 7.84 Mem 2600.00 Swap 4095.00
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	CPU 8.00 Mem 3448.00 Swap 4095.00
sys1	ConfigBlockCount	130
sys1	ConfigChecksum	46688
sys1	ConfigDiskState	CURRENT
sys1	ConfigFile	/etc/VRTSvcs/conf/config
sys1	ConfigInfoCnt	0
sys1	ConfigModDate	Wed 23 Mar 2016 12:36:17 PM IST
sys1	ConnectorState	Down
sys1	CurrentLimits	
sys1	DiskHbStatus	
sys1	DynamicLoad	
sys1	EngineRestarted	0
sys1	EngineVersion	7.1.00.0
sys1	FencingWeight	0
sys1	Frozen	0
sys1	GUIIPAddr	
sys1	HostUtilization	CPU 2 Mem 24 Swap 0
sys1	LLTNodeId	0
sys1	LicenseType	PERMANENT_SITE

```

sys1      Limits

sys1      LinkHbStatus      bge1 UP bge2 UP

sys1      LoadTimeCounter   0

sys1      LoadTimeThreshold 600

sys1      LoadWarningLevel  80

sys1      MemThresholdLevel Critical 90 Warning 80
                        Note 70 Info 60

sys1      MeterRecord       AvailableGC 2 ForecastGC 0

sys1      NoAutoDisable     0

sys1      NodeId            0

sys1      OnGrpCnt          1

sys1      PhysicalServer

sys1      ShutdownTimeout   600

sys1      SourceFile        ./main.cf

sys1      SwapThresholdLevel Critical 90 Warning 80 Note 70
                        Info 60

sys1      Site

sys1      SysInfo            Solaris:sys1,Generic_
                        118558-11,5.9,SUN4u

sys1      SysName            sys1

sys1      HostAvailableForecast

sys1      ReservedCapacity

sys1      ServerAvailableCapacity

sys1      ServerAvailableForecast

sys1      ServerCapacity

sys1      ServerReservedCapacity

```

Performing additional post-installation and configuration tasks

This chapter includes the following topics:

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

Installing language packages

To install SF Sybase CE in a language other than English, install the required language packages after installing the English packages.

To install the language packages on the server

- 1 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.
- 2 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0
# ./install_lp
```

About enabling LDAP authentication for clusters that run in secure mode

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

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

Figure 8-1 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 Veritas 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: 6.1.12.8
```

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 `catatldapconf` command to view the entries in the attributes file.

- 2 Run the LDAP configuration tool 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 LDAP_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 14149

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: 6.1.12.8

# /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=domain,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:LDAP_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 Add the non-root user to the VCS configuration.

```
# haconf -makerw
# hauser -add user1
# haconf -dump -makero
```

9 Log in as non-root user and run VCS 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 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 *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.veritas.com/land>

For more information:

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

Upgrade of SF Sybase CE

- [Chapter 9. Planning to upgrade SF Sybase CE](#)
- [Chapter 10. Performing a full upgrade of SF Sybase CE using the product installer](#)
- [Chapter 11. Performing an automated full upgrade of SF Sybase CE using response files](#)
- [Chapter 12. Performing a phased upgrade of SF Sybase CE](#)
- [Chapter 13. Performing a rolling upgrade of SF Sybase CE](#)
- [Chapter 14. Upgrading SF Sybase CE using Live Upgrade](#)
- [Chapter 15. Performing post-upgrade tasks](#)

Planning to upgrade SF Sybase CE

This chapter includes the following topics:

- [About the upgrade](#)
- [Supported upgrade paths](#)
- [Using Install Bundles to simultaneously install or upgrade full releases \(base, maintenance, rolling patch\), and individual patches](#)

About the upgrade

This release supports upgrades from 6.0 and later versions. If your existing installation is from a pre-6.0 version, you must first upgrade to version 6.0, then follow the procedures mentioned in this document to upgrade the product.

The installer supports the following types of upgrade:

- Full upgrade
- Automated upgrade using response files
- Phased Upgrade
- Rolling Upgrade

[Table 9-1](#) describes the product mapping after an upgrade.

Table 9-1 Veritas InfoScale product mapping after upgrade

Product (6.2.x and earlier)	Product (7.0 and later)	Component (7.0 and later)
SF Sybase CE	Veritas InfoScale Enterprise	SF Sybase CE

Note: From 7.0 onwards, the existing Veritas InfoScale product upgrades to the higher version of the same product. For example, Veritas InfoScale Enterprise 7.0 gets upgraded to Veritas InfoScale Enterprise 7.1.

During the upgrade, the installation program performs the following tasks:

1. Stops the product before starting the upgrade
2. Upgrades the installed packages and installs additional packages
- If your current installation uses a permanent license key, you will be prompted to update the license to 7.1. If you choose not to update, you can continue to use the old license, limiting the capability of your product to the corresponding component. For example, if you choose not to update the permanent license of your existing SFCFSHA installation, the installer after upgrade will enable SFCFSHA component. The capabilities of other components in the product Veritas InfoScale Enterprise will not be available to you. If your installation uses a keyless license, the installer registers the new keys for the new product with full product capabilities.
3. Restores the existing configuration.
- For example, if your setup contains an SF Sybase CE installation, the installer upgrades and restores the configuration to SF Sybase CE. If your setup included multiple components, the installer upgrades and restores the configuration of the components.
4. Starts the configured components.

Note: If the root disk is encapsulated, you need not unencapsulate the root disk. Reboot the system after the upgrade.

Supported upgrade paths

Table 9-2 lists the supported upgrade paths.

Table 9-2 Supported upgrade paths

From product version	From OS version	To OS version	To product version	To component
6.0	Solaris 10	Solaris 10 Update 9 or later	Veritas InfoScale Enterprise 7.1	SF Sybase CE

Table 9-2 Supported upgrade paths *(continued)*

From product version	From OS version	To OS version	To product version	To component
6.0.1, 6.0.3, 6.0.5	Solaris 10	Solaris 10 Update 9 or later	Veritas InfoScale Enterprise 7.1	SF Sybase CE
6.1, 6.1.1 6.2, 6.2.1	Solaris 10 (SPARC)	Solaris 10 Update 9 or later	Veritas InfoScale Enterprise 7.1	SF Sybase CE
7.0, 7.0.1	Solaris 10 (SPARC)	Solaris 10 Update 9 or later	Veritas InfoScale Enterprise 7.1	SF Sybase CE

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Beginning with version 6.1, you can easily install or upgrade your systems directly to a base, maintenance, patch level or a combination of multiple patches and packages together in one step using Install Bundles. With Install Bundles, the installer has the ability to merge so that customers can install or upgrade directly to maintenance or patch levels in one execution. The various scripts, packages, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or patch levels.

Releases are divided into the following categories:

Table 9-3 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	packages	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	packages	All products	Maintenance Release (MR), Rolling Patch (RP)	Veritas Services and Operations Readiness Tools (SORT)

Table 9-3 Release Levels (*continued*)

Level	Content	Form factor	Applies to	Release types	Download location
Patch	Fixes	packages	Single product	P-Patch, Private Patch, Public patch	SORT, Support site

When you install or upgrade using Install Bundles:

- Veritas InfoScale products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more patches applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and patch releases can be automatically downloaded from SORT.
- Patches can be installed using automated installers from the 6.0.1 version or later.
- Patches can now be detected to prevent upgrade conflict. Patch releases are not offered as a combined release. They are only available from Veritas Technical Support on a need basis.

You can use the `-base_path` and `-patch_path` options to import installation code from multiple releases. You can find packages and patches from different media paths, and merge package and patch definitions for multiple releases. You can use these options to use new task and phase functionality to correctly perform required operations for each release component. You can install the packages and patches in defined phases using these options, which helps you when you want to perform a single start or stop process and perform pre and post operations for all level in a single operation.

Four possible methods of integration exist. All commands must be executed from the highest base or maintenance level install script.

In the example below:

- 7.1 is the base version
- 7.1.1 is the maintenance version
- 7.1.1.100 is the patch version for 7.1.1
- 7.1.0.100 is the patch version for 7.1

1. Base + maintenance:

This integration method can be used when you install or upgrade from a lower version to 7.1.1.

Enter the following command:

```
# installmr -base_path <path_to_base>
```

2. Base + patch:

This integration method can be used when you install or upgrade from a lower version to 7.1.0.100.

Enter the following command:

```
# installer -patch_path <path_to_patch>
```

3. Maintenance + patch:

This integration method can be used when you upgrade from version 7.1 to 7.1.1.100.

Enter the following command:

```
# installmr -patch_path <path_to_patch>
```

4. Base + maintenance + patch:

This integration method can be used when you install or upgrade from a lower version to 7.1.1.100.

Enter the following command:

```
# installmr -base_path <path_to_base>  
-patch_path <path_to_patch>
```

Note: From the 6.1 or later release, you can add a maximum of five patches using `-patch_path <path_to_patch> -patch2_path <path_to_patch> ... -patch5_path <path_to_patch>`

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 7.1](#)
- [Upgrading to SF Sybase CE 7.1](#)

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 7.1.

You can perform the upgrade using Veritas script-based installation program, `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.

Preparing to perform a full upgrade to SF Sybase CE 7.1

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 each node: `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 \
/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 From any node in the cluster, 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).

```
# hagr -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
# hagr -modify sybase_group AutoStart 0
# haconf -dump -makero
```

- 8** Stop VCS on all nodes:

```
# hastop -all
```

- 9** Unmount the VxFS file system on each node, which is not under VCS control.

```
# mount -v |grep vxfs

# fuser -c /mount_point

# umount /mount_point
```

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

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

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

Stop all the ports using the installer as follows:

For 6.0 version:

```
# /opt/VRTS/install/installsfbasece -stop sys1 sys2
```

For 6.0.x, 6.1.x, and 6.2.x versions:

```
# /opt/VRTS/install/installsfbasece<version> -stop sys1 sys2
```

Where <version> is the specific release version.

For 7.0.x and later:

```
# /opt/VRTS/install/installer -stop sys1 sys2
```

Upgrading to SF Sybase CE 7.1

Perform the steps in the following procedure to upgrade to SF Sybase CE 7.1.

To upgrade to SF Sybase CE 7.1

- 1 If you want to upgrade the operating system, perform the following steps:
 - On each node, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

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

- If you are upgrading to Solaris 10 Update 10, apply the following Oracle (Solaris) patches. For instructions, see Oracle documentation.
 - For SPARC: 144524-02
- Upgrade the operating system on all nodes in the cluster. For instructions, see the operating system documentation.
- 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 7.1 using the script-based installer.

See [“Upgrading SF Sybase CE using the Veritas installation program”](#) on page 100.

You can also perform a silent upgrade:

See [“Upgrading SF Sybase CE using a response file”](#) on page 103.

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

```
# hagrps -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
# hagrps -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 144.
- See “[Reverting to the backup boot disk group after an unsuccessful upgrade](#)” on page 145.
- See “[Setting or changing the product license level](#)” on page 145.
- See “[Upgrading disk layout versions](#)” on page 146.
- See “[Upgrading the disk group version](#)” on page 72.

8 Upgrade Sybase ASE CE, if required.

For instructions, see the section *Upgrading Sybase ASE CE* in this document.

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

```
# vxdctl -c mode
```

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

```
# vxdctl upgrade
```

Upgrading SF Sybase CE using the Veritas installation program

Use the `installer` Veritas installation programs to upgrade SF Sybase CE:

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 Veritas InfoScale Enterprise 7.1 packages on the nodes.
- Starts SF Sybase CE 7.1 on all the nodes.
- Displays the location of the log files, summary file, and response file.

Note: The SF Sybase CE processes are started automatically after the upgrade completes successfully.

To upgrade to SF Sybase CE 7.1 using the installer program

- 1 Start the installation program:

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

- 2 From the menu, choose **Upgrade a product > Full upgrade**.

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.

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

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

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

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

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

You are prompted to start the split operation.

- 5 Enter `y` to continue with the split operation.

The split operation can take some time to complete.

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

- 6** 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 Veritas InfoScale 7.1 and starts SF Sybase CE 7.1 on all the nodes.

- 7** Install the language packages and patches if you would like to run SF Sybase CE in a language other than English.

See [“Installing language packages”](#) on page 82.

- 8** Complete the remaining tasks to finish the upgrade:

See [“Upgrading to SF Sybase CE 7.1”](#) on page 99.

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 SF Sybase 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 7.1 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:

For response file variable definitions:

See [“Response file variables to upgrade SF Sybase CE”](#) on page 104.

For a sample response file:

See [“Sample response file for upgrading SF Sybase CE”](#) on page 107.

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

```
# ./installer -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 SF Sybase CE

[Table 11-1](#) lists the response file variables that you can define to configure SF Sybase CE.

