

Veritas™ Cluster Server Implementation Guide for Oracle

Windows Server 2003

5.1

Veritas Cluster Server Implementation Guide for Oracle

Copyright © 2008 Symantec Corporation. All rights reserved.

Veritas Cluster Server

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

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

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

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Windows is a registered trademark of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation.

Licensing and registration

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation and Upgrade Guide* for license installation instructions.

Technical support

For technical assistance, visit <http://www.symantec.com/business/support/index.jsp> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1	Introducing the VCS agents for Oracle and NetApp	
	About the VCS agents for Oracle and NetApp	10
	About the VCS agent for Oracle	11
	Oracle agent	11
	Netsnr agent	12
	Info entry point	12
	Action entry point	13
	About VCS hardware replication agents for NetApp	14
	NetApp Filer agent	14
	NetApp SnapDrive agent	15
	NetApp SnapMirror agent	16
	How the agent makes Oracle highly available	17
	Typical Oracle configuration in a VCS cluster	17
Chapter 2	Installing the VCS agent for Oracle	
	About installing the VCS agent for Oracle	20
	Before installing the agent	20
	Installing the agent	20
	Configuring the cluster	21
	Configuring Web console	33
	Configuring notification	34
	Configuring Wide-Area Connector process for global clusters	38
Chapter 3	Installing and configuring Oracle	
	About installing and configuring Oracle	40
	Before installing and configuring Oracle	40
	Privilege requirements	40
	Configuring Microsoft iSCSI Initiator	41
	Managing shared storage	42
	If you use NetApp filers	42
	If you use Windows Logical Disk Manager	45
	Creating a temporary service group	47
	Installing and configuring Oracle on the first node	48
	Bringing the temporary service group online	48
	Installing Oracle on the first node	48

- Creating the Oracle database on shared disks 50
- Creating the listener service 51
- Configuring the listeners to use the virtual IP address 52
- Associating the database with the listener 52
- Configuring the Oracle database 53
- Configuring the Oracle and listener services 58
- Disconnecting virtual disks from the first node 59
- Releasing disks (if you use Windows LDM) 59
- Installing and configuring Oracle on failover nodes 60
 - Installing Oracle on failover nodes 60
 - Configuring Oracle on failover nodes 60

Chapter 4 Configuring the Oracle service group

- About configuring the Oracle service group 64
- Deleting the temporary service group 64
- Configuring the Oracle service group using the wizard 65
 - Before configuring the Oracle service group 65
 - Configuration instructions 66
 - Configuring dependent services 72
- Verifying the Oracle service group configuration 73
 - Bringing the service group online 73
 - Taking the service group offline 73
 - Switching the service group 74
 - Disabling the service group 74

Chapter 5 Administering the Oracle service group

- About administering the Oracle service group 78
- Modifying an Oracle service group 78
 - Before modifying the Oracle service group 78
 - Instructions 78
- Deleting an Oracle service group 79
- Configuring detail monitoring 80
 - Enabling detail monitoring 80
 - Disabling detail monitoring 82

Chapter 6	Troubleshooting the agents	
	About troubleshooting VCS agents	84
	Troubleshooting common problems	84
	VCS logging	85
	VCS Cluster Configuration Wizard logs	86
	VCWsilent logs	87
	Agent error messages and descriptions	88
	Oracle agent	88
	Netlsnr agent	90
	NetApp agents	91
Appendix A	Resource type definitions	
	About resource type definitions	94
	Oracle agent	94
	Resource type definition	94
	Attribute definitions	95
	Netlsnr agent	97
	Resource type definition	97
	Attribute definitions	98
	NetAppFiler agent	98
	Resource type definition	98
	Attribute definitions	99
	NetApp SnapDrive agent	100
	Resource type definition	100
	Attribute definitions	100
	NetApp SnapMirror agent	101
	Resource type definition	101
	Attribute definitions	101
Appendix B	Sample configuration	
	About sample configurations	104
	Resource dependency graph	104
	Sample configuration file	105

Appendix C	Oracle Enterprise Manager 10g Grid Control with VCS	
	About Oracle Enterprise Manager 10g Grid Control with VCS	110
	Supported software	110
	Installing the VCS agent for Oracle	110
	Installing Oracle Enterprise Manager server	111
	Installing Oracle	111
	Installing the Oracle management agent	111
	Creating and configuring Oracle database and listener on the first node	112
	Configuring Oracle on failover nodes	113
	Configuring the Oracle service group	113
	Configuring a virtual network name	114
	Configuring the listener for virtual network name	114
	Configuring the database service for virtual network name	115
	Configuring an additional Oracle management agent	117
	Adding the database to the OEM 10g Grid Control	118
	Making the additional agent highly available	119
	Sample Configuration	120
Index		123

Introducing the VCS agents for Oracle and NetApp

This chapter contains the following topics:

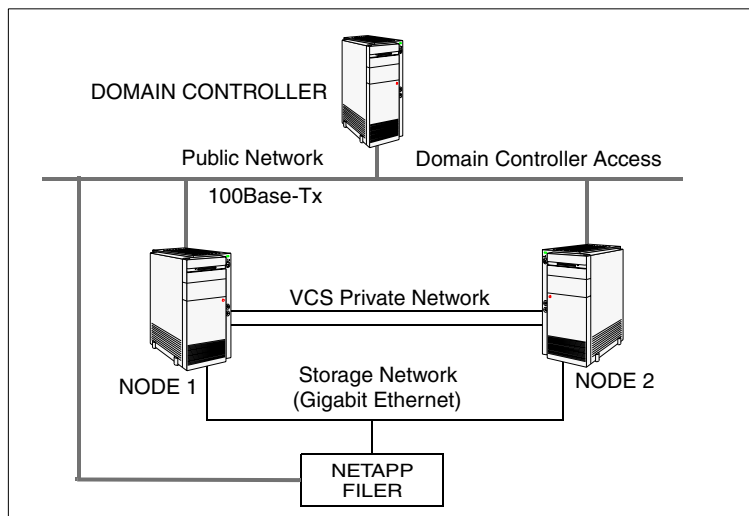
- [“About the VCS agents for Oracle and NetApp”](#) on page 10
- [“About the VCS agent for Oracle”](#) on page 11
- [“About VCS hardware replication agents for NetApp”](#) on page 14
- [“How the agent makes Oracle highly available”](#) on page 17
- [“Typical Oracle configuration in a VCS cluster”](#) on page 17

About the VCS agents for Oracle and NetApp

The Veritas Cluster Server (VCS) agent for Oracle provides high availability to Oracle in a VCS environment. The VCS agent for Network Appliance (NetApp) SnapMirror enables configuring NetApp filers in a VCS environment. Both the agents work together to provide high availability for Oracle, in environments that use Network Appliance filers for shared storage.

In a typical configuration, the agents are installed on each node in the cluster. The nodes are connected to the NetApp filers through a dedicated (private) storage network. VCS nodes are physically attached to the Network Appliance filer via an ethernet cable supporting iSCSI or Fibre Channel (FC) as the transport protocol.

Figure 1-1 Typical VCS configuration in a NetApp storage environment



This chapter provides an overview of the agents. For more information about the agents, including their resource type definitions and attribute definitions, see [Appendix A, “Resource type definitions”](#) on page 93.

About the VCS agent for Oracle

The VCS Database agent for Oracle monitors the Oracle database and listener services, brings them online, and takes them offline.

The package contains two VCS agents:

- The Oracle agent, which monitors the Oracle database services.
- The Netlsnr agent, which monitors the listener services.

Both agents work together to provide high availability for Oracle.

Oracle agent

The Oracle agent monitors the database processes. The agent brings the Oracle database services online, monitors their status, and takes them offline. Specific agent operations include:

- **Online**—Starts the Oracle service.
- **Offline**—Stops the Oracle service. The operation also stops all the services that are dependent on the Oracle service.
- **Monitor**—Verifies the status of Oracle services. The agent provides two levels of application monitoring: basic and detail.
 - **Basic monitoring**: The agent queries the Service Control Manager (SCM) to verify that Oracle services are continuously active.
 - **Detail monitoring**: The agent runs a SQL script to verify the status of the database. Detail monitoring is optional. It updates a temporary table in the Oracle database to verify the availability of the database instance. See “[Configuring detail monitoring](#)” on page 80 for instructions on configuring detail monitoring for an Oracle database.
- **Clean**—Forcibly stops the Oracle database.
- **Info**—Provides static and dynamic information about the state of the database. See “[Info entry point](#)” on page 12 for more information.
- **Action**—Performs predefined actions on a resource. See “[Action entry point](#)” on page 13 for more information.

Netlsnr agent

The Netlsnr agent monitors the listener process. The agent controls the listener for the database. The listener is a server process that listens to incoming client connection requests and manages traffic to the database.

The Netlsnr agent brings the listener services online, monitors their status, and takes them offline. Specific agent operations include:

- Online—Starts the listener service.
- Offline—Stops the listener service.
- Monitor—Queries the SCM for the status of the listener service.
- Clean—Forcibly stops the listener service.

Info entry point

The Info entry point provides static and dynamic information about the state of the database.

The entry point retrieves the following static information:

Version	InstanceNo	InstanceName
HostName	StartupTime	Parallel
Thread	InstanceRole	

The entry point retrieves the following dynamic information:

InstanceStatus	Logins	DatabaseStatus
ShutdownPending		

You can add additional attributes by adding `sql` statements to the file `%VCS_HOME%\bin\Oracle\scripts\info\db_info.sql`.

For example:

```
select 'static:HostName:'||host_name from v$instance;  
select 'dynamic:ShutdownPending:'||shutdown_pending from  
v$instance;
```

The format of the selected record must be as follows:

```
attribute_type:userkey_name:userkey_value
```

The variable `attribute_type` can take the value static and/or dynamic.

Action entry point

The Action entry point enables you to perform predefined actions on a resource. To perform an action on a resource, type the following command:

```
C:\> hares -action <res> <token> [-actionargs <arg1> ...] [-sys
<system>] [-clus <cluster>]
```

Table 1-1 includes the actions that the agent supports.

Table 1-1 Actions supported by Oracle agent

Token for Action	Description
DBRestrict	Changes the database session to enable the RESTRICTED mode.
DBUndoRestrict	Changes the database session to disable the RESTRICTED mode.
DBSuspend	Suspends a database.
DBResume	Resumes a suspended database.
DBQuiesce	Places the database into a quiesced state. Note: You must set the <i>resource_manager_plan</i> parameter and restart the database before quiescing a database.
DBUnQuiesce	Restores the database from a quiesced state to a normal operation state.

You can also add custom actions for the agent. For further information, refer to the *Veritas Cluster Server Agent Developer's Guide*.

About VCS hardware replication agents for NetApp

The VCS enterprise agent for Network Appliance provides failover support and recovery in environments employing Network Appliance filers for storage and SnapMirror for replication. The agent monitors and manages the state of replicated filer devices and ensures that at a time only one system has safe and exclusive access to the configured devices.

The agent can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments set up using the VCS Global Cluster Option. The package contains three agents; each agent is described in subsequent sections:

- NetAppFiler agent—Monitors the state of the filer.
- NetAppSnapDrive agent—Connects and disconnects virtual disks (LUNs) using the iSCSI or the FC protocol.
- NetAppSnapMirror agent—Determines the role of the filer volumes with respect to replication and promotes a read-only snapmirrored volume to a read-write source volume during a wide-area failover.

NetApp Filer agent

The NetApp Filer agent monitors the state of the filer device. The agent is represented by the NetAppFiler resource type in VCS. NetAppFiler resources are persistent, meaning that they are not brought online or taken offline.

Agent operation

- Monitor—Verifies the state of the filer attached to the host by sending an ICMP ping command to the filer. If the filer does not respond, the agent reports the state of the filer as faulted.

NetApp SnapDrive agent

The NetApp SnapDrive agent monitors, connects, and disconnects filer volumes. The agent can be configured to use the iSCSI or the FC protocol.

Agent operations

- **Online**—Connects a virtual disk (LUN) using an iSCSI or FC Initiator. The agent presents the LUN as a locally-attached drive to the host. The agent also removes LUN-host mappings made before the online operation.
- **Offline**—Disconnects the virtual disk (LUN) from the host.
- **Monitor**—Verifies that the specified virtual disk (LUN) is connected to the host.
- **Open**—Verifies that there is connectivity to the filer. It also checks that the VCS Helper service is running with the same privileges as the SnapDrive service.
- **Clean**—Attempts to forcibly disconnect a virtual disk (LUN).

NetApp SnapMirror agent

The NetApp SnapMirror agent monitors the replication state of filer devices. When a failover occurs, the agent reverses the direction of replication.

Agent operations

- **Online**—If the state of the local filer device is `SOURCE`, the agent creates a lock file to indicate that the resource can come online. This effectively makes the devices writable for the application.
If the state of the local filer is `SNAPMIRRORED`, the agent attempts to reverse the direction of replication by changing the state of the local filer to `SOURCE` and that of the original source to `SNAPMIRRORED`.
If the original source filer is down, the agent performs a mirror breakoff to enable local write access, if the filer is not already broken off.
If the original source returns to life, you must re-synchronize the data manually. The Online entry point touches a lock file if Read Write access is enabled successfully.
- **Offline**—Removes the lock file. The agent does not perform any filer operations because an offline entry point does not necessarily indicate an intention to give up the devices.
- **Monitor**—Verifies the lock file exists. If the lock file exists, the monitor entry point reports the status of the resource as online. If the lock file does not exist, the monitor entry point reports the status of the resource as offline.
- **Open**—Removes the lock file, thereby preventing potential concurrency violation if the group fails over to another node.

Note: The agent does not remove the lock file if the agent was started after an `hastop -force` command.

- **Clean**—Removes the lock file. No filer operations are performed since offlining this resource is no indication of a pending role swap.

How the agent makes Oracle highly available

In the basic monitoring mode, the database agent for Oracle detects a failure if the database or a listener service is not running. In the detail monitoring mode, the agent detects failure if a preconfigured SQL script cannot execute or access the database or if the listener service is not running.

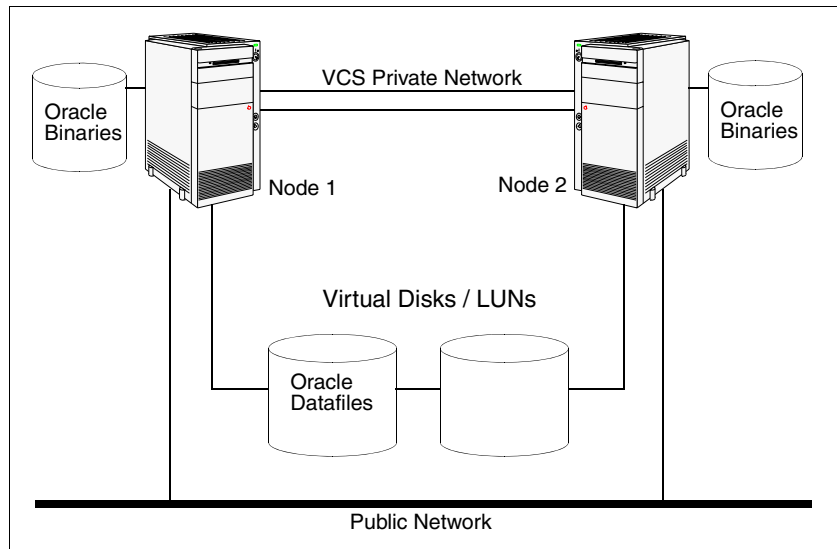
The detail monitoring interval can be defined using the *IndepthMonitorCount* attribute. See [Appendix A, “Resource type definitions”](#) on page 93 for more information.

When the agent detects a failure, the Oracle service group switches to the next available system listed in the service group’s SystemList attribute. VCS starts the configured Oracle and listener services on the new system, thus ensuring high availability for the Oracle services and data.

Typical Oracle configuration in a VCS cluster

In a typical configuration, VCS is configured in a two-node cluster (System 1 and System 2). Oracle is installed on both systems, and the data to be used by Oracle is installed on shared storage. The VCS Database agent for Oracle is installed on both systems.

