

Cluster Server 7.4.1 Implementation Guide for Oracle - Windows

Last updated: 2019-02-01

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the VCS agents for Oracle and NetApp	8
	About the VCS agents for Oracle and NetApp	8
	About the VCS agent for Oracle	9
	About the Oracle agent	9
	About the Netlsnr agent	10
	About the Info function	10
	About the Action function	11
	About the VCS hardware replication agent for NetApp	12
	About the NetApp Filer agent	13
	About the NetApp SnapDrive agent	14
	About the NetApp SnapMirror agent	15
	How the agent makes Oracle highly available	19
	Typical Oracle configuration in a VCS cluster	20
 Chapter 2	 Installing and configuring VCS	 21
	About installing the VCS agents	21
	Configuring the cluster using the Cluster Configuration Wizard	21
	Configuring notification	30
	Configuring Wide-Area Connector process for global clusters	33
 Chapter 3	 Installing and configuring Oracle	 36
	About installing Oracle	36
	Prerequisites for installing Oracle	37
	Privilege requirements	37
	Installing Oracle	37
	About creating and configuring Oracle	39
	Prerequisites for configuring Oracle	39
	Configuring Microsoft iSCSI initiator	40
	Managing storage using Windows Logical Disk Manager	40
	Reserving disks (if you use Windows LDM)	42
	Creating volumes (if you use Windows LDM)	42
	Mounting volumes (if you use Windows LDM)	43

Unassigning a drive letter	43
Releasing disks (if you use Windows LDM)	44
Managing storage using NetApp filer	44
Connecting virtual disks to the cluster node	46
Disconnecting virtual disks from the cluster nodes	47
Creating a temporary service group	47
Configuring Oracle on the first node	48
Dismounting a volume	48
Bringing the temporary service group online	48
Creating the Oracle database on shared disks	49
Creating the listener service	50
Configuring the listeners to use the virtual IP address	51
Associating the database with the listener	51
Configuring the Oracle database	52
Configuring the Oracle and listener services	58
Disconnecting virtual disks from the first node	59
Releasing disks (if you use Windows LDM)	59
Configuring Oracle on failover nodes	59

Chapter 4 Configuring the Oracle service group 62

About configuring the Oracle service group	62
Deleting the temporary service group	63
About configuring the Oracle service group using the wizard	63
Prerequisites for configuring the Oracle service group	63
Creating an Oracle service group using the wizard	64
Configuring dependent services	69
Verifying the Oracle service group configuration	69
Bringing the service group online	69
Taking the service group offline	70
Switching the service group	70
Disabling the service group	71
Reconfiguring the Oracle service group	71

Chapter 5 Administering the Oracle service group 72

About administering the Oracle service group	72
About modifying an Oracle service group	72
Prerequisites for modifying the Oracle service group	72
Modifying an Oracle service group	73
Deleting an Oracle service group	73
Configuring detail monitoring	74
Enabling detail monitoring	74
Disabling detail monitoring	75

Chapter 6	Troubleshooting	77
	About troubleshooting VCS agents	77
	Troubleshooting common problems	77
	VCS logging	79
	VCS Cluster Configuration Wizard (VCW) logs	80
	VCWsilent logs	80
	Agent error messages and descriptions	81
	Oracle agent	81
	Netlsnr agent	83
	NetApp agents	84
Appendix A	Resource type definitions	86
	About resource type definitions	86
	Oracle agent	86
	Resource type definition	86
	Attribute definitions	87
	Netlsnr agent	89
	Resource type definition	89
	Attribute definitions	90
Appendix B	Sample configuration	91
	About Oracle sample configurations	91
	Sample configuration file	91
	Sample configuration file	91
	Sample configuration file using iSCSI initiator	93
	Sample configuration file using FC initiator	94
Appendix C	Oracle Enterprise Manager 10g Grid Control with VCS	97
	About Oracle Enterprise Manager 10g Grid Control with VCS	98
	Installing the VCS agent for Oracle	98
	Installing Oracle Enterprise Manager server	98
	Installing Oracle	99
	Installing the Oracle Management Agent	99
	Creating and configuring Oracle database and listener on the first node	100
	Configuring Oracle on failover nodes	101
	Configuring the Oracle service group	101
	Configuring a virtual network name	101
	Configuring the listener for virtual network name	102

Configuring the database service for virtual network name	102
Configuring an additional Oracle management agent	104
Adding the database to the OEM 10g Grid Control	105
Making the additional agent highly available	106
Sample Configuration	106

Introducing the VCS agents for Oracle and NetApp

This chapter includes the following topics:

- [About the VCS agents for Oracle and NetApp](#)
- [About the VCS agent for Oracle](#)
- [About the VCS hardware replication agent for NetApp](#)
- [How the agent makes Oracle highly available](#)
- [Typical Oracle configuration in a VCS cluster](#)


About the VCS agents for Oracle and NetApp

The Cluster Server (VCS) agent for Oracle provides high availability to Oracle in a VCS environment. The VCS agent for 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 NetApp 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 NetApp filer via an ethernet cable supporting iSCSI or Fibre Channel (FC) as the transport protocol.