Table 11-1 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

Table 11-1 Response file variables for upgrading SF Sybase CE (*continued*)

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
CFG{upgrade}	Upgrades all packages installed. List or scalar: list Optional or required: required
CFG{keys}{keyless} CFG{keys}{license}	CFG{keys}{keyless} gives a list of keyless keys to be registered on the system. CFG{keys}{license} gives a list of user defined keys to be registered on the system. 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 packages 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

Table 11-1 Response file variables for upgrading SF Sybase CE (*continued*)

Variable	Description
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
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
CFG{opt}{disable_dmp_native_support}	If it is set to 1, Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools is disabled after upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system. List or scalar: scalar Optional or required: optional
CFG{opt}{patch_path}	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed . List or scalar: scalar Optional or required: optional
CFG{opt}{patch2_path}	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional

Table 11-1
Response file variables for upgrading SF Sybase CE *(continued)*

Variable	Description
CFG{opt}{patch3_path}	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{patch4_path}	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{patch5_path}	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{rootsecusrgrps}	Defines if the user chooses to grant read access to the cluster only for root and other users/usergroups which are granted explicit privileges on VCS objects. List or scalar: scalar Optional or required: optional
CFG{secusrgrps}	Defines the usergroup names that are granted read access to the cluster. 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{keys}{keyless}=[ qw(ENTERPRISE) ];  
$CFG{opt}{gco}=1;  
$CFG{opt}{upgrade}=1;  
$CFG{opt}{vr}=1;  
$CFG{systems}=[ qw(sys1) ];  
$CFG{vcs_allowcomms}=1;  
1;
```

Performing a phased upgrade of SF Sybase CE

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing a phased upgrade of SF Sybase CE from version 6.0 and later release](#)

About phased upgrade

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

For supported upgrade paths:

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 a phased upgrade of SF Sybase CE from version 6.0 and later release

Table 12-1 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 12-1 Summary of phased upgrade



First half of the cluster	Second half of the cluster
SF Sybase CE cluster before the upgrade:	
First half of the cluster  Second half of the cluster 	
STEP 1: Perform the following pre-upgrade steps: ■ Switch failover applications. ■ Stop all parallel applications. See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 111. STEP 2: Upgrade SF Sybase CE. See “Step 2: Upgrading the first half of the cluster” on page 114.	The second half of the cluster is up and running. 
The first half of the cluster is not running. 	STEP 3: Perform the following pre-upgrade steps: ■ Stop all parallel and failover applications. ■ Stop SF Sybase CE. See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 116. The downtime starts now.

Table 12-1 Summary of phased upgrade (continued)

First half of the cluster	Second half of the cluster
STEP 4: Perform the following post-upgrade steps: ■ Start SF Sybase CE. ■ Start all applications. See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 118. The downtime ends here.	The second half of the cluster is not running. 
The first half of the cluster is up and running. 	STEP 5: Upgrade SF Sybase CE. See “Step 5: Upgrading the second half of the cluster” on page 118. STEP 6: Perform the following post-upgrade steps: ■ Start SF Sybase CE. ■ Start all applications. See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 119.
The phased upgrade is complete and both the first and the second half of the cluster are running. 	

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
```

The installer verifies that recent backups of configuration files in the VxVM private region are saved in `/etc/vx/cbr/bk`.

If not, the following warning message is displayed: `Warning: Backup /etc/vx/cbr/bk directory.`

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

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

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

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

- 5 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
```

- 6 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 CFS mount point:

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

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

- Unmount the CFS file system:

```
# umount /mount_point
```

- 7 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
```

- 8 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 mount point:

```
# mount | grep vxfs
```

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

- Unmount the VxFS file system:

```
# umount /mount_point
```

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

```
# vxvol -g diskgroup stopall
# vxprint -Aht -e v_open
```

- 10 Stop all the ports using the installer as follows:

For 6.0 version:

```
# /opt/VRTS/install/installsfsybasece -stop sys1 sys2
```

For 6.0.x, 6.1.x, 6.2.x :

```
# /opt/VRTS/install/installsfsybasece<version> -stop sys1 sys2
```

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

For 7.0 and later:

```
# /opt/VRTS/install/installer -stop sys1 sys2
```

Use the `modinfo` command to check for loaded Veritas kernel modules and unload the modules, if any.

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
```

If you are upgrading to Solaris 10 Update 10:

Apply the following Oracle patch for SPARC: 144524-02

See the Oracle documentation for instructions.

- 2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

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

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

You may see some errors in the system log file when the nodes restart. This is because LLT is disabled. Ignore these messages.

```
svc.startd[7]: [ID 652011 daemon.warning] svc:/system/llt:default:
Method "/lib/svc/method/llt start" failed with exit status 2.
gab: [ID 438192 kern.notice] GAB WARNING V-15-1-20115
Port d registration failed, GAB not configured
```

- 4 After you finish upgrading the operating system, change the LLT configuration to its original configuration.

```
# 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 using the product installer. 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 image folder
```

```
# cd /dvd_mount/
```

Using the product installer:

```
# ./installer -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.

- 7 Change the `/etc/default/llt` file to prevent LLT from starting automatically after reboot by setting the `LLT_START` attribute to 0:

```
LLT_START=0
```

- 8 Restart the nodes:

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

You may see some errors in the system log file when the nodes restart. This is because LLT is disabled. Ignore these messages.

```
svc.startd[7]: [ID 652011 daemon.warning] svc:/system/llt:default:
Method "/lib/svc/method/llt start" failed with exit status 2.
gab: [ID 438192 kern.notice] GAB WARNING V-15-1-20115
Port d registration failed, GAB not configured
```

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

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

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

- 1 Stop all applications that are not configured under VCS 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.

```
# hagrpp -offline sybase_group -sys sys3
```

```
# hagrpp -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:

```
# hstop -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.

```
# vxvol -g diskgroup stopall
```

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

- 7 Stop all the ports using installer as follows:

For 6.0 version:

```
# /opt/VRTS/install/installsfbasece -stop sys1 sys2
```

For 6.0.1x and later versions:

```
# /opt/VRTS/install/installsfbasece<version> -stop sys1 sys2
```

Where<version> is the specific release version.

For 7.0 and later.

```
# /opt/VRTS/install/installer -stop sys1 sys2
```

Use the `modinfo` command to check loaded Veritas kernel modules and unload the modules, if any.

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

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

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

- 1 Change `/etc/default/llt` to start LLT on the nodes by setting the `LLT_START` attribute to 1:

Run the following command to bring LLT online, if it is in maintenance mode:

```
# svcadm clear llt

LLT_START=1
```

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

```
# gabconfig -x
```

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

```
# cd /opt/VRTS/install

# ./installer -start sys1 sys2
```

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

- 5 Bring the sybasece group online.

```
# hagrps -online sybasece -sys sys1

# hagrps -online sybasece -sys sys2
```

Note: The downtime ends here.

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

Step 5: Upgrading the second half of the cluster

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

To upgrade the second half of the cluster

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

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

If you are upgrading to Solaris 10 Update 10:

Apply the following Oracle patch on SPARC systems: 144524-02

- 2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

- 3

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

If you upgraded the operating system, restart the nodes:

- 4 After you finish upgrading the operating system, change the LLT configuration to its original configuration.

```
# 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 using the product installer. Navigate to the product directory on the installation media.

```
# cd /dvd_mount/
```

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

Using the product installer:

```
# ./installer -upgrade sys3 sys4
```

- 7 Restart the nodes, if prompted by the installer:

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

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

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

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

- 1 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
# ./installer -start sys3 sys4
```

- 3 Upgrade VxVM disk group version.

See [“Upgrading CVM protocol version and VxVM disk group version”](#) on page 147.

- 4 Upgrade disk layout version.

See [“Upgrading disk layout versions”](#) on page 146.

- 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 145.

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 168.

Performing a rolling upgrade of SF Sybase CE

This chapter includes the following topics:

- [About rolling upgrade](#)
- [Supported rolling upgrade paths](#)
- [Preparing to perform a rolling upgrade to SF Sybase CE 7.1](#)

About rolling upgrade

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 packages in phase 1 and VCS agent related packages 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 failover. 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 Cluster Server (VCS) engine HAD, but does not include application downtime.

Figure 13-1 illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 13-1 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.
- You can perform a rolling upgrade from 6.0 and later versions.
- The rolling upgrade procedures support only minor operating system upgrades.
- The rolling upgrade procedure requires the product to be started before and after upgrade. If the current release does not support your current operating system version and the installed old release version does not support the operating system version that the current release supports, then rolling upgrade is not supported.

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 13-1](#) shows the versions of SF Sybase CE for which you can perform a rolling upgrade to SF Sybase CE 7.1.

Table 13-1 Supported rolling upgrade paths

Platform	SF Sybase CE version
Solaris 10 SPARC	6.0, 6.0RP1, 6.0.5, 6.1, 6.2, 6.2.1, 7.0, and 7.0.1

Preparing to perform a rolling upgrade to SF Sybase CE 7.1

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.

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: Back up the `/etc/vx/cbr/bk` directory.

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

```
# hagrps -switch grp_name -to sys_name
```

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

```
# mount -v |grep vxfs | grep cluster
```

```
# fuser -c /mount_point
```

```
# umount /mount_point
```

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

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

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

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

Upgrading SF Sybase CE using Live Upgrade

This chapter includes the following topics:

- [About Live Upgrade](#)
- [Supported upgrade paths for Live Upgrade and Boot Environment upgrade](#)
- [Before you upgrade SF Sybase CE using Solaris Live Upgrade](#)
- [Performing Live Upgrade on Solaris 10 systems](#)
- [Administering boot environments in Solaris 10 Live Upgrade](#)
- [Upgrading the operating system and SF Sybase CE using Live Upgrade](#)
- [Upgrading SF Sybase CE only using Live Upgrade](#)
- [Upgrading Solaris only using Live Upgrade](#)
- [Completing the Live Upgrade](#)

About Live Upgrade


Solaris Live Upgrade provides a method of upgrading a system while the system continues to operate. This is done by creating an alternate boot environment (ABE) from the current boot environment and then upgrading the ABE. Once the ABE is upgraded, you can activate the ABE and then reboot the system.

On Solaris 10, you can use Live Upgrade technology to reduce downtime associated with the OS upgrade and SF Sybase CE product upgrade by creating a boot environment on an alternate boot disk.

- See [“Performing Live Upgrade on Solaris 10 systems”](#) on page 127.

Figure 14-1 illustrates an example of an upgrade from 6.0 to 7.1, and the operating system from Solaris 10 Update9 to Solaris 10 Update11 using Live Upgrade.

Figure 14-1 Live Upgrade process



Some service groups (failover and parallel) may be online in this cluster and the Live Upgrade process does not affect them. Downtime is experienced only when the server is restarted to boot into the alternate boot environment.

Supported upgrade paths for Live Upgrade and Boot Environment upgrade

The systems where you plan to use Live Upgrade must run Solaris 10. SF Sybase CE version must be at least 6.0.

Before you upgrade SF Sybase CE using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

- 1 Make sure that the SF Sybase CE installation media and the operating system installation images are available and on hand.
- 2 On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk

If the primary boot disk is mirrored, you need to break off the mirror for the alternate boot disk.

- 3 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you upgrade the Solaris operating system, do the following steps:

- Remove the installed Live Upgrade packages for the current operating system version:

All Solaris versions: `SUNWluu`, `SUNWlur` packages.

Solaris 10 update 7 or later also requires: `SUNWlucfg` package.

- From the new Solaris installation image, install the new versions of the following Live Upgrade packages:

All Solaris versions: `SUNWluu`, `SUNWlur`, and `SUNWlucfg` packages.

Solaris installation media comes with a script for this purpose named `liveupgrade20`. Find the script at

`/cdrom/solaris_release/Tools/Installers/liveupgrade20`. If scripting, you can use:

```
# /cdrom/solaris_release/Tools/Installers/liveupgrade20 \
-nodisplay -noconsole
```

If the specified image has some missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you have to install any missing patches on the alternate boot disk.

Performing Live Upgrade on Solaris 10 systems

Perform the Live Upgrade using the installer.

Table 14-1 Upgrading SF Sybase CE using Solaris 10 Live Upgrade

Step	Description
Step 1	Prepare to upgrade using Solaris Live Upgrade. See “Before you upgrade SF Sybase CE using Solaris Live Upgrade” on page 126.

Table 14-1 Upgrading SF Sybase CE using Solaris 10 Live Upgrade
(continued)

Step	Description
Step 2	Create a new boot environment on the alternate boot disk. See “Creating a new Solaris 10 boot environment on the alternate boot disk” on page 128.
Step 3	Upgrade SF Sybase CE using the installer. See “Upgrading SF Sybase CE using the installer for Solaris 10 Live Upgrade” on page 133. To upgrade only Solaris See the Oracle documentation on Solaris 10 operating system Note: A new boot environment is created on the alternate boot disk by cloning the primary boot environment. If you choose to upgrade the operating system, the Solaris operating system on the alternate boot environment is upgraded.
Step 4	Switch the alternate boot environment to be the new primary. See “Completing the Solaris 10 Live Upgrade” on page 134.
Step 5	Verify Live Upgrade of SF Sybase CE. See “Verifying the Solaris 10 Live Upgrade of SF Sybase CE” on page 135.