Figure 1-2 Typical Oracle configuration in a VCS cluster



Installing the VCS agent for Oracle

This chapter contains the following topics:

- [“About installing the VCS agent for Oracle”](#) on page 20
- [“Before installing the agent”](#) on page 20
- [“Installing the agent”](#) on page 20
- [“Configuring the cluster”](#) on page 21

About installing the VCS agent for Oracle

This chapter describes how to install the VCS database agent for Oracle and then configure the cluster. Follow these instructions if you did not install the agent while installing VCS.

Before installing the agent

- Verify that you have VCS installed on all cluster nodes. Refer to the *Veritas Cluster Server Installation and Upgrade Guide* for instructions.
- Verify that you have Local Administrator privileges on the node where you are installing the agent.

Installing the agent

Complete these steps to install the VCS database agent for Oracle on a cluster node. Repeat these steps on all systems on which you wish to configure Oracle.

To install the agent

- 1 Start the installer for VCS. In the Add/Remove Programs window, click **VCS 5.1 for Windows** and click **Change**.
- 2 In the VCS for Windows dialog box, choose the **Add or Remove** option and click **Next**.
- 3 In the VCS product options dialog box, click **Next**.
- 4 Check the **Veritas Cluster Server Database Agent for Oracle** check box and click **Next**.
The disk space required for the installation is displayed toward the bottom of the screen. When you add or remove an option, the total space changes.
- 5 The installer validates the system for prerequisites. After the system is accepted, click **Next**.
If a system is rejected, the Comments column displays the cause for rejecting the system. Highlight the system to view a detailed information about the failure in the Details box. Resolve the error, highlight the system from the list, and click **Validate Again**.
- 6 Review the summary of your selections and click **Update** to start the installation. The installer displays the status of installation.
- 7 After the installation is complete, review the installation report and click **Next**.

8 Click **Finish**.

After the agent is successfully installed, proceed to configuring the cluster. See “[Configuring the cluster](#)” on page 21.

Configuring the cluster

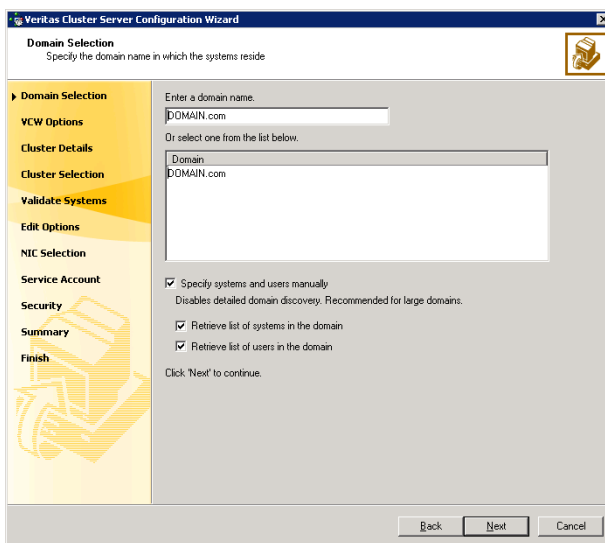
After installing VCS using the product installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console, notification, and global clusters.

- If you plan to set up a disaster recovery environment, configure the wide-area connector process for global clusters.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper Service user context to access the network. This account does not require domain admin privileges.
- In case of NetApp, the user account for the VCS Helper service must have administrative privileges on the NetApp filer.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



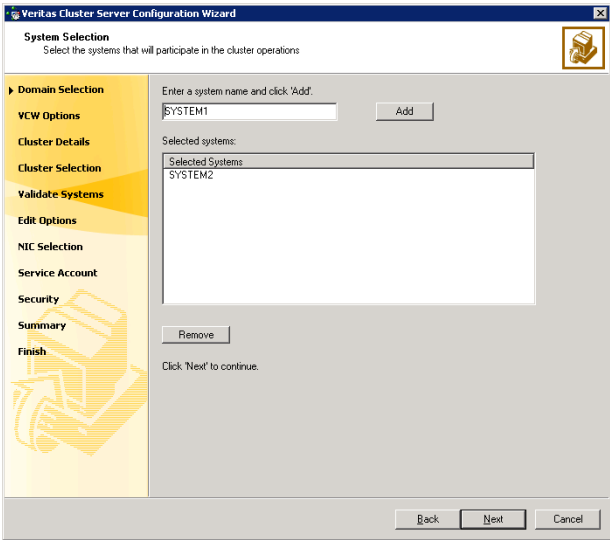
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 8](#) on page 24.

To specify systems and user names manually (recommended for large domains):

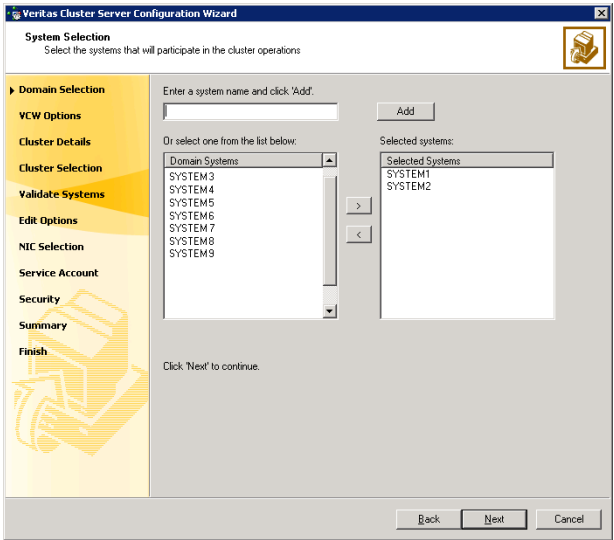
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 23.
Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 24.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier.

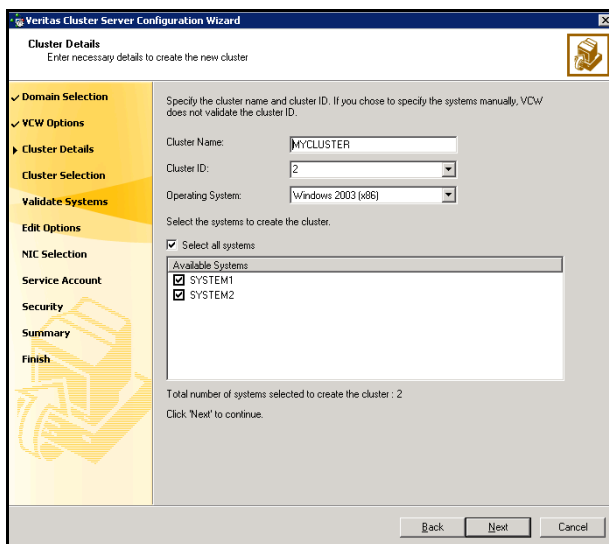
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) on page 22 or if you share a private network between more than one domain, make sure that the cluster ID is unique.

Operating System From the drop-down list, select the operating system that the systems are running.

Available Systems Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat. Check the **Select all systems** check box to select all the systems simultaneously.

10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

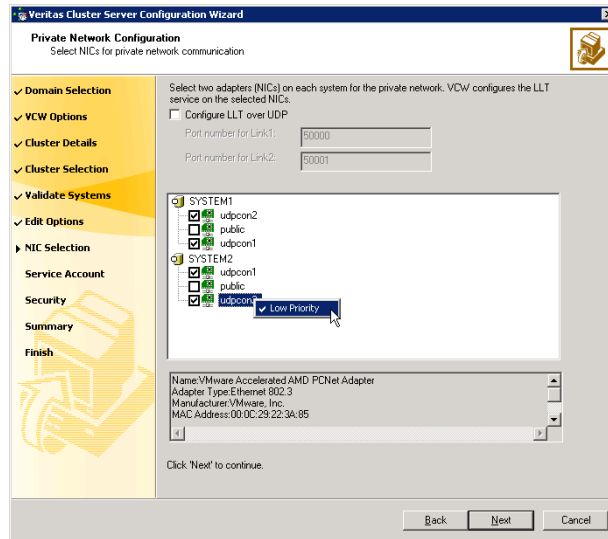
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 24, proceed to the next step. Otherwise, proceed to [step 12](#) on page 28.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:


- To configure the VCS private network over Ethernet



- Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard will configure the LLT service (over Ethernet) on the selected network adapters.

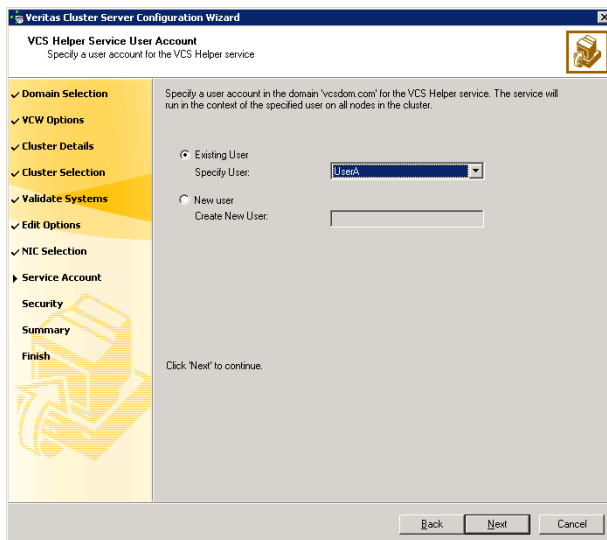
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper Service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper Service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 22, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

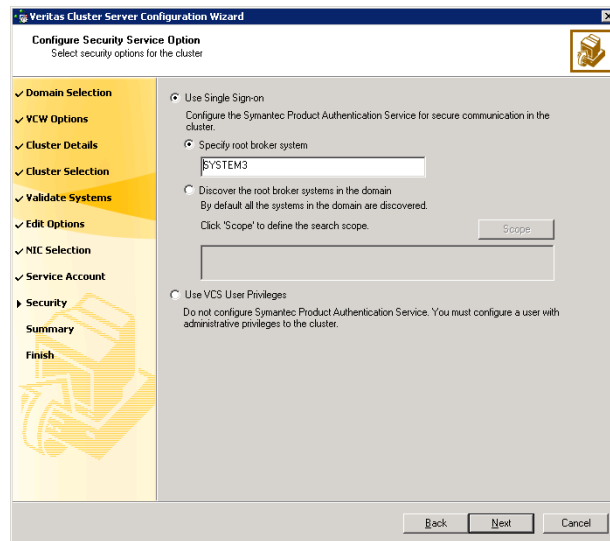
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers.

Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 2-1](#) contains some more examples of search criteria.

Table 2-1 Search criteria examples

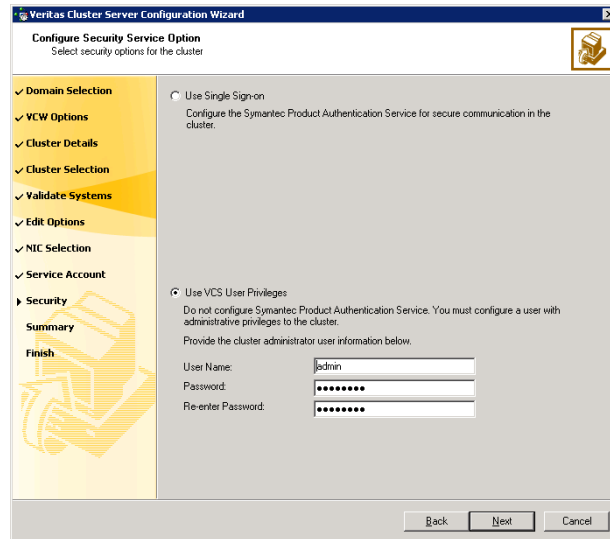
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server 2003 Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match *ALL* the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.

It is recommended that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

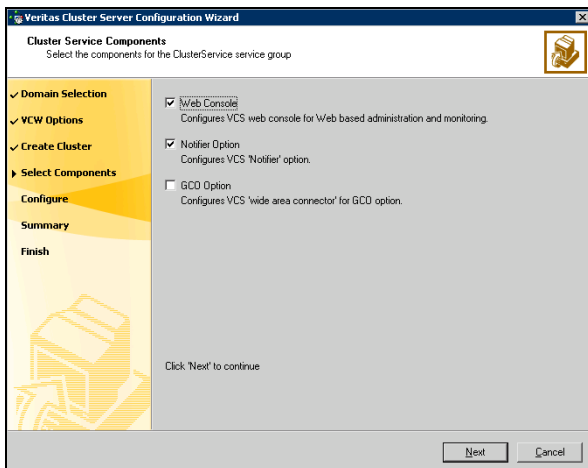
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add the system to the cluster.

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource. Although many of the operations through Cluster Manager (Java Console) can also take place through Cluster Management Console (Single Cluster Mode), this guide only provides instructions from the Java Console when VCS GUI procedures are required.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.



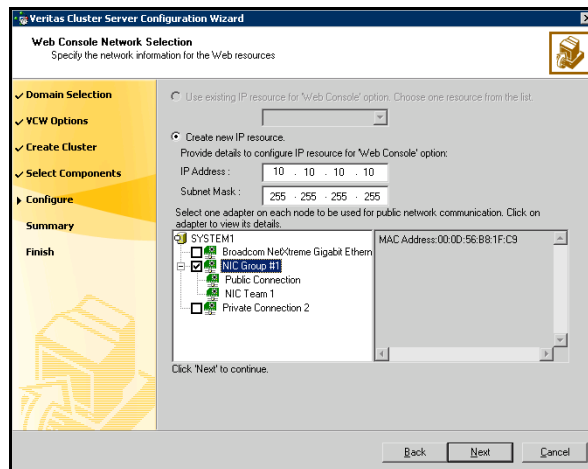
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 33.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 34.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

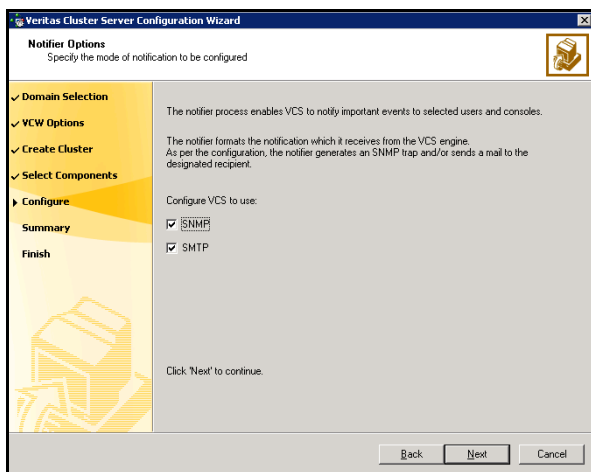
- 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 34. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.



You can configure VCS to generate SNMP (V2) traps on a designated server and/or send emails to designated recipients in response to certain events.

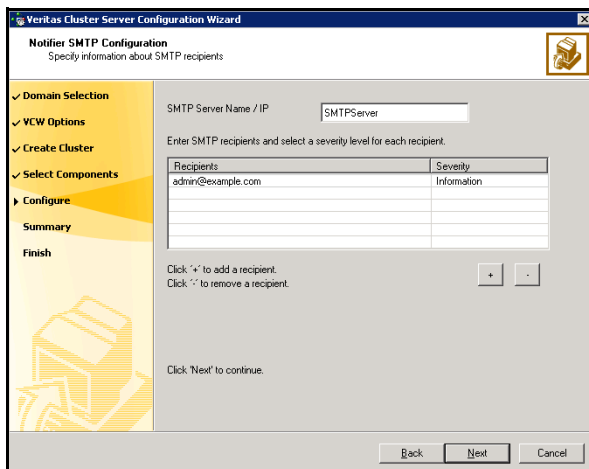
- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

The screenshot shows the 'Notifier SNMP Configuration' window. On the left is a navigation pane with options: Domain Selection, VCW Options, Create Cluster, Select Components, Configure (selected), Summary, and Finish. The main area is titled 'Specify information about SNMP console' and contains a table for console configuration. Below the table are instructions on how to add or remove consoles, a text input for the 'SNMP Trap Port' (set to 162), a note about MIB 2.0 compliance, and a 'Next' button. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

SNMP Console	Severity
snmposev	Information
snmpsev1	SevereError


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



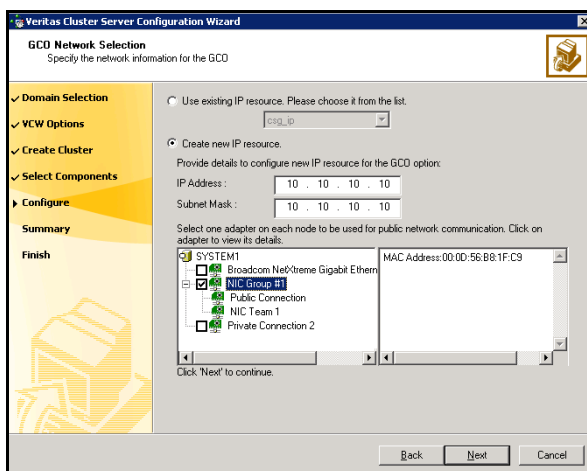
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

Configuring Wide-Area Connector process for global clusters

Configure the wide-area connector process only if you are configuring a disaster recovery environment.