The following figure illustrates a typical VCS cluster configuration in a NetApp storage environment.

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



This chapter provides an overview of the agents. Review the information about agents, including their resource type definitions and attribute definitions.

See [“About resource type definitions”](#) on page 86.

About the VCS agent for Oracle

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

The package contains the following 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.

About the 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 functions include the following:

Online	Starts the Oracle service
Offline	Stops the Oracle service, it also stops all the services that are dependent on the Oracle service

Monitor	<p>Verifies the status of Oracle services. The agent provides two levels of application monitoring, basic and detail.</p> <ul style="list-style-type: none"> ■ Basic monitoring <p>The agent queries the Service Control Manager (SCM) to verify that Oracle services are continuously active.</p> <ul style="list-style-type: none"> ■ Detail monitoring <p>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. Review the instructions for configuring detail monitoring for an Oracle database.</p> <p>See “Configuring detail monitoring” on page 74.</p>
Clean	Forcibly stops the Oracle database
Info	<p>Provides static and dynamic information about the state of the database.</p> <p>See “About the Info function” on page 10.</p>
Action	<p>Performs predefined actions on a resource.</p> <p>See “About the Action function” on page 11.</p>

About the 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 functions include the following:

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

About the Info function

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

The function retrieves the following static information:

Version	InstanceNo
HostName	StartupTime
Parallel	Thread
InstanceRole	InstanceName

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`. Here `%VCS_Home%` is
the product installation directory, typically `C:\Program Files\Veritas\Cluster
Server`.

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.

About the Action function

The Action function 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>]
```

The following table 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 *Cluster Server Agent Developer's Guide*.

About the VCS hardware replication agent for NetApp

The VCS hardware replication agent for NetApp provides failover support and recovery in environments employing NetApp filers for storage and NetApp 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 (GCO).

The VCS agents for NetApp are as follows:

- NetAppFiler agent
- NetAppSnapDrive agent
- NetAppSnapMirror agent

About the 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.

NetApp Filer agent function

The NetApp Filer agent function is as follows:

Monitor

Performs the following tasks:

- 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.
- Opens a filer connection and checks if ONTAPI version is supported by the filer. If the connection fails or the ONTAPI version is not supported, the agent reports the state as offline.

The NetApp Filer agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppFiler agent resource type definition

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

```
type NetAppFiler (  
    static int MonitorInterval = 30  
    static il8nstr ArgList[] = { FilerName, StorageIP }  
    static str Operations = None  
    str FilerName  
    str StorageIP  
)
```

NetAppFiler agent attribute definitions

The following table describes the NetApp Filer agent attributes.

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

About the NetApp SnapDrive agent

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

NetApp SnapDrive agent functions

The NetApp SnapDrive agent functions are as follows:

Online	Connects a virtual disk (LUN) using an iSCSI or an 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 SnapDrive agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppSnapDrive agent resource type definition

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

```
type NetAppSnapDrive (  
    static int MonitorInterval = 30  
    static int NumThreads = 1  
    static il8nstr 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  
)
```

NetAppSnapDrive agent attribute definitions

The following table describes the NetApp SnapDrive agent attributes.

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

About the 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. The agent supports the replication modes supported by NetApp. The agent supports asynchronous, semi-synchronous, and synchronous modes of replication. You can set the mode of replication using the SyncMode agent attribute.

NetApp SnapMirror agent functions

The NetApp SnapMirror agent functions are as follows:

Online	<p>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.</p> <p>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.</p> <p>If the original source filer is down, the agent performs a mirror break-off to enable local write access, if the filer is not already broken off.</p> <p>If the original source returns to life, you must resynchronize the data manually. The online function touches a lock file if read-write access is enabled successfully.</p>
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 that the lock file exists. If the lock file exists, the monitor function reports the status of the resource as online. If the lock file does not exist, the monitor function reports the status of the resource as offline.
Open	<p>Removes the lock file thereby preventing potential concurrency violation if the group fails over to another node.</p> <p>Note: The agent does not remove the lock file if the agent is started after an <code>hastop -force</code> command.</p>
Clean	Removes the lock file. No filer operations are performed as taking the resource offline does not indicate a pending role swap.

Action function

Use the Action function to perform predefined actions on a resource. To perform an action on a resource, type the following command:

```
hares -action <SnapMirror_resname> <token> [-actionargs <arg1> ...]
[-sys <system>] [-clus <cluster> ]
```

The following table lists the action supported by the NetAppSnapMirror agent.

Table 1-2 Actions supported by NetAppSnapMirror agent

Token for Action	Description
fbsync	<p>Resynchronizes an original source volume with a broken-off volume. After synchronization, the original source volume becomes the target volume.</p> <p>The broken-off volume was initially the target volume, but was broken off as a result of a take over.</p>

To synchronize volumes, type the following at the command prompt:

```
hares -action SnapMirror_resname fbsync -sys node_name
```

Where, *SnapMirror_resname* represents the name of the SnapMirror resource and *node_name* represents the node on which the service group is online.

Run the action for each SnapMirror resource.

You can also add custom actions for the agents. Refer to the *Cluster Server Agent Developer's Guide* for more information.

NetApp SnapMirror agent resource type definition and attribute definitions are as follows. This information will assist you during the agent configuration.

NetAppSnapMirror agent resource type definition

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

```
type NetAppSnapMirror (
    static keylist SupportedActions = { fbsync }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static il8nstr ArgList[] = { FilerResName,
        "FilerResName:FilerName", "FilerResName:StorageIP",
        VolumeName, SnapMirrorArguments, SnapMirrorSchedule,
        AppResName, VisibilityFrequency, SyncMode }
    str FilerResName
    str VolumeName
    str SnapMirrorArguments
    str SnapMirrorSchedule
    str AppResName
    int VisibilityFrequency = 180
    str SyncMode = async
)
```

NetAppSnapMirror agent attribute definitions

The following table describes the NetApp SnapMirror agent attributes.

Attribute	Description
FilerResName	<p>Name of the VCS NetAppFiler-type resource in the group.</p> <p>Type and dimension: string-scalar</p>
VolumeName	<p>Name of the filer volume containing the virtual disk. This is the volume that is to be mounted. Define the volume name in the same case as on the filer.</p> <p>Type and dimension: string-scalar</p>
SnapMirrorArguments	<p>Specifies the SnapMirror arguments such as maximum transfer speed and restart mode.</p> <p>The format for this attribute is:</p> <p><RestartMode> MaxSpeed</p> <p>MaxSpeed is an optional parameter. The RestartMode parameter can have the following values:</p> <ul style="list-style-type: none"> ■ Always ■ Never ■ Default <p>Setting the RestartMode parameter value to default indicates schedule priority.</p> <p>Default value of RestartMode parameter is 'Default'</p> <p>Example: Always MaxSpeed</p> <p>Type and dimension: string-scalar</p>
SnapMirrorSchedule	<p>Specifies the schedule the destination uses for updating data. Do not assign a value for this attribute if you use SnapManager.</p> <p>The schedule is in the following format:</p> <p>minute hour dayofmonth dayofweek</p> <p>Each field is separated by a space.</p> <p>Refer to the NetApp documentation for more details on the rules for each of these schedule fields.</p> <p>By default, this attribute does not have any value.</p> <p>Type and dimension: string-scalar</p>

Attribute	Description
AppResName	<p>Name of the resource configured to monitor the application being made highly available.</p> <p>Type and dimension: string-scalar</p>
SyncMode	<p>Specifies the mode of replication for the mirror.</p> <p>This attribute can have the following values:</p> <ul style="list-style-type: none"> ■ async: Indicates that the mirror should be configured in the asynchronous mode. ■ semi-sync: Indicates that the mirror should be configured in the semi-synchronous mode. ■ sync: Indicates that the mirror should be configured in the synchronous mode. <p>The default is async (asynchronous) mode.</p> <p>Type and dimension: string-scalar</p>
VisibilityFrequency	<p>Specifies how often the source snapshot will be visible on the destination mirror. It controls the value of <code>visibility_interval</code> in the <code>snapmirror.conf</code> file.</p> <p>The default value is 180 seconds.</p> <p>This attribute is applicable only if the mirror is configured in synchronous or semi-synchronous mode.</p> <p>Type and dimension: string-scalar</p>

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 [“About resource type definitions”](#) on page 86.


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.

The following figure shows a typical Oracle configuration in a VCS cluster.

Figure 1-2 Typical Oracle configuration in a VCS cluster



Installing and configuring VCS

This chapter includes the following topics:

- [About installing the VCS agents](#)
- [Configuring the cluster using the Cluster Configuration Wizard](#)

About installing the VCS agents

Install Veritas InfoScale Availability on all the systems where you want to configure the application. During installation, the product installer installs the VCS agents required for making the applications highly available.

Refer to the *Veritas InfoScale Installation and Upgrade Guide* for instructions.

For the latest information about supported operating systems and software, see the Software Compatibility List at:

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

Configuring the cluster using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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 VCW to remove the node from the cluster, rename the system, and then run VCW again to add that system to the cluster.

Note the following prerequisites before you proceed:

- The required network adapters (NICs), and SCSI controllers are installed and connected to each system.

Veritas recommends the following actions for network adapters:

- Disable the ethernet auto-negotiation options on the private NICs to prevent:
 - Loss of heartbeats on the private networks
 - VCS from mistakenly declaring a system as offlineContact the NIC manufacturer for details on this process.
- Remove TCP/IP from the private NICs to lower system overhead.
- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list. When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.

- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- In case of a NetApp storage environment, the user account for the VCS Helper service must have administrative privileges on the NetApp filer.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

- 1 Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 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, do the following:
 - Clear **Specify systems and users manually**.
 - Click **Next**.
Proceed to step 8.

To specify systems and user names manually (recommended for large domains), do the following:

- Select **Specify systems and users manually**.
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. 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.

- 6 On the System Selection panel, specify the systems for the 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. Review the status and then click **Next**.

Select the system 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.

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.
- Product is either not installed or there is a version mismatch.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and then click **Next**.

- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.

Veritas Cluster Server Configuration Wizard

Cluster Details
Enter necessary details to create the new cluster

Specify the cluster name and cluster ID. If you chose to specify the systems manually, VCV does not validate the cluster ID.

Cluster Name:

Cluster ID:

Operating System:

Select the systems to create the cluster.

☒ Select all systems

Available Systems

- ☒ ROGER
- ☒ SCOOPYDU

Total number of systems selected to create the cluster : 2

Click 'Next' to continue.

VERITAS

Back Next Cancel

Specify the cluster details as follows:

- | | |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name | Type a name for the new cluster. Veritas recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | <p>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 65535.</p> <p>Note: If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique.</p> |
| Operating System | <p>From the drop-down list, select the operating system.</p> <p>All the systems in the cluster must have the same operating system and architecture.</p> |

Available Systems Select the systems that you wish to configure in the cluster.

Check the **Select all systems** check box to select all the systems simultaneously.

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.

- 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, proceed to the next step. Otherwise, proceed to step 12.

- 11** On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer using IPv4 or IPv6 network.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:



- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends 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. Veritas recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
- Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use. The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas 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. In case of IPv4, each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.

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 configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports are 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 Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details as follows:

- 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, 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 communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, click **Use VCS User Privileges** and then type a user name and password.
The wizard configures this user as a VCS Cluster Administrator. In this mode, communication between cluster nodes and clients, including Cluster Manager (Java Console), occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password.
The default user name for the VCS administrator is *admin* and the 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.
Veritas recommends that you specify a new user name and password.
- To use the single sign-on feature, click **Use Single Sign-on**.
In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.
The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

- 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 group; this group is required to set up components for 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.

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 30.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.
You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.
See [“Configuring Wide-Area Connector process for global clusters”](#) on page 33.

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 then click **Next**.

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

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

The screenshot shows the 'Notifier SNMP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window has a blue title bar and a sidebar on the left with navigation links: Domain Selection, Create Cluster, Select Components, Configure, Summary, and Finish. The main area is titled 'Notifier SNMP Configuration' with the subtitle 'Specify information about SNMP console.' Below this, it says 'Enter the name or the IP address of the SNMP console and then select the desired severity level.' There is a table with two columns: 'SNMP Console' and 'Severity Information'. The 'SNMP Console' column has a text input field with the placeholder 'Click here to change the text..'. The 'Severity Information' column has a dropdown menu. Below the table, there are instructions: 'Click on '+' button to add more consoles.' and 'Click '-' to remove a console.' with corresponding '+' and '-' buttons. There is also a text input field for 'SNMP Trap Port' with the value '162'. A note states: 'Note: SNMP console must be MIB 2.0 compliant.' and a prompt says 'Click 'Next' to continue.' At the bottom right, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a mouse cursor.

Do the following:

- 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 the + icon to add a field; click the - icon 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 then click **Next**.

Veritas Cluster Server Configuration Wizard

Notifier SMTP Configuration
Specify information about SMTP recipients.

Domain Selection

Create Cluster

Select Components

Configure

Summary

Finish

SMTP Server Name / IP

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
Click here to change the text..	Information

Click "+" to add a recipient.
Click "-" to remove a recipient.

Click "Next" to continue.

Back Next Cancel

VERITAS

Do the following:

- 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 the + icon to add fields; click the - icon to remove a field.
- 4** On the Notifier Network Card Selection panel, specify the network information and then click **Next**.

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.
 - 6 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. The GCO option configures the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication. Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.

You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

To configure the wide-area connector process for global clusters

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

If the cluster has a ClusterService group configured, you can use the IP address configured in the service group or configure a new IP address.

Do the following:

- To specify an existing IP address, select **Use existing IP resource** and then select the IP address from the drop-down list.
- To use a new IP address, do the following:
 - In case of IPv4, select **IPv4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - In case of IPv6, select **IPv6** and select the IPv6 network from the drop-down list.
 The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.
 The IPv6 option is disabled if the network does not support IPv6.
- 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 WAC resources online when VCS starts and then click **Configure**.
- 3** Click **Finish** to exit the wizard.

Installing and configuring Oracle

This chapter includes the following topics:

- [About installing Oracle](#)
- [Prerequisites for installing Oracle](#)
- [Installing Oracle](#)
- [About creating and configuring Oracle](#)
- [Prerequisites for configuring Oracle](#)
- [Configuring Microsoft iSCSI initiator](#)
- [Managing storage using Windows Logical Disk Manager](#)
- [Managing storage using NetApp filer](#)
- [Creating a temporary service group](#)
- [Configuring Oracle on the first node](#)
- [Configuring Oracle on failover nodes](#)

About installing Oracle

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. This involves installing the Oracle binaries locally on each system.

You can perform this installation simultaneously on cluster nodes.

See the Oracle documentation for specific instructions on how to install Oracle.

Prerequisites for installing Oracle

Note the following prerequisites for installing Oracle in a VCS environment:

- Verify all nodes have adequate resources to run Oracle and VCS.