Creating a new Solaris 10 boot environment on the alternate boot disk

Veritas provides the `vxlustart` script that runs a series of commands to create the alternate boot environment for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-v` option.

Veritas recommends that you preview the commands with `-v` option to ensure there are no problems before beginning the Live Upgrade process. The `vxlustart` script is located in the `scripts` directory on the distribution media.

Note: This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

```
# cd /cdrom/scripts
```

```
# ./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

Table 14-2

vxlustart option	Usage
-V	Lists the commands to be executed during the upgrade process without executing them and pre-checks the validity of the command. If the operating system is upgraded, the user is prompted to compare the patches that are installed on the image with the patches installed on the primary boot disk. This determines if any critical patches are not present from the new operating system image.
-v	Indicates verbose, print commands before executing them.
-f	Forces the vtoc creation on the disk.
-Y	Indicates a default yes with no questions asked.
-m	Uses the already existing vtoc on the disk.
-D	Prints with debug option on, and is for debugging.
-U	Specifies that only the Storage Foundation products are upgraded. The operating system is cloned from the primary boot disk.
-g	Specifies the DG to which the rootdisk belongs. Optional.
-d	Indicates the name of the alternate boot disk <code>c#t#d#s2</code> on which you intend to upgrade. The default disk is <code>mirrordisk</code> .
-u	Specifies the operating system version for the upgrade on the alternate boot disk. For example, use <code>7.9</code> for Solaris 9 and <code>7.10</code> for Solaris 10. If you want to upgrade only SF products, specify the current OS version.

Table 14-2 (continued)

vxlustart option	Usage
-F	Specifies the root disk's file system, where the default is <code>ufs</code> .
-S	Specifies the path to the Solaris image. It can be a network/directory path. If the installation uses the CD, this option must not be specified. See <i>Solaris Live Upgrade installation guide</i> for more information about the path.
-r	Specifies that if the computer crashes or restarts before the <code>vxlufinish</code> command is run, the alternate disk is remounted using this option.
-k	Specifies the location of file containing auto-registration information. This file is required by <code>luupgrade (1M)</code> for OS upgrade to Solaris 10 9/10 or a later release.
-x	Excludes file from newly created BE. (<code>lucreate -x</code> option)
-X	Excludes file list from newly created BE. (<code>lucreate -f</code> option)
-i	Includes file from newly created BE. (<code>lucreate -y</code> option)
-I	Includes file list from newly created BE. (<code>lucreate -Y</code> option)
-z	Filters file list from newly created BE. (<code>lucreate -z</code> option)
-w	Specifies additional mount points. (<code>lucreate -m</code> option)
-W	Specifies additional mount points in a file (<code>lucreate -M</code> option)

If the `-u` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.

For example, to preview the commands to upgrade only the Veritas InfoScale product:

```
# ./vxlustart -V -u 7.10 -U -d disk_name
```

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0s2) and the alternate or inactive boot environment resides on Disk1 (c0t1d0s2).

At the end of the process:

- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.
- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.

To create a new boot environment on the alternate boot disk

Perform the steps in this procedure on each node in the cluster.

- 1 Navigate to the install media for the Veritas InfoScale products:

```
# cd /dvd_mount/scripts
```

- 2 Before you upgrade, make sure that you exclude the CFS mount points that the database or applications use from being copied to the new boot environment. During Live Upgrade, the `vxlustart` utility fails to recognize the CFS mount points that are configured under VCS. As a result, the data in Sybase and the mount points for Sybase ASE CE binaries, data files, and quorum that are configured as CFS mount points under VCS get copied into the local file system of the alternate boot environment. To prevent these shared mount points from being copied to the new boot environment, you need to identify and exclude these mount points as follows:

```
# for i in `hatype -resources CFSMount`; \
do hares -display $i -attribute MountPoint | awk ' \
NR != 1 { print "-", $4}'; done > /var/tmp/file_list
# cat /var/tmp/file_list
- /sybase_binary
- /masterdb
- /sysprocdb
- /db
```

Where `/var/tmp/file_list` is a temporary file that contains the list of CFS mount points to be excluded from the new boot environment

`/sybase_binary`, `/masterdb`, `/sysprocdb`, and `/db` are CFS mount points that the database or applications use.

The items in the file list are preceded either by a + or - symbol. The + symbol indicates that the mount point is included in the new boot environment. The - symbol indicates that the mount point is excluded from the new boot environment. Apart from CFS mount points, you may choose to include or exclude other files.

- 3 Run one of the following commands to create the alternate boot environment:

For example:

To upgrade both the operating system and the Veritas InfoScale product:

```
# ./vxlustart -v -u 7.10 -s /mnt/sol10u9 -d c0t1d0s2 \  
-z /var/tmp/file_list
```

where */mnt/sol10u9* is the path to the operating system image that contains the *.cdtoc* file.

To only upgrade the Veritas InfoScale product:

```
# ./vxlustart -v -u 7.10 -U -d c0t1d0s2 -z /var/tmp/file_list
```

- 4 Create the mount points manually on the alternate boot environment as follows:

```
# for i in `cat /var/tmp/file_list` ; \  
do mkdir -p /altroot.7.10/$i; done
```

- 5 Update the permissions, user name, and group name of the mount points (created on the ABE) to match that of the existing directories on the primary boot environment.
- 6 Review the output and note the new mount points. If the system is restarted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

If you need to remount, run the command:

```
# vxlustart -r -u targetos_version -d disk_name
```

- 7 After the alternate boot disk is created and mounted on */altroot.7.10*, install any operating system patches or packages on the alternate boot disk that are required for the Veritas InfoScale product installation.

```
# pkgadd -R /altroot.7.10 -d pkg_dir
```

Upgrading SF Sybase CE using the installer for Solaris 10 Live Upgrade

You can use the Veritas InfoScale product installer to upgrade SF Sybase CE as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade SF Sybase CE on all the nodes in the cluster. The program uninstalls the existing version of SF Sybase CE on the alternate boot disk during the process.

At the end of the process, SF Sybase CE 7.1 is installed on the alternate boot disk.

To perform Live Upgrade of SF Sybase CE using the installer

- 1 Insert the product disc with SF Sybase CE 7.1 or access your copy of the software on the network.
- 2 Run the installer script specifying the root path as the alternate boot disk:

```
# ./installer -upgrade -rootpath /altroot.7.10
```
- 3 Enter the names of the nodes that you want to upgrade to SF Sybase CE 7.1. The installer displays the list of packages to be installed or upgraded on the nodes.
- 4 Press **Return** to continue with the installation.
- 5 Verify that the version of the Veritas packages on the alternate boot disk is 7.1.

```
# pkginfo -R /altroot.7.10 -l VRTSpkgname
```

Review the installation logs at `/altroot.7.10/opt/VRTS/install/logs`.

Completing the Solaris 10 Live Upgrade

At the end of the process:

- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

To complete the Live Upgrade

- 1 Complete the Live upgrade process using one of the following commands. You must enter the command on all nodes in the cluster.

If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version
Live Upgrade finish on the Solaris release <7.10>
```

If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
Live Upgrade finish on the Solaris release <7.10>
```

The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

- 2 After the successful completion of `vxlustart`, if the system crashes or restarts before Live Upgrade completes successfully, you can remount the alternate disk using the following command:

```
# ./vxlustart -r -u target_os_version
```

Then, rerun the `vxlufinish` command from step 1

```
# ./vxlufinish -u target_os_version
```

If you have enabled VVR,

- 3 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

You can ignore the following error if it appears: Error: boot environment `<dest.13445>` already mounted on `</altroot.7.10>`.

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

Verifying the Solaris 10 Live Upgrade of SF Sybase CE

To ensure that Live Upgrade has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

To verify that Live Upgrade completed successfully

- 1 Verify that the alternate boot environment is active.

```
# lustatus
```

If the alternate boot environment fails to be active, you can revert to the primary boot environment.

See [“Reverting to the primary boot environment on a Solaris 10 system”](#) on page 136.

- 2 In a cluster environment, make sure that all the GAB ports are up. Note that different ports appear for different products.

```
# gabconfig -a
Port a gen fd6a01 membership 0123
Port b gen fd6a03 membership 0123
Port f gen fd6a12 membership 0123
Port h gen fd6a06 membership 0123
Port m gen fd6a0b membership 0123
Port u gen fd6a10 membership 0123
Port v gen fd6a09 membership 0123
Port w gen fd6a0d membership 0123
Port y gen fd6a08 membership 0123
```

- 3 Perform other verification as required to ensure that the new boot environment is configured correctly.

Administering boot environments in Solaris 10 Live Upgrade

Use the following procedures to perform relevant administrative tasks for boot environments.

Reverting to the primary boot environment on a Solaris 10 system

If the alternate boot environment fails to start, you can revert to the primary boot environment.

On each node, start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

where *disk0* is the primary boot disk.

Switching the boot environment for Solaris 10 SPARC

You do not have to perform the following procedures to switch the boot environment when you use the `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See [“To switch the boot environment if the root disk is not encapsulated”](#) on page 137.
- See [“To switch the boot environment if the root disk is encapsulated”](#) on page 138.

The switching procedures for Solaris SPARC vary, depending on whether VxVM encapsulates the root disk.

To switch the boot environment if the root disk is not encapsulated

- 1 Display the status of Live Upgrade boot environments.

```
# lustatus
```

Boot Environment Name	Is Complete	Active Now	Active On Reboot	Can Delete	Copy Status
source.2657	yes	yes	yes	no	-
dest.2657	yes	no	no	yes	-

In this example, the primary boot environment is currently (source.2657). You want to activate the alternate boot environment (dest.2657).

- 2 Unmount any file systems that are mounted on the alternate boot environment (dest.2657).

```
# lufslist dest.2657
```

```
boot environment name: dest.2657
```

Filesystem	fstype	device	size	Mounted on	Mount Options
/dev/dsk/c0t0d0s1	swap	4298342400	-	-	-
/dev/dsk/c0t0d0s0	ufs	15729328128	/	-	-
/dev/dsk/c0t0d0s5	ufs	8591474688	/var	-	-
/dev/dsk/c0t0d0s3	ufs	5371625472	/vxfs	-	-

```
# luumount dest.2657
```

- 3 Activate the Live Upgrade boot environment.

```
# luactivate dest.2657
```

- 4 Restart the system.

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

The system automatically selects the boot environment entry that was activated.

To switch the boot environment if the root disk is encapsulated

- 1 Display the current boot disk device and device aliases

```
# eeprom
boot-device=vx-rootdg vx-int_disk
use-nvramrc?=true
nvramrc=devalias vx-int_disk /pci@1c,600000/scsi@2/disk@0,0:a
devalias vx-rootdg01 /pci@1c,600000/scsi@2/disk@1,0:a
```

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

```
# eeprom boot-device=vx-rootdg01
```

- 3 Restart the system.

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

Upgrading the operating system and SF Sybase CE using Live Upgrade

Perform the following steps to upgrade both the operating system and SF Sybase CE using Live Upgrade.

To upgrade the operating system and SF Sybase CE using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.

See [“Before you upgrade SF Sybase CE using Solaris Live Upgrade”](#) on page 126.

- 2 Create a new boot environment on the alternate boot disk.

See [“Creating a new Solaris 10 boot environment on the alternate boot disk”](#) on page 128.

- 3 Upgrade SF Sybase CE using the installer or manually.
See [“Upgrading SF Sybase CE using the installer for Solaris 10 Live Upgrade”](#) on page 133.
- 4 Complete the Live Upgrade.
See [“Completing the Live Upgrade”](#) on page 140.
- 5 Verify Live Upgrade of SF Sybase CE.
See [“Verifying the Solaris 10 Live Upgrade of SF Sybase CE”](#) on page 135.

Upgrading SF Sybase CE only using Live Upgrade

Perform the following steps to upgrade only SF Sybase CE using Live Upgrade.

To upgrade only SF Sybase CE using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.
See [“Before you upgrade SF Sybase CE using Solaris Live Upgrade”](#) on page 126.
- 2 Create a new boot environment on the alternate boot disk.
See [“Creating a new Solaris 10 boot environment on the alternate boot disk”](#) on page 128.
- 3 Upgrade SF Sybase CE using the installer or manually.
See [“Upgrading SF Sybase CE using the installer for Solaris 10 Live Upgrade”](#) on page 133.
- 4 Complete the Live Upgrade.
See [“Completing the Live Upgrade”](#) on page 140.
- 5 Verify Live Upgrade of SF Sybase CE.
See [“Verifying the Solaris 10 Live Upgrade of SF Sybase CE”](#) on page 135.

Upgrading Solaris only using Live Upgrade

Perform the following steps to upgrade only Solaris using Live Upgrade.

To upgrade only Solaris using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.
See [“Before you upgrade SF Sybase CE using Solaris Live Upgrade”](#) on page 126.
- 2 Create a new boot environment on the alternate boot disk.
See [“Creating a new Solaris 10 boot environment on the alternate boot disk”](#) on page 128.
- 3 Complete the Live Upgrade.
See [“Completing the Live Upgrade”](#) on page 140.
- 4 Verify Live Upgrade of SF Sybase CE.
See [“Verifying the Solaris 10 Live Upgrade of SF Sybase CE”](#) on page 135.

Completing the Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

When completing the Live Upgrade process, take the following limitations into consideration for Solaris 10 Update 10:

- In a shared disk group environment, extra CFS mount entries are ignored when the `vxlustart` command is run, as they are included in `/etc/vfstab`. The entries must be manually removed before booting from the alternate boot environment.

See the *Symantec Storage Foundation for Sybase ASE CE Release notes* for more details.

To complete the Live Upgrade

- 1 Complete the Live Upgrade process using one of the following commands:

If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version
Live Upgrade finish on the Solaris release <5.10>
```

If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
Live Upgrade finish on the Solaris release <5.10>
```

The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

- 2 If the system crashes or reboots before Live Upgrade completes successfully, you may remount the alternate disk using the following command:

```
# ./vxlustart -r -u target_os_version
```

Then, rerun the `vxlufinish` command:

```
# ./vxlufinish -u target_os_version
```

- 3 If the Sybase database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the `AutoStart` value to 0. This prevents the database service group from starting automatically when VCS starts:

```
group sybasece (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStart = 0
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sys1, sys2 }
)

.
.
```

- 4 Perform the following steps on the primary boot environment:
 - Stop the applications using native application commands.
 - Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu mount-point
```

- Take offline all sybase binary groups (binmnt) and sybase database groups (sybasece) that contain CFSMount and CVMVolDg:

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

- Unmount the VxFS file systems:

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

- Deport CVM disk groups:

```
# vxvg deport diskgroup_name
```

- Make sure that no disk groups are imported:

```
# vxvg list  
NAME STATE ID
```

- 5 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: DO NOT use the reboot, halt, or uadmin commands to reboot the system. Use either the init or the shutdown commands to enable the system to boot using the alternate boot environment.