To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
 - 3 Click **Finish** to exit the wizard.

After configuring the cluster, proceed to installing and configuring Oracle in the cluster.

See [Chapter 3, “Installing and configuring Oracle”](#) on page 39.

Installing and configuring Oracle

This chapter contains the following topics:

- [“About installing and configuring Oracle”](#) on page 40
- [“Before installing and configuring Oracle”](#) on page 40
- [“Configuring Microsoft iSCSI Initiator”](#) on page 41
- [“Managing shared storage”](#) on page 42
- [“Creating a temporary service group”](#) on page 47
- [“Installing and configuring Oracle on the first node”](#) on page 48
- [“Installing and configuring Oracle on failover nodes”](#) on page 60

About installing and configuring Oracle

This chapter describes how to install and configure Oracle and its components in a VCS cluster.

The strategy for installing Oracle in a VCS cluster is aimed at ensuring that Oracle installation on all the nodes in the cluster is uniform. For specific instructions on how to install Oracle, see the Oracle documentation.

Before installing and configuring Oracle

- Verify all nodes have adequate resources to run Oracle and VCS.
- Verify that VCS is installed on all nodes in the cluster. Refer to the *Veritas Cluster Server Installation and Upgrade Guide* for installation steps.
- Verify you have run the VCS Cluster Configuration Wizard (VCW) to configure VCS on all nodes.
See “[Configuring the cluster](#)” on page 21.
- If using iSCSI, verify that the Microsoft iSCSI Initiator is configured to establish a persistent connection between the Network Appliance filer and the cluster nodes. See the Microsoft documentation for instructions.
- If using FC, verify that you install the NetApp FCP Attach Kit or Windows Host Utilities on all the cluster nodes. Refer to the NetApp documentation for more information.
- Create the required volumes or LUNs on the shared cluster disk groups before installing Oracle.
See “[Managing shared storage](#)” on page 42 for instructions.

Privilege requirements

- Log on as domain administrator to install Oracle binaries.
- You must be a member of the Local Administrators group on all nodes on which Oracle will be installed.
- You must have write permissions for the Active Directory objects corresponding to these nodes.
- You must have write permissions on the DNS server to perform DNS updates.
- Domain Administrator or Account Operator privileges: if you plan to create a new user account for VCS Helper Service, you must have Domain Administrator privileges or belong to the Account Operators group. If you

plan to use an existing user account context for the VCS Helper Service, you must know the password for the user account.

- The user account for the VCS Helper Service must be a domain user account and must be a member of the administrators group on the filers.

Configuring Microsoft iSCSI Initiator

The Microsoft iSCSI initiator enables communication between Windows systems and Network Appliance filers. The initiator uses the iSCSI protocol to present the filer volume as a local block device to the system.

To configure Microsoft iSCSI initiator

- 1 Make sure the Microsoft iSCSI Initiator software is installed on all cluster nodes. Refer to Microsoft documentation for further information.
- 2 Start the Microsoft iSCSI initiator. Double-click the Microsoft iSCSI Initiator icon from the desktop.
- 3 Click the **Target Portals** tab, if not already selected.
- 4 Click **Add...**
- 5 In the Add Target Portals dialog box, specify the IP address or the DNS name for the Network Appliance Filer and click **OK**.
- 6 Click the **Available Targets** tab and click **Log On...**
- 7 In the Log On to Target dialog box, verify the target portal name and select the **Automatically restore this connection when the system reboots** check box.
- 8 Click **OK**.
- 9 Click the **Persistent Target** tab to verify that the newly added target portal is listed under the **Select a target** box.
- 10 Click **OK**.

Managing shared storage

Create volumes or LUNs on the shared disks for the following:

- database files
- control files
- log files, including the redo logs, archive logs, `background_dump_dest` (bdump), `core_dump_dest` (cdump), and `user_dump_dest` (udump) files
- Flash Recovery Area and the Archiving location

The Oracle `background_dump_dest` (bdump), `core_dump_dest` (cdump), and `user_dump_dest` (udump) files must be on the same volume or LUN as that of the redo logs.

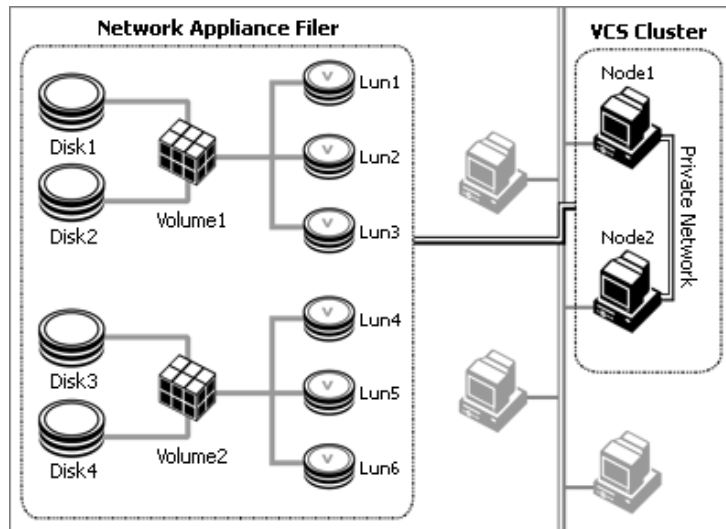
For example, you could create the following volumes or LUNs to manage your Oracle database and logs:

- `INST1_DATA_FILES`: contains the Oracle system data files (including the master, model, msdb, and tempdb databases)
- `INST1_REGREP_VOL`: contains the list of registry keys that must be replicated among cluster systems for the Oracle Service. Create a 100 MB (minimum recommended size) volume for this purpose
- `INST1_DB1_VOL`: contains the user database files
- `INST1_DB1_LOG`: contains the database log files (including redo logs, archive logs, `background_dump_dest` (bdump), `core_dump_dest` (cdump), and `user_dump_dest` (udump))

If you use NetApp filers

Network Appliance manages data by creating volumes on physical disks. These volumes can further be divided into Logical Unit Numbers (LUNs). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and Network Appliance SnapDrive installed. If you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that you install the NetApp FCP Attach Kit or Windows Host Utilities on all the cluster nodes.

Note: Symantec does not support volumes created using qtrees.

Figure 3-1 VCS cluster in a NetApp storage environment

The VCS agent for Oracle requires the following LUNs to be created on the Network Appliance filer. These LUNs must be accessible from all cluster nodes. Perform the following tasks to create LUNs on the Network Appliance filer and to make them accessible from cluster nodes:

- Create volumes on the Network Appliance filer.
- Share the volumes.
- Create LUNs on the shared volumes.

Refer to Network Appliance documentation for instructions on performing these tasks.

Connecting virtual disks to the cluster node

Once the virtual disks are created on the Network Appliance filer, they must be connected (if not connected already) to the cluster nodes using Network Appliance SnapDrive.

To connect virtual disks to the cluster node

- 1 Start the Computer Management MMC on the cluster node where you want to connect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.

- 3 Right-click **Disks** and then click **Connect Disk...** to launch the Connect Disk wizard.
- 4 Click **Next** on the Welcome page.
- 5 Specify the path of the virtual disk that you wish to connect to the cluster node and then click **Next**.
- 6 Select **Dedicated** as the Virtual Disk Type and then click **Next**.
- 7 Click **Assign a Drive Letter** and then choose a drive letter from the drop-down list.
- 8 On the Select Initiator panel, specify the initiator(s) for the virtual disk and then click **Next**.
- 9 On the igroup Management Type panel, choose the option that allows SnapDrive to perform igroup management automatically and then click **Next**.
- 10 Click **Finish** to begin connecting the specified virtual disk to the cluster node.

Disconnecting virtual disks from the cluster nodes

To disconnect virtual disks

- 1 Start the Computer Management MMC on the cluster node where you want to disconnect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Double-click **Disks** to see the LUNs that are connected to the node.
- 4 Right-click the LUN you want to disconnect and then click **Disconnect Disk...**
- 5 In the Disconnect Disk alert box, click **OK**.

If you use Windows Logical Disk Manager

If your configuration uses shared disks and volumes managed using Windows Logical Disk Manager (LDM), use the VCS DiskReservation (DiskRes) and Mount agents. Before configuring shared storage, review the resource types and attribute definitions of the Disk Reservation and Mount agents described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Perform the following tasks to create volumes and make them accessible from server farm nodes:

- Reserve disks. See “[Reserving disks](#)” on page 45.
- Create volumes. See “[Creating volumes](#)” on page 46.
- Mount volumes. See “[Mounting volumes](#)” on page 46.

About LDM support

The following restrictions apply in this release:

- Disk Reservation and Mount agents are supported on VCS only. These agents are not supported in an SFW storage environment.
- LDM support is available on Windows Server 2003 only.
- For using LDM, your storage devices must be configured to use SCSI-2 disk reservations. SCSI-3 is not supported.
- LDM support is not applicable for Disaster Recovery configurations. Currently only HA configurations are supported.

Reserving disks

Complete the following steps to reserve the disks on the node on which you are going to install Oracle.

To reserve the disks

- 1 To display all the disks, type the following on the command line:

```
C:\>haval -scsitest /l
```

You will see a table that lists all the disks that are visible from the current system. Make a note of the disk numbers (Disk# column in the table). You will need it in the next step.

- 2 To reserve a disk, type the following on the command line:

```
C:\>haval -scsitest /RES:<disk #>
```

For example, to reserve disk #4, type:

```
C:\>haval -scsitest /RES:4
```

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks while installing Oracle, and configuring the Oracle service group, on additional nodes in the cluster.

Creating volumes

Use the Windows Disk Management tool to verify that the disks are visible on the cluster nodes, and then create volumes on the reserved disks. After creating the required volumes on a node, rescan the disks on all the remaining nodes in the cluster.

Refer to Microsoft Windows documentation for more information about the Disk Management tool.

Mounting volumes

Use the Windows Disk Management tool to mount the volumes that you created earlier. After mounting the volumes on a cluster node, run the CHKDSK command and verify that there are no errors on the mounted volumes.

Make a note of the drive letters that you assign to the mounted volumes. Use the same drive letters while mounting these volumes on the remaining cluster nodes.

Refer to Microsoft Windows documentation for more information about the CHKDSK command and the Disk Management tool.

Releasing disks

To release a reserved disk, type the following on the command line:

```
C:\>havol -scsitest /REL:<disk #>
```

For example, to release disk 4, type:

```
C:\>havol -scsitest /REL:4
```

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks.

Unassigning a drive letter

While installing an application on multiple nodes, you must first unassign drive letters and release the disks from one node, and then reserve the disks, mount the volumes using the same drive letters and then install the application on the failover node.

Complete these steps to unassign the drive letters from a node.

To unassign drive letter

- 1 Log in as Administrator.

- 2 Open Disk Management.
`C:\> diskmgmt.msc`
- 3 Right-click the partition or logical drive and click **Change Drive Letter and Path**.
- 4 In the **Change Drive Letter and Paths** dialog box, click the drive letter and click **Remove**.

Note: You must run Disk Management on all systems each time you add a shared disk. This ensures each disk has a valid signature written to it, and that the device paths and symbolic links are updated.

Creating a temporary service group

Create a temporary service group with an IP resource. Assign a virtual address to the IP resource. All the configured databases and listeners must use the same virtual IP address.

To create a temporary service group

- 1 From the Cluster Manager (Java Console), click **Edit > Add > Service Group...**, and complete the following in the Add Service Group dialog box:
 - Type a name for the service group in the Service Group name field.
 - Click the nodes (Use Ctrl click to select multiple nodes) from the Available Systems box and then click the right arrow to move the systems to the Systems for Service Group box.
 - Click **OK**.
- 2 To add an IP resource to the temporary service group, right-click the service group name, click **Add Resource...**, and complete the following in the Add Resource dialog box:
 - Type a name for the IP resource in the Resource name field.
 - Select **IP** from the Resource Type drop-down list.
 - Specify appropriate values for the **Address**, **SubNetMask**, and **MACAddress** attributes. For the MACAddress attribute, select the **Per System** option, and then choose the current node name from the drop-down list. The MACAddress attribute is always local; it is different for each system. You must specify a MACAddress for each node in the service group.
 - Check the **Enabled** check box and then click **OK**.

Installing and configuring Oracle on the first node

Perform the following installation and configuration tasks on the first node in the cluster. The strategy for installing Oracle into a VCS cluster is aimed at ensuring that installations on all cluster systems are uniform. This involves installing the Oracle binaries locally on each system and installing the Oracle control files, data files, log files, bdump, cdump, udump files, the flash recovery area and the archive location, on shared storage.

Bringing the temporary service group online

Bring the temporary service group online on the node on which you installing Oracle first. The IP address specified for the IP resource in the service group is used by all the configured databases and listeners.

Complete the following steps.

To bring the temporary service group online

- 1 From the Cluster Manager (Java Console), right-click the temporary service group name, click **Online** and then select the current node from the list. This is the cluster node on which you are installing Oracle first.
- 2 Click **Yes** on the prompt that confirms whether you wish to bring the service group online.

Installing Oracle on the first node

Complete these steps to install Oracle. The installer screens and options may vary slightly for different versions of Oracle. Refer to the Oracle documentation for detailed installation steps.

While installing Oracle on the first node, note the following:

- Install the Oracle software program files on a local disk and the Oracle data files on shared storage. Install all log files, including redo logs, archive logs, bdump, cdump, and udump files on the same volume.
- Note the Oracle parameters while you install Oracle. For example, the database name, file locations, database credentials. You must use the same values for these parameters while installing Oracle on the remaining cluster nodes.
- The Oracle Home name and path should be the same on all the cluster nodes.
- Do not create the Oracle Starter database.

To install Oracle

- 1 Verify that the volumes or LUNs created for the Oracle database, control files, log files, bdump, cdump, udump files, flash recovery area, and archive location are mounted or connected on the node where you run the installer. See “[Managing shared storage](#)” on page 42.
- 2 Insert the Oracle software disc in a drive connected to the node and run the Setup program.
- 3 Complete the following instructions depending on the version of Oracle:

For Oracle 9i

- From the browser menu, click **Install/Deinstall Products**.
- Read the information in the Welcome panel and click **Next**.
- In the File Locations dialog box, type or select a name and a path on the local disk, for the Oracle Home, and click **Next**.
This will install the Oracle binaries on the local disk. Use the same Oracle Home name and path on all nodes.
- In the Available Products panel, select the **Oracle 9i Database option** and click **Next**.
- In the Installation Types panel, select the **Enterprise Edition** option and click **Next**.
- In the Database Configuration panel, select the **Customized** or **Software Only** option and click **Next**. Select the **Customized** option if you want to migrate earlier databases to Oracle 9i.
- If you chose the **Customized** option, then on the Database Identification panel, select or type a unique name and SID for the global database and click **Next**.

For Oracle 10g

- In the Available Products panel, select the **Oracle 10g Database** option and click **Next**.
- Choose the **Basic** or **Advanced** installation type.
For Basic installation:
 - In the **Oracle Home Location** box, type or select a name and a path on the local disk for the Oracle Home.
This installs the Oracle binaries on the local disk. Use the same Oracle Home name and path on all nodes.
 - Select **Enterprise Edition**.
 - Uncheck the **Create Starter Database** check box and click **Next**.For Advanced installation:

- Select the destination directory for the Oracle installation and click **Next**.
 - Select **Enterprise Edition** and click **Next**.
 - Select **Do not create a starter database**.
- 4 Follow the wizard instructions, select options as per your configuration, and complete the installation.