- Verify that the InfoScale product is installed on all nodes in the cluster. Refer to the *Veritas InfoScale 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 using the Cluster Configuration Wizard”](#) on page 21.

Privilege requirements

Before installing Oracle, make sure you have the following privileges:

- Log on as Domain Administrator to install Oracle binaries.
- You must be a member of the local Administrators group on all nodes where you wish to install Oracle.
- 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.
- You must have Domain Administrator privileges or belong to the Account Operators group, if you plan to create a new user account for VCS Helper service. 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.

Installing Oracle

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, note the following:

- 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 Oracle must be the same on all nodes.

For example, the database name, file locations, database credentials must be the same on all 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 Insert the Oracle software disc in a drive connected to the node and run the Setup program.

- 2 Complete the instructions given below, depending on the version of Oracle.

For Oracle 9i version, follow these steps:

- 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 and 11g version, follow these steps:

- Install 10g or 11g database.
- Choose the **Basic** or **Advanced** installation type.

If you choose Basic installation, follow these steps:

- 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**.

- Clear **Create Starter Database** and click **Next**.
If you choose Advanced installation, follow these steps:
 - Select the destination directory for the Oracle installation and click **Next**.
 - Select **Enterprise Edition** and click **Next**.
 - Select **Do not create a starter database**.
- 3 Follow the wizard instructions, select options as per your configuration and complete the installation.

About creating and configuring Oracle

Configuring the database involves setting parameters in Oracle configuration files. It also involves,

- setting the same virtual IP address for all the databases.
- configuring the databases to use the OS authentication model of Oracle.
- creating the service groups in VCS cluster.

Note that you must first configure the databases on the first cluster node and then subsequently configure them on the failover nodes.

Prerequisites for configuring Oracle

Ensure that you meet the following prerequisites before you configure Oracle.

- If using iSCSI, verify that the Microsoft iSCSI Initiator is configured to establish a persistent connection between the NetApp 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.
- Verify that the volumes or LUNs are created on the shared disk.
See [“Managing storage using Windows Logical Disk Manager”](#) on page 40.
See [“Managing storage using NetApp filer”](#) on page 44.
- Verify that the volumes 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 Oracle installer.
- Create a temporary service group to configure an IP (in case of IPv4) or an IPv6 (in case of IPv6) resource.

Configuring Microsoft iSCSI initiator

The Microsoft iSCSI initiator enables communication between Windows systems and NetApp 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 on a Windows Server system

- 1 Start the Microsoft iSCSI initiator.
- 2 On the Discovery tab, click **Discover Portal**.
- 3 On the Discover Target Portal dialog box, specify the DNS name for the NetApp filer and then click **OK**.
- 4 On the Target tab, click **Connect**.
- 5 On the Connect to Target dialog box, clear the **Add this connection to list of Favorite Targets** check box and then click **Ok**.
- 6 On the Targets tab, verify that the newly added portal is listed under the **Select a target** box and the status appears as 'connected.' Then, click **OK**.

Managing storage using Windows Logical Disk Manager

If your configuration uses shared disks and volumes that are managed using Windows Logical Disk Manager (LDM), use the VCS Mount and DiskReservation (DiskRes) agents. If you use LDM to manage non-shared local storage, use the VCS Mount and NativeDisks agents.

Before configuring the storage, review the resource types and attribute definitions of these VCS storage agents (Mount, DiskRes, NativeDisks) described in the *Cluster Server Bundled Agents Reference Guide*.

The following restrictions apply for storage managed using LDM:

- Mount, DiskRes, and NativeDisks agents are supported on VCS for Windows only. These agents are not supported if the storage is managed using Storage Foundation (SFW).
- If you are using shared storage, your storage devices must be configured to use SCSI-2 disk reservations. SCSI-3 is not supported.
SCSI support is not required if you are using non-shared storage.
- LDM support is not applicable for Disaster Recovery configurations. Currently only HA configurations are supported.

The VCS agent for Oracle requires that you create volumes 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 as that of the redo logs.

Veritas recommends that you create separate volumes for the following:

- 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 are using a shared storage configuration, ensure that these volumes are created on shared storage and are accessible from all cluster nodes.

If you are using a non-shared storage configuration, create these volumes separately on the local storage attached to each cluster node.

Perform the following tasks to configure your storage:

- Reserve disks
See [“Reserving disks \(if you use Windows LDM\)”](#) on page 42.
- Create volumes
See [“Creating volumes \(if you use Windows LDM\)”](#) on page 42.
- Mount volumes
See [“Mounting volumes \(if you use Windows LDM\)”](#) on page 43.
- Unassign the drive letter
See [“Unassigning a drive letter”](#) on page 43.

- Release the disks
See [“Releasing disks \(if you use Windows LDM\)”](#) on page 44.

Reserving disks (if you use Windows LDM)

Complete the following steps to reserve the disks on the node on which you are going to perform the application installation.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To reserve the disks

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

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

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 during installation and while configuring the service group, on additional nodes in the cluster.

Creating volumes (if you use Windows LDM)

Perform the following steps to create volumes.

To create volumes

- 1 Use the Windows Disk Management tool to verify that the disks are visible on the cluster nodes, and then create volumes on the disks.
- 2 In case of shared storage, after creating the required volumes on a node, release the reserved disks from that node.
See [“Releasing disks \(if you use Windows LDM\)”](#) on page 44.
- 3 In case of shared storage, 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 (if you use Windows LDM)

Perform the following steps to mount volumes on a cluster node.

To mount a volume

- 1 Use the Windows Disk Management tool to mount the volumes that you created earlier.
- 2 After mounting the volumes on a cluster node, run the CHKDSK command and verify that there are no errors on the mounted volumes.
- 3 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.

Unassigning a drive letter

In case of a shared storage configuration, 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.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

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.

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

To unassign drive letter

- 1 Log in as Administrator.
- 2 Open Disk Management. Type the following at the command prompt:

```
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**.

Releasing disks (if you use Windows LDM)

Perform the following steps to release reserved disks from a cluster node.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To release disks

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

```
C:\>havol -scsittest /l
```

Make a note of the disk numbers (Disk# column in the table) of the disk that you wish to release. You will need it in the next step.

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

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

For example, to release disk 4, type:

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

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


Managing storage using NetApp filer

NetApp manages data by creating volumes on physical disks. These volumes can further be divided into LUNs (Logical Unit Numbers). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and NetApp SnapDrive installed. However, if you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that filer is connected to the nodes and the LUNs are shared between all the cluster nodes.

Refer to the NetApp documentation for more information.

The following figure illustrates a typical VCS cluster in a NetApp storage environment.

Figure 3-1 VCS cluster in a NetApp storage environment



The VCS agent for Oracle requires that you create LUNs 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 LUN as that of the redo logs.

Veritas recommends that you create separate LUNs (virtual disks) for the following:

- **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)).

These LUNs must be accessible from all cluster nodes.

Perform the following tasks to create LUNs on the NetApp filer and to make them accessible from cluster nodes:

- Add the filer storage system to the SnapDrive Storage System Management snap-in on the cluster nodes.
- Create volumes on the NetApp filer.
- Share the volumes.
- Create LUNs or virtual disks on the shared volumes.
Refer to NetApp documentation for instructions on performing these tasks.

Connecting virtual disks to the cluster node

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

To connect virtual disks to the cluster node

- 1 On the cluster node where you want to connect the LUN, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen.
- 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

Perform the following steps to disconnect the virtual disks from a cluster node.