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

- 6 Start the database group on all nodes:

```
# hagr -online sybasece -any
```

- 7 If the Sybase database is managed by VCS, modify the VCS configuration file (/etc/VRTSvcs/conf/config/main.cf) to set the AutoStart value to 1.

```
group sybasece (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStart = 1
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sys1, sys2 }
)

.
.
```

- 8 Complete the post-upgrade tasks.
See the chapter "Performing post-upgrade tasks" in this document.
- 9 If you are on an unsupported version of Sybase ASE CE, upgrade to ASE CE 15.5.

See ["Upgrading Sybase ASE CE"](#) on page 168.

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](#)

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.

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.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vx dg` command to find the boot disk group where you are currently booted.

```
# vx dg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Setting or changing the product license level

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

After you upgrade, perform one of the following steps:

- Obtain a valid license key and run the `vxlicinst` command to add it to your system.
- Use the `vxkeyless` command to update the license keys to the keyless license model.

For more information and instructions, see the chapter *Licensing SF Sybase CE*.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 7, 8, 9, and 10. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Note: If you plan to use 64-bit quotas, you must upgrade to the latest disk layout Version 10. The use of 64-bit quota on earlier disk layout versions is deprecated in this release.

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.

To upgrade the disk layout versions

- ◆ To get to disk layout Version 10 from Version 6. You must incrementally upgrade the disk layout of this file system. For example:

```
# vxupgrade -n 7 /mnt
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
```

See the `vxupgrade(1M)` manual page.

You must upgrade any existing file systems with disk layout Version 4 or 5 to disk layout Version 7 or later using the `vxfsconvert` command.

See the `vxfsconvert(1M)` manual page.

Note: Veritas recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release. Once a disk layout version has been upgraded, it is not possible to downgrade to the previous version.

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
```

Upgrading CVM protocol version and VxVM disk group version

The default Cluster Volume Manager protocol version is 160.

Run the following command to verify the CVM protocol version:

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

If the protocol version is not 160, 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 220.

Check the existing disk group version:

```
# vxdg list dg_name | grep -i version
```

If the disk group version is not 190, run the following command on the master node to upgrade the version:

```
# vxdg -T 220 upgrade dg_name
```

Installation and upgrade of Sybase ASE CE

- [Chapter 16. 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 Veritas InfoScale Enterprise
- 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=cddisk
```

- 2 Create a diskgroup.

For example:

```
# vxdg init sybbin_dg Disk_1 Disk_2
```

- 3 Create a mirrored volume in the group:

```
# vxassist -g sybbin_dg make sybbin_vol  
12G layout=mirrored nmirrors=2
```

- 4 Create a VxFS file system on which to install the Sybase binaries:

```
# mkfs -F vxfs /dev/vx/rdisk/sybbin_dg/sybbin_vol
```

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 -F vxfs /dev/vx/dsk/sybbin_dg/sybbin_vol /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=cddisk
```

- 2 Create a CVM diskgroup.

For example:

```
# vxdg -s init sybbin_dg Disk_1 Disk_2
```

- 3 Create a mirrored volume in the group:

```
# vxassist -g sybbin_dg make sybbin_vol  
12G layout=mirrored nmirrors=2
```

- 4 Create a VxFS file system on which to install the Sybase binaries:

```
# mkfs -F vxfs -o largefiles /dev/vx/rdsk/sybbin_dg/sybbin_vol
```

For a binary installation on a shared file system, you may run the command on any one node.

- 5 Create a Sybase ASE CEhome directory (\$SYBASE) on all nodes:

```
# mkdir /sybase
```

- 6 Mount the directory:

```
# mount -F vxfs -o cluster /dev/vx/dsk/sybbin_dg/sybbin_vol /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 151.

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=cddisk  
# vxdisksetup -i Disk_4 format=cddisk
```

- 2 As root user, from the CVM master, create a shared VxVM diskgroup for the quorum device.

```
# vxdg -s init quorum_dg Disk_3 Disk_4
```

- 3 As root user, from the CVM master, create a mirrored volume, *quorum_vol*:

```
# vxassist -g quorum_dg make quorum_vol  
1G layout=mirrored \  
nmirrors=2
```

- 4 As root user, from the CVM master, create a filesystem with the volume, *quorum_vol*.

```
# mkfs -F vxfs /dev/vx/rdisk/quorum_dg/quorum_vol
```

- 5 On each system, create a directory, */quorum*:

```
# mkdir /quorum
```

- 6 On each system, mount */quorum*

```
# mount -F vxfs -o cluster /dev/vx/dsk/quorum_dg/quorum_vol  
/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 dbdata_dg Disk_5 Disk_6
```

- 3 As root user, create a mirrored volume, *sybvol*:

```
# vxassist -g dbdata_dg make sybvol 1G layout=mirrored \  
nmirrors=2
```

- 4 As root user, create a filesystem with the volume, *sybvol*.

```
# mkfs -F vxfs /dev/vx/rdisk/dbdata_dg/sybvol
```

- 5 On each system, create a directory, */sybdata*:

```
# mkdir /sybdata
```

- 6 On each system, mount */sybdata*

```
# mount -F vxfs -o cluster /dev/vx/dsk/dbdata_dg/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 152.
- 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 153.

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 agent for Sybase](#)

Language settings for the Sybase agent

For the VCS 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

LC_CTYPE=en_US.UTF-8
export LC_CTYPE

LC_ALL=en_US.UTF-8
export LC_ALL
```

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 158.

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 utility *vcscrypt* 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 secure password for the agent. See the *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.
Save the encrypted password for future reference.

About setting up detail monitoring for the agent for Sybase

The VCS agent for Sybase provides two levels of application monitoring: basic and detail. In basic monitoring, Sybase resource monitors the Sybase dataserver 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. Veritas 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.

Note: All Sybase resources in the cluster gets modified when you enable detail monitoring for agents.

To enable detail monitoring

- 1 Make sure the Sybase server is configured for detail monitoring.
See [“Configuring Sybase for detail monitoring”](#) on page 155.
- 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
```

Note: If the LevelTwoMonitorFreq attribute is overridden at resource level, then use `hares` command and set the LevelTwoMonitorFreq attribute to 0 at resource level.

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 maybe 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 229.

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 16-1](#) displays the service group dependencies for an SF Sybase CE configuration on local disk group with VxFS.

Figure 16-1 Service group dependencies for an SF Sybase CE configuration on local disk group with VxFS



[Figure 16-2](#) displays the service group dependencies for an SF Sybase CE configuration on shared disk group with CFS.

Figure 16-2 *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 Veritas InfoScale Enterprise.
- 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 16-1 lists the required resources for configuring Sybase ASE CE under VCS control.

Table 16-1 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:
■ Disk group ■ Mount point ■ Volume	■ <i>sybbin_dg</i> ■ <i>/sybase</i> ■ <i>sybbin_vol</i>	■ <i>sybbin_dg_voldg</i> ■ <i>/sybase</i> ■ <i>sybbin_dg_vol</i>
Resources for the Sybase ASE CE quorum device:	Example values for shared mount point:	
■ Disk group ■ Mount point ■ Volume	■ <i>quorum_dg</i> ■ <i>/quorum</i> ■ <i>quorum_vol</i>	
Resources for the Sybase ASE CE datafiles:	Example values for shared mount point:	

Table 16-1 Required resources for configuring Sybase ASE CE under VCS control (*continued*)

Required resources	Example values
■ Disk group ■ Mount point ■ Volume	■ <i>dbdata_dg</i> ■ <i>/sybdata</i> ■ <i>sybvol</i>
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 16-1](#)

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 19.
- 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 16-1](#) on page 164.
- 3 To select the type of file system where Sybase ASE CE binaries reside, choose one of the options.

Veritas 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 151. for shared mount point.

See [“Preparing for local mount point on VxFS for Sybase ASE CE binary installation”](#) on page 150. 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 152.

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 accessed now.
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
cvm	State	system1	ONLINE
cvm	State	system2	ONLINE
quorum_voldgmnts	State	system1	ONLINE
quorum_voldgmnts	State	system2	ONLINE
sybasece	State	system1	ONLINE
sybasece	State	system2	ONLINE
sybbin_voldgmnts	State	system1	ONLINE
sybbin_voldgmnts	State	system2	ONLINE

Upgrading Sybase ASE CE

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

For information on upgrading Sybase ASE CE software, see the Sybase ASE CE product documentation.

Adding and removing nodes

- [Chapter 17. Adding a node to SF Sybase CE clusters](#)
- [Chapter 18. 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 Veritas InfoScale 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 17-1 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 171.

Table 17-1 Tasks for adding a node to a cluster *(continued)*

Step	Description
Add a new node to the cluster.	See “Adding a node to a cluster using the Veritas InfoScale installer” on page 183. See “Adding the node to a cluster manually” on page 174.
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 180.
Prepare the new node for installing Sybase.	See “Adding the new instance to the Sybase ASE CE cluster” on page 186.
Add the node to Sybase.	See “Adding the new instance to the Sybase ASE CE cluster” on page 186.

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 installed with Enterprise 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:


```
# vxdctl protocolversion
Cluster running at protocol 160
```

- 5 If the cluster protocol on the master node is below 160, upgrade it using:

```
# vxdctl upgrade [version]
```

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 17-1](#).

Figure 17-1 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 17-1 illustrates a new node being added to an existing two-node cluster using two independent hubs.

2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.
- The node must have private network connections to two independent switches for the cluster.
For more information, see the *Cluster Server Configuration and Upgrade 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** Navigate to the folder that contains the installer program. Verify that the new node meets installation requirements. Verify that the new node meets installation requirements.

```
# ./installer -precheck
```

- 2** Install Veritas InfoScale Enterprise packages only without configuration on the new system. Make sure all the VRTS packages available on the existing nodes are also available on the new node.

```
# ./installer
```

Do not configure SF Sybase CE when prompted.

```
Would you like to configure InfoScale Enterprise after installation?  
[y,n,q] (n) n
```

- 3** Restart the node if the installer asks for a reboot.

Adding the node to a cluster manually

Perform this procedure after you install Veritas InfoScale Enterprise only if you plan to add the node to the cluster manually.

Table 17-2 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 174.
Configure the cluster processes on the new node.	See “Configuring cluster processes on the new node” on page 175.
Configure fencing for the new node to match the fencing configuration on the existing cluster.	See “Starting fencing on the new node” on page 180.
Start VCS.	See “To start VCS on the new node” on page 182.
Configure CVM and CFS.	See “Configuring Cluster Volume Manager (CVM) and Cluster File System (CFS) on the new node” on page 180.
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 182.

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 `installer` 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 bge1 /dev/bge:1 - ether - -
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where N represents the number of systems in the cluster 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/default/llt  
/etc/default/gab  
/etc/default/vxfen  
/etc/default/vcs
```

Verify if the attributes in each file are set as follows before using `smf` on Solaris 10 to start the related processes and to load drivers:

```
LLT_START/LLT_STOP=1  
GAB_START/GAB_STOP=1  
VXFEN_START/VXFEN_STOP=1  
VCS_START/VCS_STOP=1
```

- 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:

```
# svcadm enable lltd
# svcadm enable gab
```

10 On the new node, verify GAB port membership:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen df204 membership 012
```

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

[Table 17-3](#) uses the following information for the following command examples.

Table 17-3 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/VRTSvc/vcsauth/bkup  
# cd /tmp; gunzip -c /opt/VRTSvc/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvc/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvc/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \
./broker_setup.sh/tmp/eat_setup

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \
/VRTSatlocal.conf -b 'Security\Authentication \
\Authentication Broker' -k UpdatedDebugLogFileName \
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

4 Copy the broker credentials from one node in the cluster to sys5 by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

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`, 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/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

2 Start fencing on the new node:

```
# svcadm enable vxfen
```

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:

```
# hagrps -modify cvm SystemList -add system3 2
# hagrps -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 group online.