Creating the Oracle database on shared disks

Create an Oracle database on the shared disk using the Oracle Database Configuration Assistant. You need to create the database only once, from the node on which you installed Oracle first. You do not need to repeat these steps on the remaining cluster nodes.

To create a database on a shared disks

- 1 Start the Database Configuration Assistant (**Start > All Programs > Oracle-Oracle_Home > Configuration and Migration Tools > Database Configuration Assistant**). This launches the database installer.
- 2 In the Welcome panel, click **Next**.
- 3 In the Operation panel, select **Create a database** and click **Next**.
- 4 Complete the following instructions depending on the Oracle version.
For Oracle 9i
 - In the Database Templates panel, select **New Database** and click **Next**.
 - In the Database Identification panel, specify the **Global Database Name** and **SID** in the respective fields.
 - Install the initialization file on the shared virtual disk. In the Initialization Parameters dialog box, click the **File Locations** tab. In the **Initialization Parameter Filename** box, enter a path on the shared virtual disk for the PFILE. If you create an SPFILE, in the **Server Parameters Filename** field, enter a path on the shared virtual disk for the SPFILE.**For Oracle 10g**
 - In the Database Templates panel, select the desired database template and click **Next**.
 - In the Database Identification panel, specify the **Global Database Name** and **SID** in the respective fields.
 - Choose to create the database on the shared disk.

- 5 Follow the wizard instructions and select other options as per your desired database configuration. Install the control files, data files, log files, bdump, cdump, and udump files on the shared disk.
 - On the Recovery Configuration panel, specify the Flash Recovery Area and the Archiving location on the shared disk.
 - On the Initialization Parameters panel, click **All Initialization Parameters**, then click **Show All Parameters** on the All Initialization Parameters window, and then specify the path for the `background_dump_dest` (bdump), `core_dump_dest` (cdump), and `user_dump_dest` (udump) files.
 - In the Database Storage panel, specify the locations of control files, data files, and redo log files. In the left pane, expand the **Storage** folder.
 - Enter a path on the shared disk for each control file.
To do so, click **Controlfile** and click the **General** tab in the right pane. In the **File Directory** field, enter a path on the shared disk.
 - Enter a path on the shared disk for each data file.
To do so, expand the **Datafiles** folder in the left pane. Select a data file, and click the **General** tab in the right pane. In the **File Directory** field, enter a path on the shared disk.
 - Enter a path on the shared disk for each redo log file.
To do so, expand the **Redo Log Groups** folder in the left pane. Select a redo log group and click the **General** tab in the right pane. In the **File Directory** field, enter a path on the shared disk.
The path for the redo log files must be the same as that of the bdump, cdump, and udump files, specified earlier.
 - Click **Next**.
- 6 Select the **Create database** and the **Generate Database Creation Scripts** options and click **Finish**.
- 7 Follow the wizard instructions and complete the installation.

Creating the listener service

Create a listener service using the Net Configuration Assistant. Refer to the Oracle documentation for instructions.

Make a note of the listener name and port number you assign to the listener; you will need it in subsequent configuration tasks.

Configuring the listeners to use the virtual IP address

In the file listener.ora file, change the host name to the virtual IP address that you created. This is the same IP address that you specified for the IP resource in the temporary service group.

To configure listeners to use the virtual IP address

- 1 From the Oracle Home, open the file listener.ora.
Typically, the default location of the file listener.ora is
`<ORACLE_HOME>\NETWORK\ADMIN\`
- 2 Edit the entry in listener.ora as follows:
Replace
`HOST=machine_name`
with
`HOST=virtual_IP_address`
- 3 Save and close the listener.ora file.
- 4 Restart the listener service from the Services snap-in.

Associating the database with the listener

Complete these steps to associate all the databases with the listeners. Before making these changes make a backup copy of the SPFILE.

Prerequisite

Ensure that the initialization parameter file contains the following entries:

- SERVICE_NAMES (the name of the database service)
- INSTANCE_NAME (the name of the database instance)

These parameters are created during installation or database creation.

To associate the database with the listener

- 1 Configure a new attribute *listener_alias*. The listener_alias value should be unique; it should be different from the listener name. Run the following SQL commands.
 - Set the SID name.
Type the following on the command prompt:
`C:\>set ORACLE_SID=<SID NAME>`
 - Log on to the Oracle database as an administrator.
Type the following on the command prompt:
`C:\>sqlplus "/nolog"`

```
C:\>sqlplus "/as sysdba"
```

■ **Configure listener_alias.**

Type the following on the command prompt:

```
SQL> ALTER SYSTEM SET LOCAL_LISTENER='<listener_alias>'  
scope=spfile;
```

If you are using a pfile, add the following entry to the initialization parameter file (pfile):

```
LOCAL_LISTENER = <listener_alias>
```

- 2 Define the parameter listener_alias. If your Oracle configuration uses the file tnsnames.ora, edit the file as instructed below. The default location of tnsnames.ora is <ORACLE_HOME>\NETWORK\ADMIN.

Add the following to tnsnames.ora file:

```
<listener_alias>=  
(DESCRIPTION =  
(ADDRESS=(Protocol=TCP)(HOST=virtual_IP_address)(Port=port))  
)
```

Here, virtual IP address is the IP address you assigned to the IP resource in the temporary service group and port is the port number you specified while creating the listener. You can check the listener.ora file to find out the port number used. The default value is 1521.

- 3 Stop and restart the database.

The listener_alias parameter gets appended by the default domain name specified in the file sqlnet.ora.

Configuring the Oracle database

Configuring the database involves setting parameters in Oracle configuration files. It involves setting the same virtual IP address for all the databases and configuring the database to use the OS authentication model of Oracle.

Updating the tnsnames.ora files

If your configuration has a tnsnames.ora file, edit the file as instructed below. For all the databases that you wish to make highly available using VCS, change the respective host names entries for all the TCP protocol address databases to the virtual IP address. The virtual IP address is the one you specified for the IP resource of the temporary service group that you created earlier.

To update the tnsnames.ora file

- 1 From the Oracle Home, open the file tnsnames.ora.
Typically, the default location of the file tnsnames.ora is
<ORACLE_HOME>\NETWORK\ADMIN\.

- 2 Edit the entry in `tnsnames.ora` as follows:

Replace

```
HOST=machine_name
```

with

```
HOST=virtual_IP_address
```

- 3 Save and close the `tnsnames.ora` file.

If you use a PFILE (initSID.ora)

Configure the PFILE to use the virtual IP address for the service group and set the database to use the OS authentication model of Oracle.

- 1 Edit the `dispatchers` parameter by adding or editing the `HOST` parameter. If the `ADDRESS` tag does not appear, add it. The edited `dispatchers` parameter should look like:

```
dispatchers = '(ADDRESS=(Protocol=TCP)
(HOST=virtual_IP_address) (other entries))'
```

The variable *virtual_IP_address* represents the virtual IP address that you created. For example:

```
dispatchers = '(ADDRESS=(Protocol=TCP)
(HOST=10.100.101.102) (SERVICE=Data1XDB)'
```

- 2 Configure the attribute *remote_login_passwordfile*. Symantec recommends setting the attribute to `NONE` so that the database uses OS authentication and the configuration does not require a password file.

A sample entry appears as: `remote_login_passwordfile = 'NONE'`

If you use an SPFILE

- If you use an SPFILE on Oracle 9i, first modify the PFILE to point to the SPFILE as instructed in “[Modifying the PFILE \(Oracle 9i only\)](#)” on page 54, then update the SPFILE as instructed in “[Modifying the SPFILE \(Oracle 9i and 10g\)](#)” on page 55.
- If you use an SPFILE on Oracle 10g, proceed to “[Modifying the SPFILE \(Oracle 9i and 10g\)](#)” on page 55.

Modifying the PFILE (Oracle 9i only)

On Oracle 9i, edit the PFILE so that it points to the SPFILE. The PFILE must contain only the following entry:

```
SPFILE = spfile_location
```

The variable *spfile_location* represents the fully qualified path of the SPFILE on the shared virtual disk.

Modifying the SPFILE (Oracle 9i and 10g)

VCS requires a customized version of the SPFILE. Edit the SPFILE as instructed below.

Ensure that the Oracle database is mounted and is running. Make a backup copy of the SPFILE before you proceed.

To modify the SPFILE

- 1 Log on to the Oracle database as administrator and edit the dispatchers parameter by adding or editing the HOST parameter. Type the following commands on the command prompt:

```
C:\>sqlplus "/as sysdba"  
C:\>alter system set dispatchers =  
'(ADDRESS=(Protocol=TCP) (HOST=virtual_IP_address)  
(other entries))' scope=spfile;
```

The edited dispatchers parameter should look like:

```
dispatchers = '(ADDRESS=(Protocol=TCP) (HOST=virtual_IP_address)  
(other entries))'
```

The variable *virtual_IP_address* represents the virtual IP address that you specified for the IP resource in the temporary service group.

For example:

```
dispatchers = '(ADDRESS=(Protocol=TCP) (HOST=10.212.60.102)  
(SERVICE=Data1XDB)'
```

- 2 Configure the attribute *remote_login_passwordfile*.

Type the following command on the command prompt:

```
C:\>alter system set remote_login_passwordfile='NONE'  
scope=spfile;
```

Symantec recommends setting the attribute to NONE so that the database uses OS authentication and the configuration does not require a password file.

- 3 Restart the database service.

Moving the initialization file to shared disks (Oracle 10g)

For Oracle 10g, move the initialization files to the shared disk and configure the SPFILE to start a database instance.

- 1 Verify which of the following initialization files exist in your configuration. The default location of these files is `<ORACLE_HOME>\database.`

- SPFILESID.ora
- SPFILE.ora
- initSID.ora

The variable *SID* represents the database instance.

- 2 Stop the database service using the Services snap-in.
- 3 Move the initialization files from the local disk to the shared disk, preferably at the same location as the database control files.
- 4 Create a PFILE to specify the location of the SPFILE. The PFILE must be in the same directory on the shared disk as the SPFILE.

Modify the PFILE so that it contain only the following entry:

```
SPFILE = spfile_location
```

The variable *spfile_location* represents the fully qualified path of the SPFILE on the shared disk.

For example, if your SPFILE is located at path `H:\Oracle\Data`, create a PFILE `sp_init.ora` in the same location `H:\Oracle\Data` with the following parameter:

```
SPFILE = H:\Oracle\Data\spfileSID.ora
```


Setting the Oracle database parameters

Set the Oracle database parameters using the oradim utility. See the Oracle documentation for more information about this utility.

To set the Oracle database parameters

- 1 Type the following on the command prompt:

```
C:\>set ORACLE_SID=<SID_NAME>
```

The variable *SID* represents the database instance.

- 2 Run the oradim utility with parameters as specified.

Type the following on the command prompt:

```
C:\><ORACLE_HOME>\bin\oradim -edit -sid SID -pfile  
"pfile_location_on_shared_disk" -startmode auto  
-timeout 100 -shuttype srvc,inst -shutmode i
```

The variable *SID* represents the database instance.

This command performs the following tasks:

- Sets the `-startmode` parameter to start the database instance automatically when Oracle service is started.
 - Sets the `-timeout` parameter to 100. The timeout parameter specifies the maximum wait time (in seconds) for the database shutdown to complete before the service for a particular SID stops.
 - Sets the `-shuttype` parameter to stop both the database instance and the service while shutting the database.
 - Sets the `-shutmode` parameter to shut the database instance as soon as the VCS database agent for Oracle requests.
- 3 Start the Oracle database service using the Services snap-in.

Verifying the initialization file settings

If a PFILE is used:

- 1 Open the Registry Editor.
From the **Start** menu, choose **Run**. In the **Open** field, enter **regedit** and click **OK**.
- 2 Double-click the `ORA_SID_PFILE` registry key at
`HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\HOME_ID\`.
The variable `SID` represents the database instance.
- 3 Verify that the Value data field specifies the fully qualified path of the PFILE along with the file name.

If an SPFILE is used:

- 1 Stop and restart the database services.
- 2 Run `sqlplus.exe`.
- 3 Connect to the database.
- 4 Verify the following query returns the correct shared virtual disk path for the SPFILE.

```
select value from v$parameter where name = 'spfile';
```

Configuring the Oracle and listener services

Complete these steps to configure the Oracle and Listener services.

To configure the Oracle and listener services

- 1 Open the **Services** applet (**Start > All Programs > Administrative Tools > Services**).
- 2 Set the startup type of the Oracle database service to manual and configure its Log On properties:
 - Double-click the Oracle database service. In the SCM, Oracle database services appear as `OracleServiceSID`, where `SID` represents the database instance.
 - In the **Properties** window, click the **General** tab.
 - From the **Startup Type** drop-down list, select **Manual**.
 - Click **Stop** to stop the service.
 - Click the **Log On** tab.
 - Choose **This Account**.
 - Enter the credentials of the user in whose context Oracle was installed.

- Click the **General** tab and click **Start** to start the service with the new Log On properties. Click **OK**.
- 3 Set the startup type of the Listener service to manual.
- In the Services applet, double-click the listener service. Listener services appear as Oracle*Ora_Home*TNS*ListenerName*, where *Ora_Home* represents the Oracle home directory and *ListenerName* is the name of the listener set during the installation.
 - In the **Properties** window, click the **General** tab.
 - From the **Startup Type** drop-down list, select **Manual** and click **OK**.

Disconnecting virtual disks from the first node

Disconnect the virtual disks only if you plan to install Oracle on more nodes. If it's the last node for the Oracle installation, do not disconnect the virtual disks. Before disconnecting the virtual disks, stop the database and listener services.

See "[Disconnecting virtual disks from the cluster nodes](#)" on page 44.

Releasing disks (if you use Windows LDM)

Unassign the drive letters and then release the reserved disks only if you plan to install Oracle on more nodes. If it's the last node for the Oracle installation, do not disconnect the virtual disks. Before disconnecting the virtual disks, stop the database and listener services.

See "[Unassigning a drive letter](#)" on page 46.

See "[Releasing disks](#)" on page 46.

Installing and configuring Oracle on failover nodes

Complete the following steps if you wish to install Oracle on failover nodes.

Installing Oracle on failover nodes

Install Oracle from the Oracle software disc. The installer screens and options vary slightly for different versions of Oracle.

Refer to the Oracle documentation for detailed installation instructions:

Note the following before you install Oracle on additional nodes:

- Install Oracle in software only mode. Do not configure the Oracle database while installing Oracle on additional cluster nodes.
- Install the Oracle software program files on a local disk.
To ensure application installations are identical, the values you specify while installing and configuring Oracle on failover nodes must match the values specified during the installation on the first node.
- The Oracle Home name and path should be the same on all the cluster nodes.
- Do not create the Oracle Starter database.

Configuring Oracle on failover nodes

The following steps do not apply for the system on which Oracle services were installed first. For example, if the first system is system A, and subsequent installations are performed on systems B and C, perform the following steps only from systems B and C.

- 1 Verify that the volumes or LUNs created for the Oracle database, control files, log files, bdump, cdump, udump files, flash recovery area and archive location are mounted or connected on the node where you run the installer.
See “[Managing shared storage](#)” on page 42.
- 2 Bring the temporary service group online on the node where you run the installer.
See “[Bringing the temporary service group online](#)” on page 48.
- 3 Verify that the Oracle log files and the background_dump_dest (bdump), core_dump_dest (cdump), and user_dump_dest (udump) files are on the same volume on the shared storage.
- 4 Type the following on the command prompt:

```
C:\>set ORACLE_SID=<SID NAME>
```


The variable *SID* represents the database instance.

- 5 Run the oradim utility with the parameters as specified.

Type the following on the command prompt:

```
C:\><ORACLE_HOME>\bin\oradim -new -sid SID -pfile  
"pfile_location_on_shared_disk" -startmode auto  
-timeout 100 -shuttype srv,inst -shutmode i
```

The variable *SID* represents the database instance.