To disconnect virtual disks

- 1 On the cluster node where you want to disconnect the LUNs, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen .
- 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**.

Creating a temporary service group

Create a temporary service group with an IP (in case of IPv4) or an IPv6 (in case of IPv6) 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** (in case of IPv4) or **IPv6** (in case of IPv6) from the Resource Type drop-down list.

- Specify appropriate values for the **Address**, **SubNetMask** (in case of IPv4) or **Prefix** (in case of IPv6), 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**.

Configuring Oracle on the first node

Perform the following tasks on the first node in the cluster. This involves installing the control files, data files, log files, bdump, cdump, udump files, the flash recovery area and the archive location, on the shared storage.

Dismounting a volume

Dismount the volumes only if you plan to install Oracle on more nodes. If it's the last node for the Oracle installation, do not dismount the volumes. Dismounting a volume involves removing the drive letter or the folder path and deporting the disk group.

To dismount a volume

- 1 Stop the database and listener services.
- 2 Launch the VEA console from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 3 From the VEA console's tree view, right-click on the volume, and then click **File System > Change Drive Letter and Path**.
- 4 In the Drive Letter and Paths dialog box, select **Remove**.
- 5 Click **OK**.
- 6 From the VEA console's tree view, right-click the disk and then click **Deport Dynamic Group**.

Bringing the temporary service group online

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

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 where you are installing Oracle first.
- 2 Click **Yes** on the prompt that confirms whether you wish to bring the service group online.

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 from **Start > All Programs > Oracle-Oracle_Home > Configuration and Migration Tools > Database Configuration Assistant** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen. 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 instructions given below depending on the Oracle version.
For Oracle 9i version, follow these steps:
 - 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 version, follow these steps:

- 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 The wizard instructions help you to 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.
Follow these wizard instructions:
 - 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 and do the following:
 - 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 (in case of IPv4) or IPv6 (in case of IPv6) 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 Windows Services console.

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.

Prerequisites to associate the databases with the listeners

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 (in case of IPv4) or IPv6 (in case of IPv6) 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.

To configure the PFILE follow these steps:

- 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*. Veritas recommends setting the attribute to NONE so that the database uses the operating system 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. See [“Modifying the PFILE \(Oracle 9i only\)”](#) on page 54. Now update the SPFILE. See [“Modifying the SPFILE”](#) on page 54.
- If you use an SPFILE on Oracle 10g, review the procedure to modify the SPFILE. See [“Modifying the SPFILE”](#) on page 54.

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 disk.