- 2 Verify that the CVM group is 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 bge0 -sys system3
```

```
# hares -modify gconic Device bge0 -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 Veritas InfoScale installer

You can add a node to a cluster using the `-addnode` option with the Veritas InfoScale installer.

The Veritas InfoScale installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages 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
```

- Updates and copies the following files to the new node from the existing 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/default/llt`
 - `/etc/default/gab`
 - `/etc/default/vxfen`
- Configures fencing.
- Adds the new node to the CVM, ClusterService 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 Veritas InfoScale installer with the `-addnode` option.

```
# cd /opt/VRTS/install
# ./installer -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing SF Sybase CE cluster.
The installer uses the node information to identify the existing cluster.

Enter the name of any one node of the InfoScale ENTERPRISE cluster where 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 packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

If there are IP addresses already configured on the interface, confirm whether you want to use the interface as the first private heartbeat link.

Enter the NIC for the first private heartbeat
link on system3: [b,q,?] **bge1**

Enter the NIC for the second private heartbeat
link on system3: [b,q,?] **bge2**

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: **bge3**

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.

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 186.

- 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 186.
- 11** Confirm that the new node has joined the SF Sybase CE cluster using `lltstat -n` and `gabconfig -a` commands.

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 151.

See [“Preparing to create a Sybase ASE CE cluster”](#) on page 152.

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](#)
- [Modifying the Cluster Volume Manager \(CVM\) configuration on the existing nodes to remove references to the deleted node](#)
- [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 remove a node to an existing SF Sybase CE cluster.

Table 18-1 Tasks for removing a node from a cluster

Step	Description
Prepare to remove the node: ■ 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 191.

Table 18-1
 Tasks for removing a node from a cluster (*continued*)

Step	Description
Remove the node from the cluster.	See “Removing a node from a cluster” on page 191.
Modify the cluster configuration on remaining nodes. ■ Edit the /etc/llthosts file. ■ Edit the /etc/gabtab file. ■ Modify the VCS configuration to remove the node. ■ Modify the CVM configuration to remove the node.	See “Modifying the VCS configuration files on existing nodes” on page 192. See “Modifying the Cluster Volume Manager (CVM) configuration on the existing nodes to remove references to the deleted node” on page 195.
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 195.

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 File 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 Stop SF Sybase CE on the node using the Veritas InfoScale Enterprise installer.

```
# cd /opt/VRTS/install
```

```
# ./installer -stop system3
```

The installer stops all SF Sybase CE processes.
- 4 Modify the VCS configuration files on the existing nodes to remove references to the deleted node.
 See [“Modifying the VCS configuration files on existing nodes”](#) on page 192.
- 5 Modify the Cluster Volume Manager (CVM) configuration on the existing nodes to remove references to the deleted node.
 See [“Modifying the Cluster Volume Manager \(CVM\) configuration on the existing nodes to remove references to the deleted node”](#) on page 195.

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/llthosts 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
```

Modifying the Cluster Volume Manager (CVM) configuration on the existing nodes to remove references to the deleted node

- 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.

Modifying the Cluster Volume Manager (CVM) configuration on the existing nodes to remove references to the deleted node

To modify the CVM configuration on the existing nodes to remove references to the deleted node

- ◆ On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

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 19. 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 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 Volume Replicator (VVR) for replication.

For more about planning for disaster recovery environments:

See [“Supported replication technologies for global clusters”](#) on page 14.

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 198.

See [“About configuring a parallel global cluster using Volume Replicator \(VVR\) for replication”](#) on page 198.

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

See the *Veritas InfoScale™ 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 the *Veritas InfoScale Disaster Recovery Implementation Guide* to install and configure SF Sybase CE on each cluster.

- 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

About configuring a parallel global cluster using Volume Replicator (VVR) for replication

Configuring a global cluster for environment with SF Sybase CE and 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.

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 clusters using VVR:

See the *Veritas InfoScale Disaster Recovery Implementation Guide*.

See the *Storage Foundation for Sybase ASE CE Configuration and Upgrade 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.

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 InfoScale Disaster Recovery Implementation Guide*.

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Sample installation and configuration values](#)
- [Appendix C. Tunable files for installation](#)
- [Appendix D. Configuration files](#)
- [Appendix E. Configuring the secure shell or the remote shell for communications](#)
- [Appendix F. High availability agent information](#)

Installation scripts

This appendix includes the following topics:

- [Installation script options](#)
- [About using the postcheck option](#)

Installation script options

[Table A-1](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas InfoScale product scripts, except where otherwise noted.

Table A-1 Available command line options

Command Line Option	Function
-addnode	Adds a node to a high availability cluster.
-ai	The <code>-ai</code> option is supported on Solaris 11 only, and is used to generate Automated Installation manifest. This can be used by Solaris Automated Installation Server to install the Veritas InfoScale product, along with the Solaris 11 operation system.
-allpkgs	Displays all packages required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-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.
-comsetup	The <code>-comsetup</code> option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configcps	The <code>-configcps</code> option is used to configure CP server on a running system or cluster.
-configure	Configures the product after installation.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-fencing	Configures I/O fencing in a running cluster.
-fips	The <code>-fips</code> option is used to enable or disable security with fips mode on a running VCS cluster. It could only be used together with <code>-security</code> or <code>-securityonnode</code> option.
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-install	Used to install products on system
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-jumpstart <i>dir_path</i>	Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.
-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.
-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.
-noipc	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkgtable	Displays product's packages in correct installation order by group.
-postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
-component	Specifies the component for operations.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
-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 packages get upgraded to the latest version.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-rollingupgrade_phase2	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
-rootpath <i>root_path</i>	Specifies an alternative root directory on which to install packages. On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See “About configuring secure shell or remote shell communication modes before installing products” on page 253.
-security	The <code>-security</code> option is used to convert a running VCS cluster between secure and non-secure modes of operation.
-securityonenode	The <code>-securityonenode</code> option is used to configure a secure cluster node by node.
-securitytrust	The <code>-securitytrust</code> option is used to setup trust with another broker.
-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.
-settnables	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.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-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 packages 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.
-uninstall	This option is used to uninstall the products from systems
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.

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 VRTSIlt pkg version is not consistent on the nodes.
- The Ilt-linkinstall value is incorrect.
- The `/etc/llthosts` and `/etc/llttab` 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 `Vxfsn` 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 packages are installed.
- The versions of the required packages are correct.
- There are no verification issues for the required packages.

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 (`vxdg 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 `/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 `/etc/vfstab` file are mounted.
- Whether all VxFS file systems present in `/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 73.

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.

Veritas recommends using the worksheets provided to record values for your systems before you begin the installation and configuration process.

SF Sybase CE worksheet

[Table B-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 B-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 packages vs. all SF Sybase CE	Install only the required packages if you do not want to configure any optional components or features. Default option is to install all packages.	
Primary cluster name	clus1	
Primary cluster ID number	101	
Private network links You can choose a network interface card that is not part of any aggregated interface, or you can choose an aggregated interface. The interface names that are associated with each NIC for each network link must be the same on all nodes. Do not use the network interface card that is used for the public network, which is typically bge0.	bge1,bge2	
Cluster Manager NIC (Primary NIC)	bge0	
Cluster Manager IP	10.10.12.1, 10.10.12.2	
Netmask for the virtual IP address	255.255.240.0	

Table B-1 SF Sybase CE worksheet (*continued*)

Installation information	Sample value	Assigned value
Mode for Authentication Service: ■ Automatic mode ■ Semiautomatic mode using encrypted files ■ Semiautomatic mode without using encrypted files Default option is automatic mode.	Automatic mode	
User name Adding users is required if when using secure cluster mode. Otherwise it is optional.	smith	
User password	password	
User privilege VCS privilege levels include: ■ Administrators— Can perform all operations, including configuration options on the cluster, service groups, systems, resources, and users. ■ Operators—Can perform specific operations on a cluster or a service group. ■ Guests—Can view specified objects.	admin	
Domain-based address of the SMTP server The SMTP server sends notification email about the events within the cluster.	smtp.veritasexample.com	
Email address of each SMTP recipient to be notified	john@veritasexample.com	

Table B-1 SF Sybase CE worksheet (*continued*)

Installation information	Sample value	Assigned value
Minimum severity of events for SMTP email notification The severity levels are defined as follows: ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption	E	
Email address of SMTP notification recipients	admin@veritasexample.com	
SNMP trap daemon port number the console	162	
System name for the SNMP console	system2	
Minimum severity level of events for SMTP notification The severity levels are defined as follows: ■ Information - Important events that exhibit normal behavior ■ Warning - Deviation from normal behavior ■ Error - A fault ■ Severe Error -Critical error that can lead to data loss or corruption	i	
CVM enclosure-based naming Requires Dynamic Multi-pathing (DMP).	yes	

Table B-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.	vxfscoordg	
The name of three disks that form the coordinator disk group.	■ c1t1d0s2 ■ c2t1d0s2 ■ c3t1d0s2	
Vxfs disk group	vxfscoordg	

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 216.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys1 sys2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 217.

- 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 218.

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 220.

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 220.

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 219.
- 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
-settunables [sys1 sys2 ...]
```

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 220.

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 219.
- 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 `-settunables` option.

```
# ./installer -tunablesfile tunables_file_name -settunables [
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 220.

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 219.
- 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 220.

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 *Storage Foundation Cluster File System High Availability Administrator's Guide* for detailed information on product tunable ranges and recommendations.

[Table C-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table C-1 Supported tunable parameters

Tunable	Description
autoreminor	(Veritas Volume Manager) Enable reminoring in case of conflicts during disk group import.
autostartvolumes	(Veritas Volume Manager) Enable the automatic recovery of volumes.
dmp_cache_open	(Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached.
dmp_daemon_count	(Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_delayq_interval	(Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy.
dmp_fast_recovery	(Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Dynamic Multi-Pathing is started.
dmp_health_time	(Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy.
dmp_log_level	(Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed.
dmp_low_impact_probe	(Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled.
dmp_lun_retry_timeout	(Dynamic Multi-Pathing) The retry period for handling transient errors.
dmp_monitor_fabric	(Dynamic Multi-Pathing) Whether the Event Source daemon (<code>vxesd</code>) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored.
dmp_native_multipathing	(Dynamic Multi-Pathing) Whether DMP will intercept the I/Os directly on the raw OS paths or not.
dmp_native_support	(Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices.
dmp_path_age	(Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy.
dmp_pathswitch_blks_shift	(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.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_probe_idle_lun	(Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs.
dmp_probe_threshold	(Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon.
dmp_restore_cycles	(Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic.
dmp_restore_interval	(Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths.
dmp_restore_policy	(Dynamic Multi-Pathing) The policy used by DMP path restoration thread.
dmp_restore_state	(Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started.
dmp_retry_count	(Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed.
dmp_scsi_timeout	(Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP.
dmp_sfg_threshold	(Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature.
dmp_stat_interval	(Dynamic Multi-Pathing) The time interval between gathering DMP statistics.
fssmartmovethreshold	(Veritas Volume Manager) The file system usage threshold for SmartMove (percent). This tunable must be set after Veritas Volume Manager is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer can only set the system default value of max_diskq. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer can only set the system default value of read_ahead. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer can only set the system default value of read_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer can only set the system default value of read_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
reclaim_on_delete_start_time	(Veritas Volume Manager) Time of day to start reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
reclaim_on_delete_wait_period	(Veritas Volume Manager) Days to wait before starting reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
same_key_for_alldgs	(Veritas Volume Manager) Use the same fencing key for all disk groups. This tunable must be set after Veritas Volume Manager is started.
sharedminorstart	(Veritas Volume Manager) Start of range to use for minor numbers for shared disk groups. This tunable must be set after Veritas Volume Manager is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
storage_connectivity	(Veritas Volume Manager) The CVM storage connectivity type. This tunable must be set after Veritas Volume Manager is started.
usefssmartmove	(Veritas Volume Manager) Configure SmartMove feature (all, thinonly, none). This tunable must be set after Veritas Volume Manager is started.
vol_checkpoint_default	(Veritas File System) Size of VxVM storage checkpoints (kBytes). This tunable requires a system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires a 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 a 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 a 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 Volume Replicator .

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (kBytes). This tunable requires a system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires a system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires a system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (kBytes). This tunable requires a 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) Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by 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 a system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions. This tunable requires a system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires a system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (kBytes). This tunable requires a 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 a system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires a system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires a system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires a system reboot to take effect.
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires a system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires a system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires a 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 can only set the system default value of write_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer can only set the system default value of write_pref_io. Refer to the tuneftab(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/vfstab`
- 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 "OracleASMTypes.cf"
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster sybasece (
    SecureClus = 1
    DefaultGuestAccess = 1
    UseFence = SCSI3
    HacliUserLevel = COMMANDROOT
)

system sflorat52-1-v03 (
)

system sflorat52-2-v04 (
)

group cvm (
    SystemList = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sflorat52-1-v03, sflorat52-2-v04 }
```

```

    )

    CFSfsckd vxfsckd (
        ActivationMode @sflorat52-1-v03 = { sybbin_dg = sw,
quorum_dg = sw }
        ActivationMode @sflorat52-2-v04 = { sybbin_dg = sw,
quorum_dg = sw }
    )

    CVMCluster cvm_clus (
        CVMClustName = sybasece
        CVMNodeId = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
        CVMTransport = gab
        CVMTimeout = 200
    )

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

    ProcessOnOnly vxattachd (
        Critical = 0
        PathName = "/sbin/sh"
        Arguments = "- /usr/lib/vxvm/bin/vxattachd root"
        RestartLimit = 3
    )

    cvm_clus requires cvm_vxconfigd
    vxfsckd requires cvm_clus

    // resource dependency tree
    //
    //     group cvm
    //     {
    //         ProcessOnOnly vxattachd
    //         CFSfsckd vxfsckd
    //         {
    //             CVMCluster cvm_clus
    //             {
    //                 CVMVxconfigd cvm_vxconfigd
    //             }
    //         }
    //     }