This command performs the following tasks:

- Sets the `-startmode` parameter to start the database instance automatically when Oracle service is started.
 - Sets the `-timeout` parameter to 100. The timeout parameter specifies the maximum wait time (in seconds) for the database shutdown to complete before the service for a particular SID stops.
 - Sets the `-shuttype` parameter to stop both the database instance and the service while shutting the database.
 - Sets the `-shutmode` parameter to shut the database instance as soon as the VCS database agent for Oracle requests.
- 6 Create a listener service using the Oracle Net Configuration Assistant. Refer to the Oracle documentation for instructions.
 - 7 Stop the listener service and make a backup copy of the original files, `tnsnames.ora` and `listener.ora`, from the first node.
 - 8 Copy the files, `tnsnames.ora` and `listener.ora` from the first node to the current node, at the same directory as that on the first node. The path for `tnsnames.ora` and `listener.ora` must be the same on all failover nodes.
 - 9 From the Services snap-in, set the startup type of the Listener service to Manual and then restart the Listener and the Oracle service for the database instance on the node.
In the Services snap-in, the names of Oracle services appear as `OracleServiceSID`, where *SID* represents the database instance.
 - 10 Verify access to the database by connecting to it and running a `select` query.

Repeat these tasks on all the remaining nodes that will participate in the Oracle service group.

Configuring the Oracle service group

This chapter contains the following topics:

- [“About configuring the Oracle service group”](#) on page 64
- [“Deleting the temporary service group”](#) on page 64
- [“Configuring the Oracle service group using the wizard”](#) on page 65
- [“Verifying the Oracle service group configuration”](#) on page 73

About configuring the Oracle service group

Configuring the Oracle service group involves creating the Oracle service group, its resources, and defining attribute values for the configured resources.

VCS provides several ways to configure the Oracle service group; the Oracle configuration wizard, the Java and Web consoles, and the command line. This chapter provides instructions on how to use the Oracle configuration wizard to configure the service group.

Deleting the temporary service group

Delete the temporary service group before configuring the Oracle service group.

To delete the temporary service group

- 1 Bring the temporary service group online on a cluster node.
From the Cluster Manager (Java Console), right-click the temporary service group name in the left pane, click **Online** and then click the node name from the list.
- 2 Click **Yes** on the confirmation dialog box.
- 3 Delete the temporary service group.
From the Cluster Manager (Java Console), right-click the temporary service group name in the left pane and click **Delete**.
- 4 Click **Yes** on the confirmation dialog box.

Configuring the Oracle service group using the wizard

VCS provides a configuration wizard that guides you through the process of configuring the Oracle service group. You can use the wizard to create and modify Oracle service groups.

This section describes the steps required to create a new Oracle service group using the wizard. To modify an existing service group, see [“Modifying an Oracle service group”](#) on page 78.

Before configuring the Oracle agent, review the resource type and the attribute definitions of the Oracle agents, described in [Appendix A, “Resource type definitions”](#) on page 93.

For sample configuration file and resource dependency graph of the Oracle service group, see [Appendix B, “Sample configuration”](#) on page 103.

Before configuring the Oracle service group

- If you have configured Windows Firewall, add the following to the Firewall Exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`.
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141For a detailed list of services and ports used by VCS, refer to the *Veritas Cluster Server Installation and Upgrade Guide*.
- Verify that VCS along with the VCS agent for Oracle is installed on all cluster nodes.
See [Chapter 2, “Installing the VCS agent for Oracle”](#) on page 19.
- Verify a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
See [“Configuring the cluster”](#) on page 21.
- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be a Domain Administrator on the node where you run the wizard.
- Verify that the Veritas High Availability engine (HAD) is running on the system from where you run the wizard.

- Connect the LUNs or mount the volumes containing the data files, control files, redo log files, bdump, cdump, and udump files. Unmount or disconnect the volumes or LUNs from other nodes in the cluster.
- Mount the database and start the Oracle instance on the node running the wizard.
- Make sure that you have the following information ready; the wizard will prompt you for this information:
 - The databases and listeners to be monitored by VCS.
 - For the instances to be monitored in detail, name and location of the respective SQL files.
 - A valid domain name, user name, and password with which the database service was configured for the database.

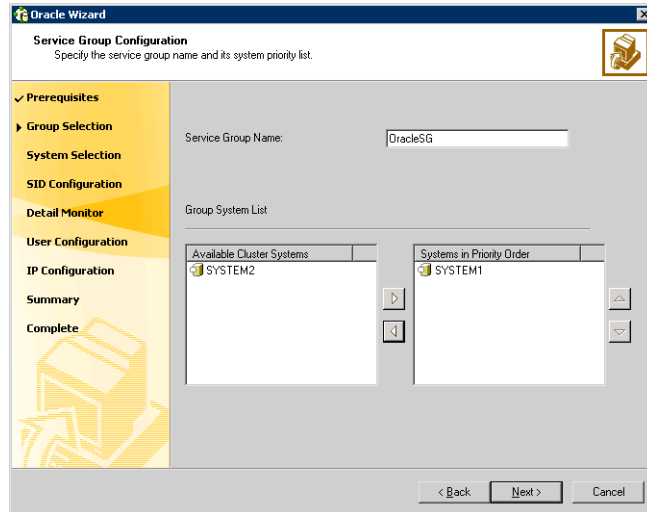
Configuration instructions

This section describes how to create an Oracle service group.

To create an Oracle service group

- 1 Start the Oracle configuration wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard**.
- 2 In the Welcome panel click **Next**.
- 3 In the Wizard Options panel, select **Create service group** and click **Next**.

- 4 In the Service Group Configuration panel, complete the following and click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the Oracle service group.

Available Cluster Systems

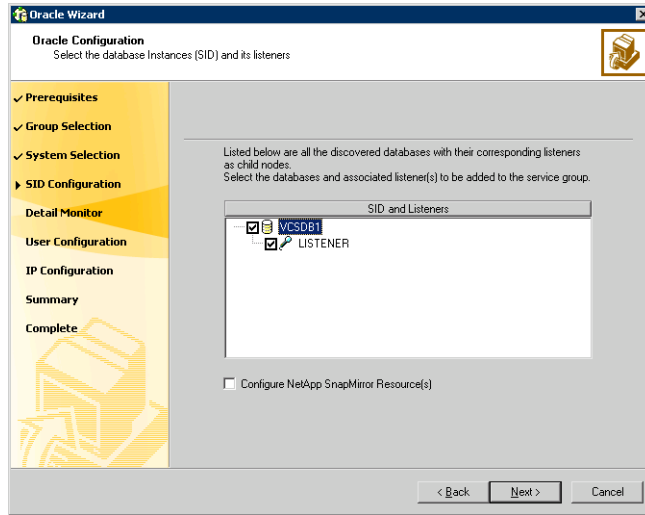
Select the systems on which to configure the service group and click the right arrow to move the systems to the Systems in Priority Order box.

The Systems in Priority Order box represents the service group's system list. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

To remove a system from the service group's system list, select a system in the Systems in Priority Order box and click the left arrow.


To change a system's priority in the service group's system list, select the system from the Systems in Priority Order box, and click the up and down arrows.

- 5 In the Oracle Configuration panel, complete the following and then click **Next**.



- Select the SIDs and the associated listeners to be added to the service group. The SID is a system identifier that uniquely identifies the Oracle database instance, while the listener is the name of the corresponding listener service.
- Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

- 6 On the Detail Monitoring panel, configure detail monitoring for the Oracle database if required, and click **Next**.

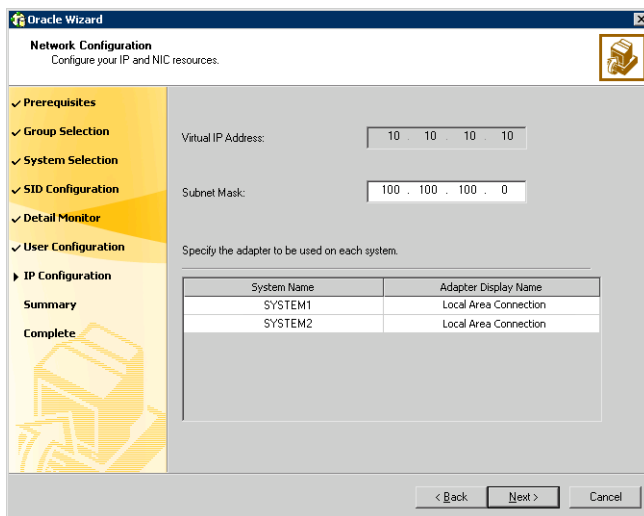


Detail Monitor Check the Detail Monitor option corresponding to each database that you want to configure detail monitoring for.

SQL Path Type the path of the SQL file that will query the database to validate the status. Click the icon next to the field to browse for the SQL file.
 A sample SQL file, check.sql, is located at
`%VCS_HOME%\bin\Oracle\.`

- 7 In the Domain and User selection panel, type a valid domain name, user name, and password with which the database service was configured for the database and click **Next**.
 See “[Configuring the Oracle and listener services](#)” on page 58 for more information.
- 8 On the Initiator Selection panel, select the initiators for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.
 If you are configuring MPIO over FC, you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.
- 9 On the Network Configuration panel, specify the network related information and click **Next**.

The wizard discovers and displays the virtual IP address for the Oracle Server.

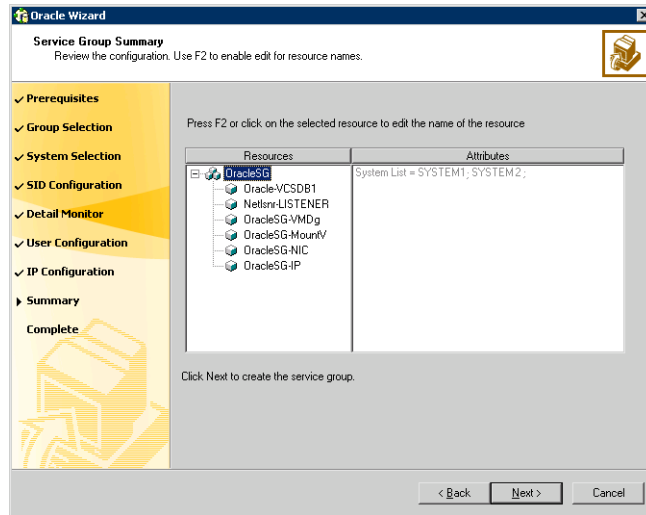


Subnet Mask Type the subnet mask to which the virtual IP belongs.

Adapter Display Name For each system in the cluster, select the public network adapter. Select the Adapter Name field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- 10 Review the configuration on the Summary panel and click **Next**.



Resources Lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box. The wizard assigns unique names to the resources. To edit a resource name, select the resource name and click on it, or press the F2 key. After the edit, press the Esc key to cancel the changes, or press the Enter key to confirm the changes.

- 11 In the confirmation dialog box, click **Yes**. Click **No** if you wish to review your settings.
 The wizard starts running commands to create the Oracle service group.
- 12 On the Completing the Oracle Configuration panel, check **Bring the service group online** to bring the service group online on the local system, and click **Finish**. The Oracle service group is created in your cluster.

Configuring dependent services

If the database service has other dependent services, make sure the dependent services are running on the node where the database service is online. Note that the online agent operation brings only the database service online and not the dependent services.

For example, on Oracle 10g, the DBConsole service corresponding to an Oracle database has a dependency on the database service. That is, for the Enterprise Manager to manage the databases, you must make sure the DBConsole service is running on the node where the database service is online.

To configure a dependent service

- 1 For the dependent service, add a GenericService resource manually.
- 2 Make the GenericService resource dependent on the corresponding Oracle resource.
- 3 Set the Critical attribute to False if the Oracle service group must not fail over when the GenericService resource faults.

Refer to the *Veritas Cluster Server Administrator's Guide* for configuration instructions.

Verifying the Oracle service group configuration

This section provides steps to verify an Oracle service group configuration by bringing the service group online, taking it offline, and switching the service group to another cluster node.

Bringing the service group online

Perform the following steps to bring the service group online from the VCS Java or Web Console.

To bring a service group online from the Java Console

- 1 In the Cluster Explorer configuration tree, select the Oracle service group to be taken online.
- 2 Right-click the service group and select to online the service group on the system. (Right-click > Online > *system_name*)

To bring a service group online from the Web Console

- 1 On the **Service Group** page (**Cluster Summary > All Groups > Service Group**), click **Online**.
- 2 In the Online Group dialog box, select the system on which to bring the service group online.
- 3 To run PreOnline script, select the **Run PreOnline Script** check box.
- 4 Click **OK**.

Taking the service group offline

Perform the following steps to take the service group offline from the VCS Java or Web Console.

To take a service group offline from the Java Console

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
or
Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Offline**, and choose the appropriate system from the pop-up menu. (Right-click > Offline > *system_name*)

To take a service group offline from the Web Console

- 1 On the **Service Group** page (**Cluster Summary > All Groups > Service Group**), click **Offline**. This opens the Offline Group dialog box.
- 2 Select the system on which to take the service group offline.
- 3 Click **OK**.

Switching the service group

The process of switching a service group involves taking it offline on its current system and bringing it online on another system.

To switch a service group from the Java Console

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
or
Select the cluster in the Cluster Explorer configuration tree, select the Service Groups tab, and right-click the service group icon in the view panel.
- 2 Choose **Switch To**, and choose the appropriate system from the pop-up menu. (Right-click > Switch To > *system_name*)

To switch a service group from the Web Console

- 1 From the Service Group page (**Cluster Summary > All Groups > Service Group**), click **Switch**.
- 2 In the Switch Group dialog box, select the system to switch the service group to.
- 3 Click **OK**.

Disabling the service group

To disable the agent, you must change the Oracle service group to an OFFLINE state. You can stop the application completely or switch the agent to another system.

To disable a service group from the Java Console

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
or
Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Disable**, and choose the appropriate system from the pop-up menu.

To disable a service group from the Web Console

- 1 In the **Service Group** page (**Cluster Summary > All Groups > Service Group**), click **Disable**. This opens the Disable Group dialog box.
- 2 Select the system on which to disable the service group. To disable the service group on all systems, click **All Systems**.

Administering the Oracle service group

This chapter contains the following topics:

- [“About administering the Oracle service group”](#) on page 78
- [“Modifying an Oracle service group”](#) on page 78
- [“Deleting an Oracle service group”](#) on page 79
- [“Configuring detail monitoring”](#) on page 80

About administering the Oracle service group

This chapter describes the administrative tasks you can perform on Oracle service groups such as modifying the service group configuration and deleting a service group. The chapter also describes steps to configure detail monitoring for an Oracle database.

Modifying an Oracle service group

The following steps describe how to modify the configuration of the service groups using the configuration wizard.

Before modifying the Oracle service group

- If the Oracle service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
You can however enable and disable the SIDs and listeners to be monitored in detail and change their detail monitoring options when the Oracle service group is online.
- To change the resource attributes, you must take the service group offline. However, the NetAppFiler and NetAppSnapDrive resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

Instructions

To modify an Oracle service group

- 1 Start the Oracle configuration wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group to modify, and click **Next**.
- 4 Follow the wizard instructions and make desired modifications to the service group configuration. See [“Configuring the Oracle service group using the wizard”](#) on page 65 for instructions.

Deleting an Oracle service group

The following steps describe how to delete an Oracle service group using the configuration wizard.

To delete an Oracle service group

- 1 Start the Oracle configuration wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Delete service group**, select the service group to delete, and click **Next**.
- 4 In the Service Group Summary panel, click **Next**.
- 5 On the message that informs you that the wizard will run commands to delete the service group, click **Yes** to delete the service group and then click **Finish**.

Configuring detail monitoring

Use the detail monitoring capability of VCS database agent for Oracle to monitor the status of a database. Before setting up detail monitoring, you must have the agent running at the basic level of monitoring, that is, the DetailMonitor attribute must be set to False.

The Oracle agent uses a script to monitor the status of the database. A sample SQL script, located at %VCS_HOME%\bin\Oracle\check.sql, is provided with the agent for the purpose. If the script is successfully executed during monitoring, the agent considers the database instance available. If the execution fails, the database instance is considered not available and the service group faults and fails over to the failover nodes. You can customize the script to meet your configuration requirements.

Note: You must use a separate script for each Oracle service group that exists in the cluster. The script must exist on all the nodes in the service group.

Symantec recommends that before configuring detail monitoring, you must ensure that Oracle is configured correctly and you are able to connect to the database.

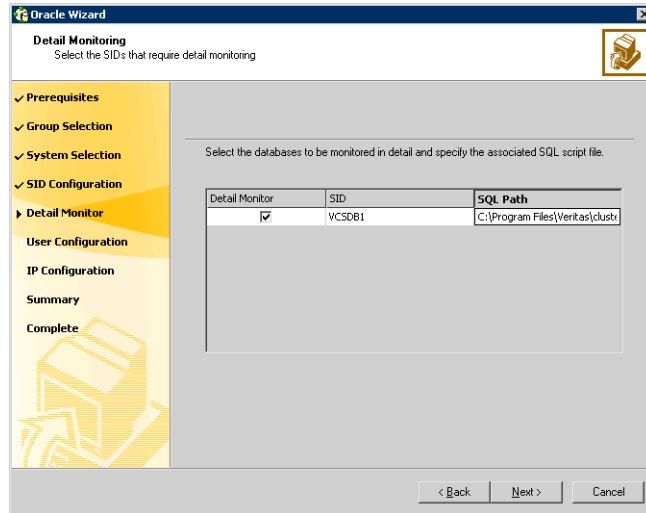
Enabling detail monitoring

The following steps describe how to enable detail monitoring using the configuration wizard.

To enable detail monitoring

- 1 Start the Oracle configuration wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group configured for the instance to be monitored in detail, and click **Next**.
- 4 In the Oracle Configuration panel, select the SIDs to be configured along with their respective listeners and click **Next**.

- 5 In the Detail Monitoring dialog box, specify information to enable detail monitoring, and click **Next**.



Detail Monitor Check the Detail Monitor option corresponding to each database that you want to configure detail monitoring for.

SQL Path Type the path of the SQL file that will query the database to validate the status. Click the icon next to the field to browse for the SQL file.
 A sample SQL file, check.sql, is located at
`%VCS_HOME%\bin\Oracle\.`

- 6 In the Domain and User selection panel, type a valid domain name, user name, password, and then click **Next**.
- 7 Follow the wizard instructions and accept the default values in the subsequent dialog boxes. See [“Configuring the Oracle service group using the wizard”](#) on page 65 for instructions.

Disabling detail monitoring

The following steps describe how to disable detail monitoring using the configuration wizard.

To disable detail monitoring

- 1 Start the Oracle configuration wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group configured for the instance for which detail monitoring is being disabled, and click **Next**.
- 4 In the Oracle Configuration panel, click **Next**.
- 5 In the Detail Monitoring Configuration panel, uncheck the check box corresponding to the Oracle Server instance for which detail monitoring is being disabled and click **Next**.
- 6 Follow the wizard instructions and accept the default values in the subsequent dialog boxes. See [“Configuring the Oracle service group using the wizard”](#) on page 65 for instructions.

Troubleshooting the agents

This chapter contains the following topics:

- [“About troubleshooting VCS agents”](#) on page 84
- [“Troubleshooting common problems”](#) on page 84
- [“VCS logging”](#) on page 85
- [“Agent error messages and descriptions”](#) on page 88

About troubleshooting VCS agents

This chapter lists some commonly encountered problems with the Oracle agent and some possible solutions. It also lists the error messages associated with the VCS database agent for Oracle. Each message includes a description and a recommended solution, if applicable.

Troubleshooting common problems

Occasionally you may encounter problems when using the Oracle database agent. [Table 6-1](#) describes the most common problems and provides standard workarounds.

Table 6-1 Common problems when using Oracle database agent

Error	Description
The Oracle service starts, but does not start the database.	<p>Make sure the Oracle database parameter <code>startmode</code> is set to <code>auto</code>.</p> <p>Make sure the path of the initialization file is set correctly.</p> <p>Verify access to the shared virtual disk.</p>
Clients cannot connect to the virtual IP address set for the service group.	Make sure that the listener service is running. Verify that the listener is configured with virtual IP address in the host location.
Resource returns offline during detail monitoring.	<p>This could occur because the database is offline or has not been mounted or opened properly. Check the error in the log file for further details.</p> <p>Verify the configuration definition of the agent. Make sure that the correct file name is specified in the configuration. Verify access to sqlplus from the command prompt.</p>
The sqlplus script cannot be run because sqlplus returns a process error.	If the script timed out, increase the timeout value set in the attribute <code>SQLTimeOut</code> . Otherwise, try recreating the Oracle service using the <code>oradim</code> utility.
The Oracle wizard could not locate PFILE or SPFILE for a database.	<p>Make sure the initialization file settings are configured correctly.</p> <p>See “Verifying the initialization file settings” on page 58.</p>

Table 6-1 Common problems when using Oracle database agent (Continued)

Error	Description
<p>Listener service does not start after configuring the virtual IP address.</p>	<p>Make sure the virtual IP address is configured correctly for the listeners.</p> <p>Another possible cause is that the listener.ora file is corrupted.</p> <p>Solution:</p> <ol style="list-style-type: none"> 1 Stop the listener services on the node. 2 Rename the listener.ora file. 3 Recreate the listeners using the Net Configuration Assistant. 4 Bring up the virtual IP on the node. 5 Configure the virtual IP address for the newly created listeners using the Net Manager. 6 Restart the listener services. 7 Run <code>lsnrctl status</code> command for each listener to verify the virtual IP address. <p>Refer to the Oracle documentation for instructions on using the Net Configuration Assistant and the Net Manager.</p>

VCS logging

VCS generates two error message logs: the engine logs and the agent logs. Log file names are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The agent log is located at `%VCS_HOME%\log\agent_A.txt`. The format of agent log messages is:

Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type |
Resource Name | Entry Point | Message Text

The agent log message components are defined as follows:

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Symantec product logs the message. For VCS database agent for Oracle, mnemonic is 'VCS'.
- Severity denotes the seriousness of the message. Severity of the VCS error messages is classified into the following types:
 - CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately.

- ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action.
- WARNING indicates a warning or error, but not an actual fault.
- NOTE informs the user that VCS has initiated an action.
- INFO informs the user of various state messages or comments.
Of these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.
- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the Oracle agent would resemble: V-16-20046-5
Originator ID for all VCS products is 'V-16.' Category ID for Oracle agent is 20046, for Netlsnr agent is 20047. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.

You can view these message logs using Notepad or any text editor. All messages are logged to the engine and the agent logs. Messages of type CRITICAL and ERROR are written to the Windows event log.

A typical agent log resembles:

```
2005/05/10 09:32:44 VCS NOTICE V-16-20046-5 Oracle:
Oracle-DATA1:online:Service status is Running
```

VCS Cluster Configuration Wizard logs

The VCS Cluster Configuration Wizard (VCW) log is located at
%allusersprofile%\Application Data\Veritas\Cluster
Server\vcw.log.

Here, %allusersprofile% is the file system directory containing application data for all users. A typical path is C:\Documents and Settings\All Users\.

The format of the wizard logs is

ThreadID | Message Text

- *ThreadID*: the ID of the thread initiated by the wizard.
- *Message Text*: the actual message generated by the wizard.

A typical wizard log resembles:

```
00000576-00000264: ExecMethod return 00000000.
00000576-00000110: CRegistry::Query for VCS License failed.
Error=0x00000000
00000576-00000264: ExecMethod return 00000000.
00000576-00000264: ExecMethod return 00000001.
00000576-00000127: QueryDWORDValue returned 0x00000001
```

```
00000576-00000132: CRegistry::Query for VxSS Root information
failed. Error=0x00000001
```

VCWsilent logs

The VCWsilent log is located at `<currentdirectory>\vcwsilent.log`.

Here, `<currentdirectory>` is the directory from where the VCWsilent.exe is run.

A typical VCWsilent log resembles:

```
00005540-00000064: 5540: STARTING - Discovering NICs on the
selected machines...
00009956-00000064: 9956: STARTING - Generating private network
related files...
00009956-00000048: 9956: COMPLETED - Gererating LLT host
files...
00009956-00000048: 9956: COMPLETED - Generating GAB tab files...
00009956-00000048: 9956: COMPLETED - Generating main.cf file...
00009956-00000064: 9956: STARTING - Configuring LLT on all the
nodes.
00009956-00000048: 9956: COMPLETED - Configuring LLT on all the
nodes.
```

Agent error messages and descriptions

Table 6-2 and Table 6-3 lists the messages of type ERROR and WARNING. Each message includes a description and a recommended solution, if applicable.

Oracle agent

Table 6-2 Oracle agent error messages

Message	Description
Cannot open service handle.	<p>The agent could not open the Service Control Manager handle. This may occur because of an invalid service name.</p> <p>Solution: Verify that the correct service name is specified in the configuration definition of the agent.</p>
Cannot close service handle.	<p>The agent could not close the Service Control Manager handle.</p> <p>Solution: Verify that the specified service is running.</p>
Cannot start service.	<p>The agent could not start the specified service. This may occur because of the following:</p> <ul style="list-style-type: none"> ■ The specified SCM handle does not have access to the service. or ■ The specified handle is invalid. <p>Solution: Verify that the service name in the configuration file is the same as the name displayed by the SCM. Verify that the specified service is supported by the host. Verify that the service is listed in the SCM window.</p>
Cannot stop the service.	<p>The agent could not stop the specified service. This may occur because of the following:</p> <ul style="list-style-type: none"> ■ The specified service was not opened successfully. or ■ The specified service was not opened with the required access. <p>Solution: Verify that the service was opened successfully. Check the access with which the service was opened.</p>

Table 6-2 Oracle agent error messages (Continued)

Message	Description
Wait for mutex failed.	The wait function for mutex failed. This message is followed by an error code. Check the error code for more information.
Wait for mutex abandoned.	The mutex object is was not released by the thread that owned the mutex object before the thread terminated. This indicates a possibility that the SQL script is not terminated. Solution: Verify that the SQL script executed correctly.
Cannot wait for mutex object.	The agent could not acquire mutex for the resource. This may occur if two or more Oracle resources configured. Solution: Verify that the SQL script executed correctly.
Failed to decrypt password.	The agent could not decrypt the domain password. Solution: Verify that the password was encrypted using the VCS Encrypt utility.
No domain specified in argument list.	The attribute Domain is not configured for the resource. Solution: Verify that the correct Domain is defined in the configuration file.
No username specified in argument list.	The attribute UserName is not configured for the resource. Solution: Verify that the correct UserName is defined in the configuration file.
No password specified in argument list.	The attribute Password is not configured for the resource. Solution: Verify that the password is encrypted using the VCS Encrypt utility. Verify that the correct encrypted password is defined in the configuration file.
No SID specified in argument list.	The configuration definition of the agent does not contain a value for the attribute SID. Solution: Verify that the attribute SID is correctly defined in the configuration file.
SQL*Plus process exited with error <i>error_code</i> .	The agent could not execute the SQL*Plus process correctly. See the associated error code for more information.

Table 6-2 Oracle agent error messages (Continued)

Message	Description
Invalid filename in argument list file_name.	The SQL script for the resource is incorrectly defined in the filename. Solution: Verify that the attribute SQLFile is correctly defined in the configuration file. Verify the validity of the file name and path.
An Oracle error occurred.	An error occurred while generating the SQL script. This is an Oracle error and the corresponding error number and error message is logged to the agent log.

Netlsnr agent

Table 6-3 Netlsnr agent error messages

Message	Description
Cannot open service handle.	The agent could not open the Service Control Manager handle. This may occur because of an invalid service name. Solution: Verify that the correct service name is specified in the configuration definition of the agent.
Cannot start service.	The agent could not start the specified service. This may occur because of the following: <ul style="list-style-type: none"> ■ The specified SCM handle does not have access to the service. or ■ The specified handle is invalid. Solution: Verify that the service name in the configuration file is the same as the name displayed by the SCM. Verify that the specified service is supported by the host. Verify that the service is listed in the SCM window.

Table 6-3 Netlsnr agent error messages (Continued)

Message	Description
Cannot stop the service.	<p>The agent could not stop the specified service. This may occur because of the following:</p> <ul style="list-style-type: none"> ■ The specified service was not opened successfully. ■ The specified service was not opened with the required access. <p>Solution: Verify that the service was opened successfully. Check the access with which the service was opened.</p>
Cannot close service handle.	<p>The agent could not close the Service Control Manager handle.</p> <p>Solution: Verify that the specified service is running.</p>

NetApp agents

This section describes the error messages for the NetApp agents.

Table 6-4 NetApp agents error messages

Message	Description
Failed to open connection to filer %s.	<p>Make sure that the VCS Helper service account is part of the administrator's group on the local host and the filer.</p> <p>Make sure the private network is functioning properly. Verify you can ping the IP used for the private storage network. This is the IP defined the StorageIP attribute of the NetAppFiler resource.</p>
Failed to initialize ONTAPI on system	<p>The agent could not find the file NTAPADMIN.DLL on the system. Verify the file exists in the %VCS_HOME%\bin directory</p>
Invalid attributes exist in the configuration	<p>Some agent attributes have not been defined or have been defined incorrectly. Verify the configuration definition for the agent.</p>
ONTAP API called failed for object_name on filer_name.	<p>The specified API failed on the specified object. See the NetApp ONTAP API documentation for information about the associated error message</p>
Volume %s on filer %s is not a SnapMirror replicated volume	<p>Verify replication is set up on the specified volume.</p>

Table 6-4 NetApp agents error messages (Continued)

Message	Description
<p>Multiple snapmirror destinations for a volume is not supported by this agent. 'snapmirror status' for volume %s on filer %s returned multiple status entries. Administrative intervention required</p>	<p>There should be only one destination per source volume.</p>
<p>Initialize VLibNetAppHost::Initialize() failed. (error_type: %s, error_code: 0x%s)</p>	<p>The agent could not detect the iSCSI or FC Initiator on the host. Make sure that you have installed and configured Microsoft iSCSI Initiator or an FC Initiator on each node.</p>
<p>Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s, error_message: %s)</p>	<p>This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration:</p> <ul style="list-style-type: none"> ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP <p>Verify the configuration definition of the resource. Make sure each attribute is defined correctly.</p>
<p>Unable to create/delete online lock file %s. Error code %s,</p>	<p>Make sure you have write permissions on the specified directory.</p>

Resource type definitions

This chapter contains the following topics:

- [“About resource type definitions”](#) on page 94
- [“Oracle agent”](#) on page 94
- [“Netlsnr agent”](#) on page 97
- [“NetAppFiler agent”](#) on page 98
- [“NetApp SnapDrive agent”](#) on page 100
- [“NetApp SnapMirror agent”](#) on page 101

About resource type definitions

This appendix lists the resource type definitions and attribute definitions of the Oracle and the NetApp agents. The resource type represents the VCS configuration definition of the agent and specifies how the agent is defined in the configuration file main.cf. The Attribute Definitions table explains the attributes associated with the agent. The Required attributes table lists the attributes that must be configured for the agent to function properly.

Oracle agent

The Oracle agent is configured as a resource of type Oracle.

Resource type definition

```
type Oracle (  
    static keylist SupportedActions = { DBRestrict, DBUndoRestrict,  
    DBResume, DBSuspend, DBQuiesce }  
    static i18nstr ArgList[] = { ServiceName, DelayAfterOnline,  
        DelayAfterOffline, DetailMonitor, IndepthMonitorCount, SID,  
        Domain, UserName, EncryptedPasswd, SQLFile, SQLTimeOut }  
    str ServiceName  
    int DelayAfterOnline = 20  
    int DelayAfterOffline = 20  
    boolean DetailMonitor = 0  
    int IndepthMonitorCount = 5  
    str SID  
    str Domain  
    str UserName  
    str EncryptedPasswd  
    str SQLFile  
    int SQLTimeOut = 20  
)
```

Attribute definitions

Table A-1 Oracle agent - required attributes

Required Attribute	Type and Dimension	Definition
ServiceName	string-scalar	The name of the Oracle service set during the installation of Oracle. The service name is specified in the format OracleService <i>SID</i> , where <i>SID</i> represents the database instance.
Domain	string-scalar	The name of the domain to which the user belongs in whose context Oracle was installed.
SID	string-scalar	A system identifier that uniquely identifies the Oracle database instance to be monitored in detail.
UserName	string-scalar	The name of the domain user in whose context Oracle was installed.
EncryptedPasswd	string-scalar	The encrypted password for the user identified by UserName. The password must be encrypted with the VCS Encrypt utility.