```

```

//          }
//          }

group quorum_voldgmnts (
    SystemList = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sflorat52-1-v03, sflorat52-2-v04 }
)

CFSMount cfsmount2 (
    Critical = 0
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_dg/quorum_vol"
    MountOpt @sflorat52-1-v03 = rw
    MountOpt @sflorat52-2-v04 = rw
    NodeList = { sflorat52-1-v03, sflorat52-2-v04 }
)

CFSMount cfsmount3 (
    Critical = 0
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/quorum_dg/sybvvol"
    MountOpt @sflorat52-1-v03 = rw
    MountOpt @sflorat52-2-v04 = rw
    NodeList = { sflorat52-1-v03, sflorat52-2-v04 }
)

CVMVolDg cvmvoldg2 (
    Critical = 0
    CVMDiskGroup = quorum_dg
    CVMVolume = { quorum_vol, sybvvol }
    CVMActivation @sflorat52-1-v03 = sw
    CVMActivation @sflorat52-2-v04 = sw
)

requires group cvm online local firm
cfsmount2 requires cvmvoldg2
cfsmount3 requires cvmvoldg2

// resource dependency tree

```

```
//
//      group quorum_voldgmnts
//      {
//      CFSMount cfsmount2
//      {
//      CVMVolDg cvmvoldg2
//      }
//      CFSMount cfsmount3
//      {
//      CVMVolDg cvmvoldg2
//      }
//      }

group sybasece (
    SystemList = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sflorat52-1-v03, sflorat52-2-v04 }
    TriggerResStateChange = 1
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount quorum_dg_quorum_vol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_dg/quorum_vol"
)

CFSMount quorum_dg_sybvol_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/quorum_dg/sybvol"
)

CVMVolDg quorum_dg_voldg (
    CVMDiskGroup = quorum_dg
    CVMVolume = { sybvol, quorum_vol }
    CVMActivation = sw
)

Process vxxfend (
    PathName = "/sbin/vxxfend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
```

```

    )

Sybase ase (
    Server @sflorat52-1-v03 = asev103
    Server @sflorat52-2-v04 = asev204
    Owner = sybase
    Home = "/sybase"
    Version = sybase15
    SA = sa
    Quorum_dev = "/quorum/quorum_device"
)

requires group binmnt online local firm
ase requires quorum_dg_quorum_vol_mnt
ase requires quorum_dg_sybvol_mnt
ase requires vxfsend
quorum_dg_quorum_vol_mnt requires quorum_dg_voldg
quorum_dg_sybvol_mnt requires quorum_dg_voldg

// resource dependency tree
//
//      group sybasece
//      {
//      Sybase ase
//      {
//          CFSSMount quorum_dg_sybvol_mnt
//          {
//              CVMVolDg quorum_dg_voldg
//          }
//          CFSSMount quorum_dg_quorum_vol_mnt
//          {
//              CVMVolDg quorum_dg_voldg
//          }
//          Process vxfsend
//      }
//      }

group sybbin_voldgmnts (
    SystemList = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
    AutoFailOver = 0
    Parallel = 1

```

```

AutoStartList = { sflorat52-1-v03, sflorat52-2-v04 }
)

CFSMount cfsmount1 (
    Critical = 0
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbin_dg/sybbin_vol"
    MountOpt @sflorat52-1-v03 = rw
    MountOpt @sflorat52-2-v04 = rw
    NodeList = { sflorat52-1-v03, sflorat52-2-v04 }
)

CVMVolDg cvmvoldg1 (
    Critical = 0
    CVMDiskGroup = sybbin_dg
    CVMVolume = { sybbin_vol }
    CVMActivation @sflorat52-1-v03 = sw
    CVMActivation @sflorat52-2-v04 = sw
)

requires group cvm online local firm
cfsmount1 requires cvmvoldg1

// resource dependency tree
//
//      group sybbin_voldgmnts
//      {
//      CFSMount cfsmount1
//      {
//      CVMVolDg cvmvoldg1
//      }
//      }

group vxfen (
    SystemList = { sflorat52-1-v03 = 0, sflorat52-2-v04 = 1 }
    AutoFailOver = 0
    Parallel = 1
)

CoordPoint coordpoint (
    LevelTwoMonitorFreq = 5

```

```

    )

    Phantom RES_phantom_vxfen (
    )

// resource dependency tree
//
//      group vxfen
//      {
//      Phantom RES_phantom_vxfen
//      CoordPoint coordpoint
//      }

```

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 sybbin_dg_voldg (
    DiskGroup = sybbindg
)

Mount sybbin_dg_sybbin_vol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbin_dg/sybbin_vol"
    FSType = vxfs
    FsckOpt = "-y"
)

Volume sybbin_dg_vol (
    DiskGroup = sybbindg
    Volume = sybbin_vol
)

requires group cvm online local firm
sybbin_dg_sybbin_vol_mnt requires sybbin_dg_vol
sybbin_dg_vol requires sybbin_dg_voldgdg
```

```
// resource dependency tree
//
// group binlocalmnt
// {
// Mount sybbin_dg_sybbin_vol_mnt
// {
//     Volume sybbindg_vol
//     {
//         DiskGroup sybbin_dg_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 vxfsend process.
// 4. Sybase database instance.

group sybasece (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount quorum_dg_quorum_vol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_dg/quorum_vol"
)

CFSMount dbdata_dg_sybvol_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/dbdata_dg/sybvol"
)

CVMVolDg quorum_dg_voldg (
    CVMDiskGroup = quorum_dg
    CVMVolume = { quorum_vol }
    CVMActivation = sw
)

CVMVolDg dbdata_dg_voldg (
    CVMDiskGroup = dbdata_dg
    CVMVolume = { sybvol }
    CVMActivation = sw
)

Process vxfsend (
    PathName = "/sbin/vxfsend"
    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_dg_quorum_vol_mnt
ase requires dbdata_dg_sybvol_mnt
ase requires vxfeed
quorum_dg_quorum_vol_mnt requires quorum_dg_voldg
dbdata_dg_sybvol_mnt requires dbdata_dg_voldg

// resource dependency tree
//
// group sybasece
// {
//   Sybase ase
//   {
//     CFSSMount quorum_dg_quorum_vol_mnt
//     {
//       CVMVoldg quorum_dg_voldg
//     }
//     CFSSMount dbdata_dg_sybvol_mnt
//     {
//       CVMVoldg dbdata_dg_voldg
//     }
//     Process vxfeed
//   }
// }

```

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 "CVMTTypes.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 = bge0
        Address = "10.180.88.188"
        NetMask = "255.255.255.0"
    )

    NIC csgnic (
        Device = bge0
    )

    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 dbdata_voldg (
    CVMDiskGroup = dbdata_dg
    CVMActivation = sw
)

RVGShared dbdata_rvg (

```

```

        RVG = syb_rvg
        DiskGroup = dbdata_dg
    )

requires group binmnt online local firm
dbdata_rvg requires dbdata_voldg

group binmnt (
    SystemList = { system1 = 0, system2 = 1 }
    Parallel = 1
    AutoStartList = { system1, system2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount sybbin_dg_sybbin_vol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbin_dg/sybbin_vol"
)

CVMVolDg sybbin_dg_voldg (
    CVMDiskGroup = sybbin_dg
    CVMVolume = { sybbin_vol }
    CVMActivation = sw
)

requires group cvm online local firm
sybbin_dg_sybbin_vol_mnt requires sybbin_dg_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 = bge0
    Address = "10.10.9.101"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = bge0
)

RVGLogowner rvg_logowner (
    RVG = syb_rvg

```

```

        DiskGroup = dbdata_dg
    )

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_dg_quorum_vol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_dg/quorum_vol"
)

CFSMount dbdata_dg_sybvol_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/dbdata_dg/sybvol"
)

```

```
CVMVolDg quorum_dg_voldg (
    CVMDiskGroup = quorum_dg
    CVMVolume = { quorum_vol }
    CVMActivation = sw
)

Process vxfsend (
    PathName = "/sbin/vxfsend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

RVGSharedPri syb_vvr_shpri (
    RvgResourceName = dbdata_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
dbdata_dg_sybvol_mnt requires syb_vvr_shpri
ase requires vxfsend
ase requires dbdata_dg_sybvol_mnt
ase requires quorum_dg_quorum_vol_mnt
quorum_dg_quorum_vol_mnt requires quorum_dg_voldg

// resource dependency tree
//
//      group sybasece
//      {
//      Sybase ase
//      {
//          CFSMount dbdata_dg_sybvol_mnt
//          {
//              RVGSharedPri syb_vvr_shpri
//          }
//      }
//  }
```

```
//          Process vxfsend
//          CFSSMount quorum_dg_quorum_vol_mnt
//          {
//          CVMVolDg quorum_dg_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 "CVMTTypes.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 = bge0
            Address = "10.190.99.199"
            NetMask = "255.255.255.0"
        )

            NIC csgnic (
                Device = bge0
            )

            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 dbdata_voldg (
    CVMDiskGroup = dbdata_dg
    CVMActivation = sw
)

RVGShared dbdata_rvg (
    RVG = syb_rvg
    DiskGroup = dbdata_dg
)

requires group binmnt online local firm
dbdata_rvg requires dbdata_voldg

group binmnt (
    SystemList = { system3 = 0, system4 = 1 }
    Parallel = 1
    AutoStartList = { system3, system4 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

CFSMount sybbin_dg_sybbin_vol_mnt (
    MountPoint = "/sybase"
    BlockDevice = "/dev/vx/dsk/sybbin_dg/sybbin_vol"
)

CVMVolDg sybbin_dg_voldg (
    CVMDiskGroup = sybbin_dg
    CVMVolume = { sybbin_vol }
    CVMActivation = sw
)

requires group cvm online local firm
sybbin_dg_sybbin_vol_mnt requires sybbin_dg_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 = bge0
        Address = "10.11.9.102"
        NetMask = "255.255.255.0"
    )

    NIC nic (
        Device = bge0
    )

    RVGLogowner rvg_logowner (
        RVG = syb_rvg
        DiskGroup = dbdata_dg
    )

    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_dg_quorum_vol_mnt (
    MountPoint = "/quorum"
    BlockDevice = "/dev/vx/dsk/quorum_dg/quorum_vol"
)

CVMVolDg quorum_dg_voldg (
    CVMDiskGroup = quorum_dg
    CVMVolume = { quorum_vol }
    CVMActivation = sw
)

CFSMount dbdata_dg_sybvol_mnt (
    MountPoint = "/sybdata"
    BlockDevice = "/dev/vx/dsk/dbdata_dg/sybvol"
)

Process vxfsend (
    PathName = "/sbin/vxfsend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

RVGSharedPri syb_vvr_shpri (
    RvgResourceName = dbdata_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
dbdata_dg_sybvol_mnt requires syb_vvr_shpri
ase requires vxfsend
ase requires dbdata_dg_sybvol_mnt
ase requires quorum_dg_quorum_vol_mnt

```

```
quorum_dg_quorum_vol_mnt requires quorum_dg_voldg
```

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- [About configuring secure shell or remote shell communication modes before installing products](#)

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas InfoScale software from a remote system, or to install and configure a system. The system from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install the Veritas InfoScale software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). Veritas recommends that you use `ssh` as it is more secure than `rsh`.

You can set up `ssh` and `rsh` connections in many ways.

- You can manually set up the `ssh` and `rsh` connection with UNIX shell commands.
- You can run the `installer -comsetup` command to interactively set up `ssh` and `rsh` connection.
- You can run the password utility, `pwdutil.pl`.