Table A-2 Oracle agent - optional attributes

Optional Attributes	Type and Dimension	Definition
DelayAfterOffline	integer-scalar	Number of seconds the agent waits to start the monitoring routine after Oracle is taken offline. Default = 20
DelayAfterOnline	integer-scalar	Number of seconds the agent waits to start the monitoring routine after Oracle is brought online. Default = 20

Table A-2 Oracle agent - optional attributes (Continued)

Optional Attributes	Type and Dimension	Definition
DetailMonitor	boolean-scalar	<p>A flag that defines whether the agent monitors Oracle in detail by trying to verify access to the database by running a SQL script. The value 1 indicates that the agent will monitor Oracle in detail. The value 0 indicates that the agent will not.</p> <p>Default = 0</p> <p>If the DetailMonitor flag is set to 1, the following attributes must be defined:</p> <ul style="list-style-type: none"> ■ SID ■ SQLFile
IndepthMonitorCount	integer-scalar	<p>An integer that defines the monitor interval between two consecutive detail monitoring cycles. If this attribute is set to 5, the agent monitors the application in detail after every five monitor intervals.</p> <p>Default = 5</p> <p>Note: Do not set the value of this attribute to 0.</p>
SQLTimeOut	integer-scalar	<p>The number of seconds the agent waits for the SQL script to return a value during detail monitoring. If this limit exceeds, the resource is declared faulted.</p> <p>Default = 20</p>

Table A-2 Oracle agent - optional attributes (Continued)

Optional Attributes	Type and Dimension	Definition
SQLFile	string-scalar	<p>The name of SQL script to be used for detail monitoring, along with its complete path. The agent verifies the existence of this file.</p> <p>A sample file, check.sql, is located at <code>%VCS_HOME%\bin\Oracle\</code>. The sample file contents are:</p> <pre>select * from v\$database;</pre> <p>Do not include any spool commands within the SQL script.</p> <p>If an Oracle error occurs during execution, the error is logged to the agent log. If debug logging is also enabled for the resource, the output of the SQL script is also logged to the agent log.</p>

Netlsnr agent

The Netlsnr agent is configured as a resource of type Netlsnr.

Resource type definition

```
type Netlsnr (
    static i18nstr ArgList[] = { ServiceName, DelayAfterOnline,
        DelayAfterOffline }
    str ServiceName
    int DelayAfterOnline = 20
    int DelayAfterOffline = 20
)
```

Attribute definitions

Table A-3 Netlsnr agent - required attributes

Required Attribute	Type and Dimension	Definition
ServiceName	string-scalar	The name of the configured listener service. The service name is specified in the format <i>OracleOraHomeTNSListenerName</i> , where <i>OraHome</i> represents the Oracle home set during the installation and <i>ListenerName</i> represents the name of the configured listener.

Table A-4 Netlsnr agent - optional attributes

Optional Attributes	Type and Dimension	Definition
DelayAfterOffline	integer-scalar	Number of seconds the agent waits to start the monitor routine after the listener is taken offline. Default = 20
DelayAfterOnline	integer-scalar	Number of seconds the agent waits to start the monitor routine after the listener is brought online. Default = 20

NetAppFiler agent

The NetApp Filer agent is configured as a resource of type NetAppFiler.

Resource type definition

```

type NetAppFiler (
    static int MonitorInterval = 30
    static i18nstr ArgList[] = { FilerName, StorageIP }
    static str Operations = None
    str FilerName
    str StorageIP
)

```

Attribute definitions

Table A-5 NetApp Filer agent attributes

Attribute	Type - Dimension	Description
FilerName	string-scalar	DNS-resolvable name or IP address of the locally attached filer.
StorageIP	string -scalar	The private storage IP address of the filer.

NetApp SnapDrive agent

The NetApp SnapDrive agent is configured as a resource of type NetAppSnapDrive.

Resource type definition

```

type NetAppSnapDrive (
    static int MonitorInterval = 30
    static int NumThreads = 1
    static i18nstr ArgList[] = { FilerResName,
    "FilerResName:FilerName", "FilerResName:StorageIP", VolumeName,
    ShareName, LUN, MountPath, Initiator, InitiatorMonitorInterval }
    str FilerResName
    str VolumeName
    str ShareName
    str LUN
    str MountPath
    str Initiator[]
    int InitiatorMonitorInterval = 30
)

```

Attribute definitions

Table A-6 NetApp SnapDrive agent attributes

Attribute	Type - Dimension	Description
FilerResName	string-scalar	Name of the VCS NetAppFiler-type resource in the service group.
VolumeName	string-scalar	Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer.
ShareName	string-scalar	Name of the CIFS share containing the virtual disk.
LUN	string-scalar	Name of the LUN (virtual disk) on the filer that is presented to the host for mounting. Define the LUN name in the same case as on the filer.
MountPath	string-scalar	Drive letter to be assigned to the virtual disk.
Initiator	string-vector	Name of iSCSI or FC initiator the host uses to connect virtual disks. You can retrieve this value from the Disk Management console.

NetApp SnapMirror agent

The NetApp SnapMirror agent is configured as a resource of type NetAppSnapMirror.

Resource type definition

```
type NetAppSnapMirror (
    static keylist SupportedActions = { fbsync }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static i18nstr ArgList[] = { FilerResName,
        "FilerResName:FilerName",
        "FilerResName:StorageIP", VolumeName,
    SnapMirrorArguments,
        SnapMirrorSchedule, AppResName }
    str FilerResName
    str VolumeName
    str SnapMirrorArguments
    str SnapMirrorSchedule
    str AppResName
)
```

Attribute definitions

Table A-7 NetApp SnapMirror agent attributes

Attribute	Type - Dimension	Description
FilerResName	string-scalar	Name of the VCS NetAppFiler-type resource in the group.
VolumeName	string-scalar	Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer.
SnapMirrorArguments	string-scalar	Specifies the SnapMirror arguments such as maximum transfer speed and restart mode.

Table A-7 NetApp SnapMirror agent attributes (Continued)

Attribute	Type - Dimension	Description
SnapMirrorSchedule	string-scalar	<p>Specifies the schedule the destination uses for updating data. Do not assign a value for this attribute if you use SnapManager.</p> <p>By default, this attribute does not have any value.</p>
AppResName	string-scalar	<p>Name of the resource configured to monitor the application being made highly available.</p>

Sample configuration

This chapter contains the following topics:

- [“About sample configurations”](#) on page 104
- [“Resource dependency graph”](#) on page 104
- [“Sample configuration file”](#) on page 105


About sample configurations

This appendix describes a typical service group configured to monitor the state of Oracle in a VCS cluster. The sample dependency graph depicts the resource types, resources, and resource dependencies within the service group. The sample configuration file (`main.cf`) is also included for your reference. Review the dependency carefully before configuring the agent. For more information about VCS resource types, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

Resource dependency graph

In the sample configuration shown in the dependency graph below, the shared virtual disk group is configured using the NetAppFiler agent. The shared volume on which Oracle shared data is installed is mounted as drive `G:\` using the NetappSnapDrive agent. The service group IP address for the server is configured using the IP and NIC resource types. The Oracle resource can be started after each of these resources is brought online.

Figure B-1 Sample Oracle service group resource dependency graph



Sample configuration file

Sample configuration files are provided for your reference.

Sample configuration file using iSCSI initiator

```
include "types.cf"

cluster merlin285-286 (
    SecureClus = 1
)
system VCSNET285 (
)
system VCSNET286 (
)
group ora-iscsi (
    SystemList = { VCSNET285 = 0, VCSNET286 = 1 }
)
IP ora-iscsi-IP (
    Address = "10.255.254.253"
    SubNetMask = "255.255.255.0"
    MACAddress @VCSNET285 = "00-11-43-DD-F3-96"
    MACAddress @VCSNET286 = "00-11-43-DD-F4-5F"
)

NIC ora-iscsi-NIC (
    MACAddress @VCSNET285 = "00-11-43-DD-F3-96"
    MACAddress @VCSNET286 = "00-11-43-DD-F4-5F"
)

NetAppFiler ora-iscsi-NetAppFiler (
    FilerName = vcsnetappl
    StorageIP = "10.255.254.252"
)

NetAppSnapDrive ora-iscsi-NetAppSnapDrive (
    FilerResName = ora-iscsi-NetAppFiler
    VolumeName = voll
    ShareName = data01
    LUN = "sril.lun"
    MountPath = M
    Initiator @VCSNET285 = {
    "iqn.1991-05.com.microsoft:vcsnet285.zen.com" }
    Initiator @VCSNET286 = {
    "iqn.1991-05.com.microsoft:vcsnet286.zen.com" }
)

Netlsnr Netlsnr-LISTENER (
    ServiceName = OracleOraDb10g_home1TNSListener
)
```

```
Oracle Oracle-MerlinA (  
    ServiceName = OracleServiceMerlinA  
    DetailMonitor = 1  
    Domain = zen  
    SID = MerlinA  
    UserName = Administrator  
    EncryptedPasswd = fnsRktHrrRorTnfLglHli  
    SQLFile = "C:\\Program Files\\Veritas\\cluster  
server\\bin\\Oracle\\check.SQL"  
)
```

Oracle-MerlinA requires ora-iscsi-NetAppSnapDrive
Oracle-MerlinA requires ora-iscsi-IP
Netlsnr-LISTENER requires Oracle-MerlinA
ora-iscsi-NetAppSnapDrive requires ora-iscsi-NetAppFiler
ora-iscsi-IP requires ora-iscsi-NIC

Sample configuration file using FC initiator

```
include "types.cf"  
  
cluster merlin277 (  
    SecureClus = 1  
)  
  
system VCSNET277 (  
)  
  
system VCSNET278 (  
)  
  
group ora-fc (  
    SystemList = { VCSNET277 = 0, VCSNET278 = 1 }  
)  
  
IP ora-fc-IP (  
    Address = "10.255.254.253"  
    SubNetMask = "255.255.255.0"  
    MACAddress @VCSNET277 = "00-11-43-33-1B-07"  
    MACAddress @VCSNET278 = "00-11-43-33-1B-60"  
)  
  
NIC ora-fc-NIC (  
    MACAddress @VCSNET277 = "00-11-43-33-1B-07"  
    MACAddress @VCSNET278 = "00-11-43-33-1B-60"  
)  
  
NetAppFiler ora-fc-NetAppFiler (  
    FilerName = vcsnetappl
```

```
StorageIP = "10.255.254.252"
)

NetAppSnapDrive ora-fc-NetAppSnapDrive (
  FilerResName = ora-fc-NetAppFiler
  VolumeName = voll
  ShareName = data01
  LUN = "sri3.lun"
  MountPath = K
  Initiator @VCSNET277 = { "21:00:00:e0:8b:19:bf:eb" }
  Initiator @VCSNET278 = { "21:00:00:e0:8b:19:d0:30" }
)

Netlsnr Netlsnr-LISTENER (
  ServiceName = OracleOraDb10g_home1TNSListener
)

Oracle Oracle-MerlinB (
  ServiceName = OracleServiceMerlinB
  DetailMonitor = 1
  Domain = zen
  SID = MerlinB
  UserName = Administrator
  EncryptedPasswd = ckpOhqEooOloQkcIdiEif
  SQLFile = "C:\\Program Files\\Veritas\\cluster
server\\bin\\Oracle\\check.SQL"
)

Oracle-MerlinB requires ora-fc-NetAppSnapDrive
Oracle-MerlinB requires ora-fc-IP
Netlsnr-LISTENER requires Oracle-MerlinB
ora-fc-NetAppSnapDrive requires ora-fc-NetAppFiler
ora-fc-IP requires ora-fc-NIC
```


Oracle Enterprise Manager 10g Grid Control with VCS

This chapter contains the following topics:

- [“About Oracle Enterprise Manager 10g Grid Control with VCS”](#) on page 110
- [“Supported software”](#) on page 110
- [“Installing the VCS agent for Oracle”](#) on page 110
- [“Installing Oracle Enterprise Manager server”](#) on page 111
- [“Installing Oracle”](#) on page 111
- [“Installing the Oracle management agent”](#) on page 111
- [“Creating and configuring Oracle database and listener on the first node”](#) on page 112
- [“Configuring Oracle on failover nodes”](#) on page 113
- [“Configuring the Oracle service group”](#) on page 113
- [“Configuring a virtual network name”](#) on page 114
- [“Configuring the listener for virtual network name”](#) on page 114
- [“Configuring the database service for virtual network name”](#) on page 115
- [“Configuring an additional Oracle management agent”](#) on page 117
- [“Adding the database to the OEM 10g Grid Control”](#) on page 118
- [“Making the additional agent highly available”](#) on page 119

About Oracle Enterprise Manager 10g Grid Control with VCS

Oracle Enterprise Manager (OEM) 10g allows users to manage Oracle components, including Oracle databases, through the Grid Control Management interface.

This appendix covers general instructions on how to install the Oracle 10g Grid Control Management Agent, and detailed step-by-step instructions on how to set up the Oracle Enterprise Manager 10g Grid Control to enable users to manage Oracle databases from the Oracle Enterprise Manager console, in a VCS cluster.

Perform the steps in the given order to configure Oracle Enterprise Manager 10g Grid Control to monitor databases in a VCS environment.

Supported software

- Oracle 10g, Release 2 (10.2.0.1.0) Standard Edition or Enterprise Edition with Windows Server 2003 (Standard Edition, Enterprise Edition, or Datacenter Edition) (SP 1 required, SP2 supported, but not required for all editions)
- Oracle Enterprise Manager 10g Grid Control Release 10.1.0.3
- Veritas Cluster Server 5.1 for Windows with the Veritas Cluster Server Database Agent for Oracle

Installing the VCS agent for Oracle

Install the VCS database agent for Oracle on all the cluster nodes.

See [Chapter 2, “Installing the VCS agent for Oracle”](#) on page 19.

Installing Oracle Enterprise Manager server

Install the Enterprise Manager server on a separate server. The Enterprise Manager server should not be a part of the cluster.

To install the Enterprise Manager server

- 1 From the Oracle Enterprise Manager CD, double-click Setup.exe to launch the Oracle Universal Installer.
- 2 On the Welcome panel, click **Next**.
- 3 On the Specify File Locations panel, specify the appropriate paths required for the installation and click **Next**.
- 4 On the Select a Product to Install panel, choose **Enterprise Manager 10g Grid Control Using a New Database 10.1.0.3.0** and click **Next**.
- 5 Follow the instructions of the installer. Refer to the Oracle product documentation for more information. At the end of the installation, make a note of the information displayed on the Summary page. You may need this information to complete the configuration.

Installing Oracle

Install Oracle on all the cluster nodes. Refer to the following topics:

- See [“Before installing and configuring Oracle”](#) on page 40
- See [“Managing shared storage”](#) on page 42
- See [“Creating a temporary service group”](#) on page 47
- See [“Bringing the temporary service group online”](#) on page 48
- See [“Installing Oracle on the first node”](#) on page 48

Do not create the Oracle database yet. We will create and configure the database later in the flow.

Installing the Oracle management agent

Install the Oracle Management Agent on all the cluster nodes under the same Oracle home path.

To install the Oracle Management Agent

- 1 From the Oracle Enterprise Manager CD, double-click Setup.exe to launch the Oracle Universal Installer.

- 2 On the Welcome panel, click **Next**.
- 3 On the Specify File Locations panel, specify the appropriate paths required for the installation and click **Next**.
- 4 On the Select a Product to Install panel, choose **Additional Management Agent 10.1.0.3.0** and click **Next**.
- 5 On the Specify Oracle Management Service Location panel, specify the host name (the Enterprise Manager Server name) and the port (default port) that the Management Agent should use to communicate with the Enterprise Management service.
- 6 The installer might display a Warning. Click **OK** to continue.
- 7 Click **Next**. Follow the instructions of the installer to complete the installation.
- 8 Repeat steps 1 to 7 on all the cluster nodes.

You must install the Oracle Additional Management Agent under the same Oracle Home path on all the cluster nodes.

Creating and configuring Oracle database and listener on the first node

Perform the following steps to create and configure Oracle database and listener on the first node in the cluster. Do not perform these steps on the failover nodes yet. Steps for failover nodes are described later in the flow.

To create and configure an Oracle database and a listener

- 1 Ensure that the volumes or LUNs created for the Oracle database, control files, log files, bdump, cdump, udump files, flash recovery area, and archive location are mounted on the node.
- 2 Bring the temporary service group online on the node. This is the service group that you created while installing Oracle earlier. This service group contains an IP resource which is the virtual IP address for the databases and the listeners.
- 3 Create a Listener service.
See “[Creating the listener service](#)” on page 51 for more instructions.
- 4 Configure the listener to use the virtual IP address.
See “[Configuring the listeners to use the virtual IP address](#)” on page 52 for instructions.

- 5 Use the Oracle Database Configuration Assistant (DBCA) to create the database.
 - Start the Database Configuration Assistant (**Start > All Programs > Oracle-Oracle_Home > Configuration and Migration Tools > Database Configuration Assistant**). This launches the database installer.
 - On the Welcome panel, click **Next**.
 - On the Operation panel, click **Create a database** and click **Next**.
 - On the Database Templates panel, click **Custom Database** and click **Next**.
 - On the Management Options panel (Step 4 of 12), choose **Use Grid Control for Database Management**.
 - Click **Next**. Follow the installer instructions and select other options as per your desired database configuration.See “[Creating the Oracle database on shared disks](#)” on page 50 for instructions.
- 6 Configure the database on the shared virtual disk. Complete all the procedures described in the topic “[Configuring the Oracle database](#)” on page 53.
- 7 Configure the Oracle and listener services.

See Follow all the steps given in the procedure with one exception. In step 2, bullet 6, choose **Local System Account** instead of **This Account**.

Configuring Oracle on failover nodes

Configure Oracle on all the failover nodes that are going to be a part of the Oracle service group.

See “[Installing and configuring Oracle on failover nodes](#)” on page 60.

Configuring the Oracle service group

Configure the Oracle service group.

See [Chapter 4, “Configuring the Oracle service group](#)” on page 63.

Configuring a virtual network name

See the *Veritas Cluster Server Administrator's Guide* for more information on how to create a resource, bring the resource online, and link a resource.

To configure a virtual network name

- 1 Create a Lanman resource with the virtual server name.
- 2 Specify the attribute IPResName. The value of this attribute should be the IP resource name of the service group.
Resource names are case-sensitive in VCS.
- 3 Create a dependency with the Lanman resource as the parent resource and the IP resource as the child resource.
- 4 Enable the Lanman resource and bring it online. This will bring the virtual network name online.
- 5 Create a dependency with the Oracle resource as the parent resource and the Lanman resource as the child resource.
- 6 Remove the dependency between the Oracle resource and the IP resource.

Configuring the listener for virtual network name

Perform these steps on all the cluster nodes. Ensure that the service group is online on the node on which you perform these steps.

To configure the listener for the virtual network name

- 1 Offline the listener service resource on the node.
- 2 In the listener.ora file, change the host name to the network name from the Lanman resource that you created.
Replace
`HOST=virtual_IP_address`
with
`HOST=network_name`
The default location of the file listener.ora is
`<ORACLE_HOME>\NETWORK\ADMIN\.`
- 3 Remove the database entries in SID_LIST, if any.
- 4 Save and close the configuration.
- 5 Bring the listener service resource online on the node.

Configuring the database service for virtual network name

Perform these steps to configure the database service for the virtual network name.

To configure the database service for virtual network name

- 1 If the file `tnsnames.ora` does not contain entries for Database Service and Network Name, then create these entries using the Net Configuration Assistant. You must create these entries on all the cluster nodes.

- 2 Associate the database with the listener.

- Ensure that the initialization parameter file contains the following entries:

- `SERVICE_NAMES` (the name of the database service)
- `INSTANCE_NAME` (the name of the database instance)

Example SQL queries to ensure these entries:

```
show parameter INSTANCE_NAME
show parameter SERVICE_NAME
```

These parameters are created during installation or database creation.

- Use one of the following ways to configure the new attribute `listener_alias`:

Run the following SQL command:

```
SQL> ALTER SYSTEM SET LOCAL_LISTENER='<listener_alias>'
scope=spfile;
```

or

Add the following entry to the initialization parameter file (pfile/spfile):

```
LOCAL_LISTENER = <listener_alias>
```

- Define the parameter `listener_alias`. If your Oracle configuration uses the file `tnsnames.ora`, edit the file as instructed below. The default location of `tnsnames.ora` is `<ORACLE_HOME>\NETWORK\ADMIN`.

Add the following to `tnsnames.ora` file:

```
<listener_alias>=
(DESCRIPTION =
(AADDRESS=(Protocol=TCP)(HOST=virtual_name)(Port=default port))
)
```

The `listener_alias` parameter gets appended by the default domain name specified in the file `sqlnet.ora`.

- Take the listener service and the database service resources offline, and then bring them online.

- 3 Edit the initialization parameter file. For the dispatchers parameter, replace the virtual IP address with virtual name for dispatchers.

If you use PFILE

Edit the dispatchers parameter by adding or editing the HOST parameter.

The edited dispatchers parameter should look like:

```
dispatchers = '(ADDRESS=(Protocol=TCP)
(HOST=virtual_name) (other entries))'
```

Here *virtual_name* is the name of the Lanman resource.

If you use SPFILE

Log on to the Oracle database as administrator and edit the dispatchers parameter by adding or editing the HOST parameter. Type the following commands on the command prompt:

```
C:\>sqlplus "/as sysdba"
```

```
C:\>alter system set dispatchers =
'(ADDRESS=(Protocol=TCP) (HOST=virtual_name) (other
entries))' scope=spfile;
```

Here *virtual_name* is the name of the Lanman resource.

Configuring an additional Oracle management agent

Configure an additional Oracle Management Agent for the virtual server. Note that this agent is in addition to the agent that you installed earlier. The agent that you installed in “[Installing the Oracle management agent](#)” on page 111 is for the physical node, and this agent is for the virtual host.

Perform these steps on all the nodes. Make sure that the Oracle service group is online on the node on which you perform these steps.

To configure an additional Oracle Management Agent

1 At the command prompt, type the following commands:

- `C:\> cd [Management Agent home]/bin`
- `C:\[Management Agent home]/bin>emctl deploy agent [-n WindowsServiceName] AGENTSTATE_dir AGENTSTATE_hostname:port DeployHostName:port`

Where,

WindowsServiceName: the new service which is used for the HA of additional management agent. (We will create a GenericService resource for this later.)

AGENTSTATE_dir: a new directory

AGENTSTATE_hostname: virtual host name

port: a new port

DeployHostName: physical FQDN of the host

port: default port

Example:



```
C:\Oracle\product\10.1.0\em_1\bin>emctl deploy agent -n  
OraEm10gOneVCSHost E:\Oracle\product\10.1.0\em_1\vcshost\  
vcshost:1831 vcs2h283.vcsnet.com:1830
```

- 2 Start the Management Agent service ("OraEm10gOneVCSHost" in the above example) for the management agent and make sure it starts.
- 3 Stop the Management Agent service on this node.
- 4 Repeat steps 1 to 3 on the next node.

Adding the database to the OEM 10g Grid Control

Perform these steps to add the database to the OEM 10g Grid Control.

To add the database to the OEM 10g Grid Control

- 1 Access the OEM Grid control using the following URL:
`http://<EM server.domainName>:7777/em`
- 2 Logon with user as sysman. The Enterprise Management Server will discover the virtual network in its target host. You have to remove the listener and the database for the physical node and add the listener and database for the virtual network name.
- 3 On the Enterprise Manager console, click **Targets** and then click **All Targets**.
- 4 Choose the database for the physical node, and click **Remove**. Click **Yes** to confirm the database removal. Remove the listener for the physical node in a similar way.
- 5 On the Enterprise Manager console, click **Targets** and then click **Databases**.
- 6 Click **Add**. The Enterprise Management Server will prompt you to select a host for which you want to add the database.
- 7 Choose the virtual network name and click **Continue**. The Enterprise Manager will discover the targets on that machine. Check if the database and listener names displayed are the ones that you wish to manage from the Enterprise Management Server.
- 8 Under Databases, click on the tools icon (under Configure)  to configure the database.
- 9 Specify the Database name and the Monitor Password. This password is later used when you login to the database to monitor it.
- 10 Click **Next** and follow the instructions to complete the database configuration. After you configure the database successfully, the tools icon will appear as  .

Making the additional agent highly available

Add a GenericService resource to make the additional agent highly available.

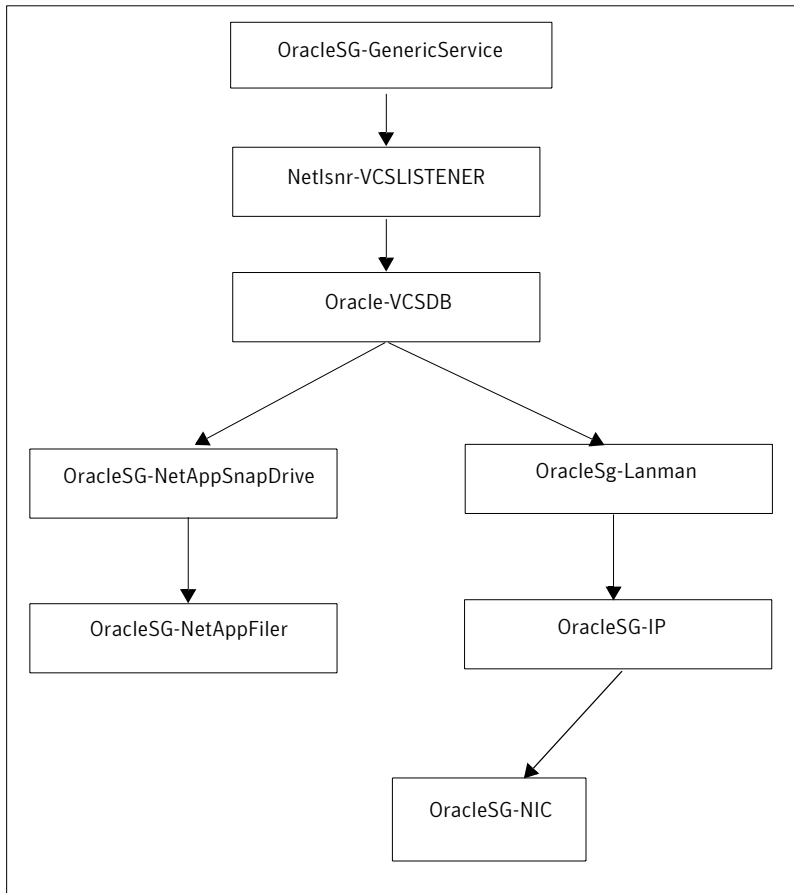
To make the additional agent highly available

- 1 Set the startup type for the additional agent service to Manual, on all the nodes. You must do this before configuring GenericService resource.
- 2 Create a GenericService resource for the Oracle Management Agent service. (This is the service for the additional agent that you created in [“Configuring an additional Oracle management agent”](#) on page 117).
- 3 Specify the attributes DelayAfterOnline and DelayAfterOffline. Set the value as 30 seconds.

Sometimes the GenericService resource may go into an unknown state. This is because the Oracle agent service takes some time in coming up. The workaround for this is to increase the attributes DelayAfterOnline and DelayAfterOffline, or probe the resource after the Windows service has started.

- 4 Create a dependency with the GenericService resource as the parent resource and the Netlsnr resource as the child resource.
- 5 Enable the GenericService resource and bring it online.

Sample Configuration



Notes:

- The sequence in which database and the additional management agent is created is not critical as long as everything is configured with the virtual server.
- The additional Oracle Management Agent service will not come online till the network name (Lanman) is online on the same node.
- The Oracle Management Agent (for the physical node) can be online on both the nodes. Only the additional Oracle Management Agent (for the virtual server) that needs to be online on one of the two nodes, as it depends on the virtual network name. This ensures that the host information is always available.

- Failover time may increase after the Management Agent is added to the VCS cluster group. After Oracle service group is failed over to the other node, the Enterprise Manager agent service should be restarted on the Enterprise Manager Server in order to connect to the database using Enterprise Manager.

Index

A

- agent operations
 - Netlsnr agent 12
 - Oracle agent 11
- attributes
 - Netlsnr agent 98
 - Oracle agent 95

C

- clusters
 - configuring (HA) 21
- configuration wizard 65
- configuration, typical setup 17
- configure
 - LLT over UDP using VCW 27
- configuring
 - cluster (HA) 21

D

- database agent
 - about 11
 - configuring using wizard 65
 - error messages 85
 - installing 20
 - troubleshooting 84
 - typical setup 17
- database services
 - configuring for Oracle 60
- database, creating 50
- DelayAfterOffline attribute
 - Netlsnr agent 98
 - Oracle agent 95
- DelayAfterOnline attribute
 - Netlsnr agent 98
 - Oracle agent 95
- dependency graph 104
- detail monitoring
 - about 11
 - configuring 80
 - disabling 82

- enabling 80

- DetailMonitor attribute 96
- Domain attribute 95

E

- EncryptedPasswd attribute 95
- error messages
 - Netlsnr agent 90
 - Oracle agent 88
- error tags 85

F

- failover 17

I

- IndepthMonitorCount attribute 96

L

- listener service
 - associating database 52
 - configuring 58
 - creating 51
- LLT over UDP
 - configuring using VCW 27
- logging
 - VCW logs 86
 - VCWsilent logs 87
- Logs
 - VCS 85
 - VCW 86
 - VCWsilent 87

M

- message logs 85
- message tags 85

N

- NetAppFiler agent

- type definition 98
- NetAppFiler agent attributes
 - definitions 99
- NetAppSnapDrive agent
 - type definition 100
- NetAppSnapDrive agent attributes
 - definitions 100
- NetAppSnapMirror agent
 - type definition 101
- NetAppSnapMirror agent attributes
 - definitions 101
- Netlsnr agent
 - attributes 98
 - error messages 90
 - operations 12
 - type definition 97
- Netlsnr agent attributes
 - definitions 98
 - DelayAfterOffline 98
 - DelayAfterOnline 98
 - ServiceName 98

O

- offlining service group 73
- operations
 - Netlsnr agent 12
 - Oracle agent 11
- Oracle
 - configuring database 60
- Oracle 9i/10g
 - associating database and listener 52
 - configuring services 58
 - creating database 50
 - creating listener service 51
 - installing and configuring 48
- Oracle agent
 - attributes 95
 - error messages 88
 - operations 11
 - type definition 94
- Oracle agent attributes
 - DelayAfterOffline 95
 - DelayAfterOnline 95
 - DetailMonitor 96
 - Domain 95
 - Encryptedpasswd 95
 - IndepthMonitorCount 96
 - ServiceName 95
 - SID 95

- SQLFile 97
 - SQLTimeOut 96
 - UserName 95
- Oracle service, configuring 58

R

- resource type
 - NetAppFiler agent 98
 - NetAppSnapDrive agent 100
 - NetAppSnapMirror agent 101
 - Netlsnr agent 97
 - Oracle agent 94

S

- Security Services
 - configuring 28
- Service group
 - creating using wizard 66
- service group
 - switching 74
 - taking offline 73
- ServiceName attribute
 - Netlsnr agent 98
 - Oracle agent 95
- SID attribute 95
- SQLFile attribute 97
- SQLTimeOut attribute 96
- switching service group 74

T

- troubleshooting information 84
- type definition
 - NetAppFiler agent 98
 - NetAppSnapDrive agent 100
 - NetAppSnapMirror agent 101
 - Netlsnr agent 97
 - Oracle agent 94

U

- UserName attribute 95

V

- VCS
 - configuring the cluster for HA 21
- VCS logs 86

W

wizard 65