This section contains an example of how to set up `ssh` password free communication. The example sets up `ssh` between a source system (`sys1`) that contains the

installation directories, and a target system (sys2). This procedure also applies to multiple target systems.

Note: The product installer supports establishing passwordless communication.


Manually configuring passwordless ssh

The ssh program enables you to log into and execute commands on a remote system. ssh enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

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

Figure E-1 illustrates this procedure.

Figure E-1 Creating the DSA key pair and appending it to target systems



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

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

To create the DSA key pair

- 1 On the source system (sys1), log in as root, and navigate to the root directory.

```
sys1 # cd /
```

- 2 Make sure the `/.ssh` directory is on all the target installation systems (sys2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

Solaris 10:

```
sys2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

Solaris 10:

```
sys2 # chmod go-w /.ssh
```

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

```
sys1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.
```

```
Enter file in which to save the key (//.ssh/id_dsa):
```

- 4 Press Enter to accept the default location of `/.ssh/id_dsa`.
- 5 When the program asks you to enter the passphrase, press the Enter key twice.

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

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

- 1** Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (sys2 in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem                sftp          /usr/lib/ssh/sftp-server
```

- 2** If the lines are not there, add them and restart ssh.

To restart ssh on Solaris 10, type the following command:

```
sys1 # svcadm restart ssh
```

- 3** From the source system (sys1), move the public key to a temporary file on the target system (sys2).

Use the secure file transfer program.

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

Use the following command for secure file transfer:

```
sys1 # sftp sys2
```

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

```
Connecting to sys2 ...
The authenticity of host 'sys2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 4** Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'sys2,10.182.00.00'
(DSA) to the list of known hosts.
root@sys2 password:
```

- 5** Enter the root password of sys2.

- 6 At the `sftp` prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

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

```
sftp> quit
```

- 8 To begin the `ssh` session on the target system (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

- 9 After you log in to sys2, enter the following command to append the `id_dsa.pub` file to the `authorized_keys` file:

```
sys2 # cat /id_dsa.pub >> /.ssh/authorized_keys
```

- 10 After the `id_dsa.pub` public key file is copied to the target system (sys2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on sys2:

```
sys2 # rm /id_dsa.pub
```

- 11 To log out of the `ssh` session, enter the following command:

```
sys2 # exit
```

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

```
sys1 # exec /usr/bin/ssh-agent $SHELL
sys1 # ssh-add
```

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

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

To verify that you can connect to a target system

- 1 On the source system (sys1), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where `sys2` is the name of the target system.

- 2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

Setting up `ssh` and `rsh` connection using the `installer -comsetup` command

You can interactively set up the `ssh` and `rsh` connections using the `installer -comsetup` command.

Enter the following:

```
# ./installer -comsetup
```

```
Input the name of the systems to set up communication:
```

```
Enter the <platform> system names separated by spaces:
```

```
[q,?] sys2
```

```
Set up communication for the system sys2:
```

```
Checking communication on sys2 ..... Failed
```

```
CPI ERROR V-9-20-1303 ssh permission was denied on sys2. rsh
```

permission was denied on sys2. Either ssh or rsh is required to be set up and ensure that it is working properly between the local node and sys2 for communication

Either ssh or rsh needs to be set up between the local system and sys2 for communication

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) y

Enter the superuser password for system sys2:

- 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) 1

Setting up communication between systems. Please wait.
 Re-verifying systems.

Checking communication on sys2 Done

Successfully set up communication for the system sys2

Setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled under the `scripts` directory. The users can run the utility in their script to set up the ssh and rsh connection automatically.

```
# ./pwdutil.pl -h
```

Usage:

Command syntax with simple format:

```
pwdutil.pl check|configure|unconfigure ssh|rsh <hostname|IP addr>
[<user>] [<password>] [<port>]
```

Command syntax with advanced format:

```
pwdutil.pl [--action|-a 'check|configure|unconfigure']
```

```
[--type|-t 'ssh|ssh|ssh']
[--user|-u '<user>']
[--password|-p '<password>']
[--port|-P '<port>']
[--hostfile|-f '<hostfile>']
[--keyfile|-k '<keyfile>']
[-debug|-d]
<host_URI>
```

```
pwdutil.pl -h | -?
```

Table E-1 Options with pwdutil.pl utility

Option	Usage
--action -a 'check configure unconfigure'	Specifies action type, default is 'check'.
--type -t 'ssh ssh ssh'	Specifies connection type, default is 'ssh'.
--user -u '<user>'	Specifies user id, default is the local user id.
--password -p '<password>'	Specifies user password, default is the user id.
--port -P '<port>'	Specifies port number for ssh connection, default is 22
--keyfile -k '<keyfile>'	Specifies the private key file.
--hostfile -f '<hostfile>'	Specifies the file which list the hosts.
-debug	Prints debug information.
-h -?	Prints help messages.
<host_URI>	Can be in the following formats: <hostname> <user>:<password>@<hostname> <user>:<password>@<hostname>: <port>

You can check, configure, and unconfigure ssh or rsh using the `pwdutil.pl` utility. For example:

- To check ssh connection for only one host:

```
pwdutil.pl check ssh hostname
```

- To configure ssh for only one host:

```
pwdutil.pl configure ssh hostname user password
```

- To unconfigure rsh for only one host:

```
pwdutil.pl unconfigure rsh hostname
```

- To configure ssh for multiple hosts with same user ID and password:

```
pwdutil.pl -a configure -t ssh -u user -p password hostname1
hostname2 hostname3
```

- To configure ssh or rsh for different hosts with different user ID and password:

```
pwdutil.pl -a configure -t ssh user1:password1@hostname1
user2:password2@hostname2
```

- To check or configure ssh or rsh for multiple hosts with one configuration file:

```
pwdutil.pl -a configure -t ssh --hostfile /tmp/sshrsh_hostfile
```

- To keep the host configuration file secret, you can use the 3rd party utility to encrypt and decrypt the host file with password.

For example:

```
### run openssl to encrypt the host file in base64 format
# openssl aes-256-cbc -a -salt -in /hostfile -out /hostfile.enc
enter aes-256-cbc encryption password: <password>
Verifying - enter aes-256-cbc encryption password: <password>
```

```
### remove the original plain text file
# rm /hostfile
```

```
### run openssl to decrypt the encrypted host file
# pwdutil.pl -a configure -t ssh `openssl aes-256-cbc -d -a
-in /hostfile.enc`
enter aes-256-cbc decryption password: <password>
```

- To use the ssh authentication keys which are not under the default `$HOME/.ssh` directory, you can use `--keyfile` option to specify the ssh keys. For example:

```
### create a directory to host the key pairs:
# mkdir /keystore
```

```
### generate private and public key pair under the directory:
# ssh-keygen -t rsa -f /keystore/id_rsa

### setup ssh connection with the new generated key pair under
the directory:
# pwdutil.pl -a configure -t ssh --keyfile /keystore/id_rsa
user:password@hostname
```

You can see the contents of the configuration file by using the following command:

```
# cat /tmp/sshrsh_hostfile
user1:password1@hostname1
user2:password2@hostname2
user3:password3@hostname3
user4:password4@hostname4

# all default: check ssh connection with local user
hostname5
The following exit values are returned:

0      Successful completion.
1      Command syntax error.
2      Ssh or rsh binaries do not exist.
3      Ssh or rsh service is down on the remote machine.
4      Ssh or rsh command execution is denied due to password is required.
5      Invalid password is provided.
255    Other unknown error.
```

Restarting the ssh session

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

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

To restart ssh

- 1 On the source installation system (sys1), bring the private key into the shell environment.

```
sys1 # exec /usr/bin/ssh-agent $SHELL
```

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

```
sys1 # ssh-add
```

Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

Veritas recommends configuring a secure shell environment for Veritas InfoScale product installations.

See [“Manually configuring passwordless ssh”](#) on page 254.

See the operating system documentation for more information on configuring remote shell.

To enable rsh

- 1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
# inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

- 2 To enable a disabled `rsh/rlogin` service, type the following command:

```
# inetadm -e rlogin
```

- 3 To disable an enabled `rsh/rlogin` service, type the following command:

```
# inetadm -d rlogin
```

- 4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using `rsh`. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `sys1` from `sys2`, you must add an entry for `sys2.companyname.com` in the `.rhosts` file on `sys1`.

```
# echo "sys2.companyname.com" >> $HOME/.rhosts
```

- 5 After you complete an installation procedure, delete the `.rhosts` file from each user's `$HOME` directory to ensure security:

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

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 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
- CFSSMount 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 CFSSMount agent for each volume in the disk group.

Use the information in this appendix about the entry points and attributes of the listed agents to make necessary configuration changes. For information on how to modify the VCS configuration:

See the *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 *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. ■ Type and dimension: string-scalar
CVMNodeAddr	List of host names and IP addresses. ■ Type and dimension: string-association
CVMNodeId	Associative list. The first part names the system; the second part contains the LLT ID number for the system. ■ Type and dimension: string-association

Table F-2 CVMCluster agent attributes (*continued*)

Attribute	Description
CVMTransport	Specifies the cluster messaging mechanism. ■ Type and dimension: string-scalar ■ Default = gab Note: Do not change this value.
PortConfigd	The port number that is used by CVM for vxconfigd-level communication. ■ Type and dimension: integer-scalar
PortKmsgd	The port number that is used by CVM for kernel-level communication. ■ Type and dimension: integer-scalar
CVMTimeout	Timeout in seconds used for CVM cluster reconfiguration. ■ Type and dimension: integer-scalar ■ Default = 200
CVMNodePreference	The preference value that determines which nodes in the CVM cluster are the most appropriate candidates to run the master role. The preference values are in the range from -2147483648 to 2147483647. If you do not specify custom preferences, CVM gives preference to the node with the maximum visibility to the storage to become the CVM master node. ■ Type and dimension: string ■ Default = 0

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```
type CVMCluster (
    static keylist RegList = { CVMNodePreference, CVMDGSubClust }
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName, CVMNodeAddr, CVMNodeId,
                            PortConfigd, PortKmsgd, CVMTimeout, CVMDGSubClust }

    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
```

```
str CVMNodePreference
int PortConfigd
int PortKmsgd
int CVMTimeout
boolean CVMDGSubClust = 0
```

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.

Veritas 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 *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
imf_init	Initializes the agent to interface with the AMF kernel module. This function runs when the agent starts up.
imf_getnotification	Gets notification about the <code>vxconfigd</code> process state. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification. If the <code>vxconfigd</code> process fails, the function initiates a traditional CVMVxconfigd monitor entry point.
imf_register	Registers or unregisters the <code>vxconfigd</code> process id (pid) with the AMF kernel module. This function runs after the resource goes into steady online state.

Attribute definition for CVMVxconfigd agent

Table F-4 describes the modifiable attributes of the CVMVxconfigd resource type.

Table F-4 CVMVxconfigd agent attribute

Attribute	Description
CVMVxconfigdArgs	List of the arguments that are sent to the <code>online</code> entry point. Veritas recommends always specifying the <code>syslog</code> option. ■ Type and dimension: keylist

Table F-4 CVMVxconfigd agent attribute (*continued*)

Attribute	Description
IMF	This resource-type level attribute determines whether the CVMVxconfigd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association For more details of IMF attribute for the agent type, refer to the <i>Cluster Server Administrator's Guide</i>.

CVMVxconfigd agent type definition

The following type definition is included in the CVMTypes.cf file:

```
type CVMVxconfigd (
    static int 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 228.

CVMVoIDg agent

The CVMVoIDg 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 CVMVoIDg agent manages the CVM disk groups and CVM volumes and volume sets within the disk groups by performing the following functions:

- Imports the shared disk group from the CVM master node
- Starts the volumes and volume sets in the disk group
- Monitors the disk group, volumes, and volume sets
- Optionally, deports the disk group when the dependent applications are taken offline. The agent deports the disk group only if the appropriate attribute is set.

Configure the CVMVoIDg agent for each disk group used by a 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 CVMVoIDg agent

[Table F-5](#) describes the entry points used by the CVMVoIDg agent.

Table F-5 CVMVolDg agent entry points

Entry Point	Description
Online	Starts all volumes in the shared disk group specified by the CVMVolume attribute. Imports the shared disk group from the CVM master node, if the disk group is not already imported. Sets the disk group activation mode to shared-write if the value of the CVMActivation attribute is sw. You can set the activation mode on both slave and master systems.
Offline	Sets the activation mode of the shared disk group to “off.” If the <code>CVMDeportOnOffline</code> attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.
Monitor	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. The agent takes a volume set offline if the file system metadata volume of a volume set is discovered to be offline in a monitor cycle. Note: If the CFSMount resource goes offline and the file system on the volume set is unmounted, the agent retains the online state of the volume set even if the file system metadata volume in the volume set is offline. This is because the CVMVolDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.
Clean	Removes the temporary files created by the online entry point.

Attribute definition for 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)	Shared disk group name. ■ Type and dimension: string-scalar

Table F-6 CVMVolDg agent attributes (*continued*)

Attribute	Description
CVMVolume (required)	Lists critical volumes in the disk group. SF Sybase CE requires specifying at least one volume in the disk group. ■ Type and dimension: string-keylist
CVMActivation (required)	Activation mode for the disk group. ■ Type and dimension: string-scalar ■ Default = <code>sw</code> (<code>shared-write</code>) This is a localized attribute.
CVMDeportOnOffline (optional)	Indicates whether or not the shared disk group must be deported when the last online CVMVolDg resource for a disk group is taken offline. The value 1 indicates that the agent will deport the shared disk group from the CVM master node, if not already deported, when the last online CVMVolDg resource for the disk group is taken offline. The value 0 indicates that the agent will not deport the shared disk group when the CVMVolDg resource is taken offline. ■ Type and dimension: integer-scalar ■ Default = 0 Note: If multiple CVMVolDg resources are configured for a shared disk group, set the value of the attribute to either 1 or 0 for all of the resources. The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the <code>CVMDeportOnOffline</code> attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported. The deport operation fails if the shared disk group contains open volumes.
ClearClone	Indicates whether or not the disk group is imported with the <code>-c</code> option. The <code>-c</code> option clears the clone and the <code>udid_mismatch</code> flags from the disks of the disk groups and also updates the UDID when you import the disk group. ■ Type and dimension: integer-scalar ■ Default = 0

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 @sys1 = sw
    CVMActivation @sys2 = sw
    CVMDeportOnOffline = 1
)

CVMVolDg sybbin_dg_voldg (
    CVMDiskGroup = sybbin_dg
    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 *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.
imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.
imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSMount agent

Table 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. Type and dimension: string-scalar
BlockDevice	Block device for the mount point. Type and dimension: string-scalar

Table F-8 CFSMount Agent attributes (*continued*)

Attribute	Description
NodeList	List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list. ■ Type and dimension: string-keylist
IMF	Resource-type level attribute that determines whether the CFSMount agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association

Table F-8 CFSMount Agent attributes (*continued*)

Attribute	Description
MountOpt (optional)	Options for the mount command. To create a valid MountOpt attribute string: ■ Use the VxFS type-specific options only. ■ Do not use the -o flag to specify the VxFS-specific options. ■ Do not use the -F vxfs file system type option. ■ Be aware the cluster option is not required. ■ Specify options in comma-separated list: <pre>ro ro,cluster blkclear,mincache=closesync</pre> ■ Type and dimension: string-scalar
Policy (optional)	List of nodes to assume the primaryship of the cluster file system if the primary fails. If set to NULL or if none of the hosts specified in the list is active when the primary fails, a node is randomly selected from the set of active nodes to assume primaryship. ■ Type and dimension: string-scalar

CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition:

```
type CFSMount (
    static 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/sybbin_vol"
    Primary = sys2;
)
```

To see CFSMount defined in a more extensive example:

See [“About sample main.cf files”](#) on page 228.

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 the process with optional arguments.
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. The agent only reports the process as online if the value configured for PathName attribute exactly matches the process listing from the ps output.
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. This attribute must not exceed 80 characters. Type and dimension: string-scalar Example: "/usr/lib/sendmail"

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. This attribute must not exceed 80 characters. Type and dimension: string-scalar Example: "bd -q1h"
ContainerType	Do not change. For internal use only.

Resource type definition

```

type Process (
    static keylist SupportedActions = { "program.vfd", getcksum }
    static str ContainerType = Zone
    static str ArgList[] = { ContainerName, PathName, Arguments }
    str ContainerName
    str PathName
    str Arguments
)

```

Sample configurations

```

Process vxfend (
    PathName = "/sbin/vxfend"
    Arguments = "-m sybase -k /tmp/vcmp_socket"
)

```

Monitoring options for the Sybase agent

The VCS agent for Sybase provides two levels of application monitoring: basic and detail.

In the basic monitoring mode, the agent for Sybase monitors the Sybase dataserver 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 server is functioning properly. The test table should be created by the user, and the table is specified in the attribute `Table` for the Sybase agent. The agent uses this test table for internal purposes. Veritas recommends that you do not perform any other transaction on the test table.

See [“About setting up detail monitoring for the agent for Sybase”](#) on page 158.

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
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,
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 int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }
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
Home	The \$SYBASE path to Sybase binaries and configuration files. Type and dimension: string-scalar Default value: No default value

Table F-11 Required attributes (*continued*)

Required Attributes	Definition
Owner	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. Type and dimension: string-scalar
Quorum_dev	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. Type and dimension: string-scalar Default value: No default value Note: This attribute should be specified only for the cluster edition. Note: 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.
SA	Sybase database administrator. This attribute is required to connect to the ASE for shutdown. Type and dimension: string-scalar Default value: No default value
SApswd	Encrypted password for Sybase database administrator. This password is required to connect to the ASE for shutdown. Type and dimension: string-scalar Default value: No default value See “Encrypting passwords for Sybase” on page 157. Note: You need not specify a value for this attribute if the SA user does not require a password.

Table F-11 Required attributes (*continued*)

Required Attributes	Definition
Server	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. Type and dimension: string-scalar Default value: No default value
Version	Version of Sybase ASE. Type and dimension: string-scalar Default value: No default value Note: After 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. Note: For Sybase cluster edition, after 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.

Table F-12 lists the optional attributes.

Table F-12 Optional attributes

Optional Attributes	Definition
DetailMonitor	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. Note: This attribute has been removed from the agent version 6.1.0 onwards. Default is 0. Type and dimension: int-scalar Note: The DetailMonitor attribute is deprecated in SF Sybase CE 7.1. Instead, LevelTwoMonitorFreq attribute at the Sybase resource type level may be used.

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
LevelTwoMonitorFreq	Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. This is a resource type level attribute. You can override the value of this attribute at the resource level. The default value of LevelTwoMonitorFreq attribute is 0 (zero). If required, the value of LevelTwoMonitorFreq attribute can be overridden at resource level.
User	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. Type and dimension: string-scalar Default value: No default value
UPword	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. See "Encrypting passwords for Sybase" on page 157. intercType and dimension: string-scalar Default value: No default value
Db	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. Type and dimension: string-scalar Default value: No default value
Table	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. Type and dimension: string-scalar Default value: No default value

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
Monscript	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. Type and dimension: string-scalar Default value: No default value 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.
Run_ServerFile	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. Type and dimension: string-scalar Default value: No default value
IMF	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. This attribute includes the following keys: ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 3

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
IMF (cont.)	■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 5 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the sybase_imf_register 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. Default: 3 Type and dimension: Integer-association.
interfaces_File	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. Type and dimension: string-scalar Default value: No default value For example: /sybase/my_interfaces_file 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.

Table F-12 Optional attributes (*continued*)

Optional Attributes	Definition
DelayAfterOnline	Specifies the number of seconds that elapse after the Online entry point is complete and before the next monitor cycle is invoked. Type and dimension: integer-scalar Default value: 10
DelayAfterOffline	Specifies the number of seconds that elapse after the Offline entry point is complete and before the next monitor cycle is invoked. Type and dimension: integer-scalar Default value: 2
ShutdownWaitLimit	Maximum number of seconds for which the agent waits for the Sybase instance to stop after issuing the <code>shutdown with wait</code> command, and before attempting to issue the <code>kill -15 <data server-pid></code> command, if required. Type and dimension: integer-scalar Default value: 60
Quorum_dev	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. Note: 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. Type and dimension: String-scalar Default value: No default value
Run_ServerFile	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. Type and dimension: String-scalar Default value: No default value

Index

A

- adding
 - users 32
- agent for SQL server
 - attribute definitions 283
 - resource type 282
 - type definition 282
- agents
 - about 265
 - CFSMount 275
 - CVMCluster 267
 - CVMVolDg 272
 - CVMVxconfigd 269
 - of VCS 266
- attributes
 - about agent attributes 265
 - CFSMount agent 276
 - CVMCluster agent 267
 - CVMVolDg agent 267, 273
 - CVMVxconfigd agent 270

B

- backup boot disk group 144–145
 - rejoining 144

C

- cables
 - cross-over Ethernet 172
- CFSMount agent 275
 - attributes 276
 - entry points 276
 - sample configuration 278–279
 - type definition 278
- CFSTypes.cf 278
- cluster
 - removing a node from 191
 - verifying operation 78
- commands
 - hastatus 78
 - hasys 79

- commands (*continued*)
 - litstat 74
 - vxdisksetup (initializing disks) 38
- configuration
 - required information
 - for SF Sybase CE, 210
- configuring VCS
 - adding users 32
 - event notification 33, 35
 - global clusters 36
- CVM
 - CVMTypes.cf file 268
 - upgrading protocol version 147
- CVMCluster agent 267
 - attributes 267
 - entry points 267
 - sample configuration 269
 - type definition 268
- CVMTypes.cf
 - definition, CVMCluster agent 268
 - definition, CVMVolDg agent 274
 - definition, CVMVxconfigd agent 271
- CVMVolDg agent 272
 - attributes 273
 - entry points 272
 - sample configuration 275
 - type definition 274
- CVMVxconfigd agent 269
 - attributes 270
 - CVMTypes.cf 271
 - entry points 270
 - sample configuration 272
 - type definition 271

D

- detail monitoring
 - disabling 159
 - enabling 158
- disks
 - adding and initializing 38
 - testing with vxfsentsthdw 41

disks (*continued*)
 verifying node access 42

E

Ethernet controllers 172

G

GAB

 port memberships 76

GAB ports 77

global clusters

 configuration 36

H

hastatus -summary command 78

hasys -display command 79

hubs

 independent 172

I

I/O fencing

 checking disks 41

 shared storage 41

Install Bundles

 integration options 93

installation

 preparation 210

installation script options 201

installation worksheets 210

installer

 upgrading SF Sybase CE 100

installing

 language packages 82

L

Live Upgrade

 administering Solaris 10 boot environments 136

 completing Solaris 10 upgrade 134

 preparing 126

 reverting to primary boot environment 136

 Solaris 10 systems 127

 supported upgrade paths 126

 Switching boot environment for Solaris

 SPARC 137

 upgrading Solaris 10 on alternate boot disk 128

 upgrading Solaris 10 using the installer 133

 verifying Solaris 10 upgrade 135

LLT

 verifying 74

lltstat command 74

M

monitoring

 basic 281

 detail 281

N

nodes

 adding Sybase ASE CE nodes

 configuring GAB 175

 configuring LLT 175

 configuring VXFEN 175

 starting Volume Manager 174

 preparing Sybase ASE CE nodes

 about 186

 configuring CVM 180

 creating Sybase user and groups 186

 preparing Sybase resource mount

 points 186

 removing a node from a cluster

 tasks 190

 removing nodes

 GAB configuration 193

 LLT configuration 193

 modifying VCS configuration 194

O

options

 SF Sybase CE configuration 18

P

PATH variable

 VCS commands 74

ports

 GAB 77

preparing

 Live Upgrade 126

R

rejoining

 backup boot disk group 144

removing a node from a cluster

 editing VCS configuration files 192

 procedure 191

- removing a node from a cluster *(continued)*
 - tasks 190
- rolling upgrade
 - about 121

S

- sample configuration files. *See* main.cf samples
- SF Sybase CE configuration
 - of components 19
- SF Sybase CE configuration
 - about 18
 - options 18
 - preparation
 - worksheets 210
 - required information 210
- SF Sybase CE installation
 - preparation
 - worksheets 210
 - verifying
 - cluster operations 74
 - GAB operations 74
 - LLT operations 74
- SF Sybase CE upgrade
 - post-upgrade tasks
 - upgrading CVM protocol version 147
 - upgrading disk group version 147
 - preparation 97, 123
 - restoring configuration files 101
 - stopping cluster resources 97, 123
 - stopping Sybase ASE CE 97
 - using installer 100
- simultaneous install or upgrade 93
- SMTP email notification 33
- SNMP trap notification 35
- SQL server agent attributes
 - Db 286
 - DelayAfterOffline 289
 - DelayAfterOnline 289
 - DetailMonitor 285–286
 - home 283
 - monscript 287
 - owner 284
 - Quorum_dev 284
 - Run_ServerFile 287
 - SA 284
 - SApswd 284
 - server 285
 - ShutdownWaitLimit 289
 - table 286

- SQL server agent attributes *(continued)*
 - UPword 286
 - user 286
 - version 285

- Storage Foundation for Sybase ASE CE configuration
 - verifying 76

- supported upgrade paths
 - Live Upgrade 126

- Sybase
 - language settings 155
 - setting up for detail monitoring 155
- Sybase agent
 - detail monitoring 158
 - monitoring options 281
- system state attribute value 78

T

- tunables file
 - about setting parameters 215
 - parameter definitions 220
 - preparing 219
 - setting for configuration 216
 - setting for installation 216
 - setting for upgrade 216
 - setting parameters 220
 - setting with no other operations 217
 - setting with un-integrated response file 218
- type definition 282

U

- unsuccessful upgrade 145
- upgrade
 - stopping Sybase ASE CE 123
- using Live Upgrade 125

V

- VCS
 - command directory path variable 74
- VVR
 - global cluster overview 198
- vxdisksetup command 38

W

- worksheets
 - for SF Sybase CE 210