Modifying the SPFILE

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 (in case of IPv4) or IPv6 (in case of IPv6) 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;
```

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

- 3 Restart the database service.

Moving the initialization file to shared disks (Oracle 10g and 11g)

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

To move the initialization files to the shared disk, use the following steps:

- 1 Verify which of the following initialization files exist in your configuration.

The default location of the following 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 Windows Services console.
- 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 Windows Services console.

Verifying the initialization file settings

To verify the initialization file settings, if a PFILE is used

- 1 Open the Registry Editor.

Press **Window+R** on the desktop to open the Run dialog box. In the Open field, type `regedit` and click **OK**.
- 2 Navigate to the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\HOME_ID\` and double-click the `ORA_SID_PFILE` registry key.

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.

To verify the initialization file settings, 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** console from **Start > All Programs > Administrative Tools > Services** or, on Windows Server 2012 operating systems from the **Start** screen, click **Administrative Tools**.
- 2 Set the startup type of the Oracle database service to manual.
Configure its Log On properties as follows:
 - Double-click the Oracle database service. In the SCM, Oracle database services appear as OracleService*SID*, 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.
Complete the following tasks:
 - In the Services applet, double-click the listener service. Listener services appear as OracleOra_HomeTNSListenerName, 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 47.

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 43.

See [“Releasing disks \(if you use Windows LDM\)”](#) on page 44.

Configuring Oracle on failover nodes

Complete the following steps to configure Oracle on failover nodes.

To configure Oracle on failover nodes

- 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 storage using Windows Logical Disk Manager”](#) on page 40.

See [“Managing storage using NetApp filer”](#) on page 44.

- 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 srvctl,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 Windows Services console, 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 includes the following topics:

- [About configuring the Oracle service group](#)
- [Deleting the temporary service group](#)
- [About configuring the Oracle service group using the wizard](#)
- [Verifying the Oracle service group configuration](#)
- [Reconfiguring the Oracle service group](#)

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.

Note: In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's `AdditionalDNSServers` attribute.

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.

About 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.

Review the instructions to modify an existing service group.

See [“About modifying an Oracle service group”](#) on page 72.

Before configuring the Oracle agent, review the resource type and the attribute definitions of the Oracle agents.

See [“About resource type definitions”](#) on page 86.

Review sample configuration file and resource dependency graph of the Oracle service group.

See [“About Oracle sample configurations”](#) on page 91.

Prerequisites for configuring the Oracle service group

Prerequisites for configuring the Oracle service group are as follows:

- If you have configured a firewall, add the required services and ports to the Firewall Exception list.

For a detailed list of the required services and ports, see the *Veritas InfoScale Installation and Upgrade Guide*.

- Verify that the appropriate InfoScale product is installed on all cluster nodes.
- Verify a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
 See [“Configuring the cluster using the Cluster Configuration Wizard”](#) on page 21.
- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be logged on as 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.
- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved in the service group configuration.
 After you modify the service group configuration you can again edit the resource attributes to set the desired value.
- Keep 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

Creating an Oracle service group using the wizard

This section describes how to create an Oracle service group.

To create an Oracle service group

- Start the Oracle configuration wizard.

In case of VCS, Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu.

In case of SFW HA, click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu to start the Solutions Configuration Center (SCC). In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.
- On the Welcome panel click **Next**.
- On the Wizard Options panel, select **Create service group** and click **Next**.
- On the Service Group Configuration panel, use the following fields as appropriate:

Service Group Name	Type a name for the Oracle service group.
Group System List	<p>In the Available Cluster Systems list, 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 list.</p> <p>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.</p> <ul style="list-style-type: none"> 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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox. For information about the AutoStartList attribute, see the <i>Cluster Server Administrator's Guide</i> .

Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

- On 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**.

Enter the appropriate information in the following fields:

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 <code>%vcs_home%\bin\Oracle\</code> . Here, <code>%vcs_home%</code> is the installation directory, typically <code>C:\Program Files\Veritas\Cluster Server</code> .

- 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.

- 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.

Do the following:

- In case of IPv4, select **IPv4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- In case of IPv6, select **IPv6** and select the IPv6 network from the drop-down list.
The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.
The IPv6 option is disabled if the network does not support IPv6.
- For each system in the cluster, select the public network adapter name.
The Adapter Display Name field displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. Verify that you select the adapters assigned to the public network, not the private.

Note: If you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that have the same MAC address, then you must edit the “MACAddress” attribute of the NIC agent and the IP agent, after you configure the application service group.

See [“About the modifications required for tagged VLAN or teamed network”](#) on page 68.

10 Review the configuration on the Summary panel.

Resources	<p>Lists the configured resources. The wizard assigns unique names to the resources.</p> <p>Click on a resource to view its attributes and their configured values in the Attributes box.</p>
Attributes	<p>Enables you to edit a resource name. Click the resource name or press the F2 key. After editing, press the Esc key to cancel the changes, or press the Enter key to confirm the changes.</p>

Click **Next**.

- 11 On 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.

About the modifications required for tagged VLAN or teamed network

Perform this task only if you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that share the same MAC address.

After you configure the application service group, you must edit the “MACAddress” attribute of the VCS NIC agent and the IP agent.

During the application service group configuration, you are required to select a network adapter for each cluster system and specify the virtual IP address for the virtual server. The application configuration wizard internally retrieves the MAC address of the specified interface and the MAC address of the interface to which the specified IP address is assigned. It then sets these MAC Addresses as the value of the “MACAddress” attribute of the VCS NIC and IP agent respectively.

If the selected interface or the interface to which the specified IP is assigned shares the MAC address with other logical interfaces, then the following issues may occur:

- NIC agent may begin to monitor an interface other than the one selected.
- The IP agent may assign the specified virtual IP address or the virtual server name to an interface other than the one selected. As a result, the IP agent may monitor an IP address other than the one specified.

As a workaround, use the VCS Java Console to edit the “MACAddress” attribute and specify its value as the interface name instead of the MAC address. You must enter the interface name in double quotes. For example, MACAddress = “InterfaceName”

Notes:

- After you specify the interface name as the “MACAddress” attribute value, if you want to use the VCS wizards to modify any settings, then you must first reset the value of the “MACAddress” attribute to the MAC address of the interface. Failing this, the VCS wizard may fail to identify and populate the selected interface. Use the VCS Java Console to edit the attribute values.
- If you change the interface name, you must update the “MACAddress” attribute value to specify the new name. Failing this, the NIC resource will go in an UNKNOWN state.
- While editing the “MACAddress” attribute to specify the interface name, you must specify the name of only one interface.

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 *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**.

Reconfiguring the Oracle service group

To re-configure the Oracle service group, delete the existing service group and recreate it using the Oracle configuration wizard.

While deleting the service group from wizard, the database and listener processes are stopped. Also, the IP resource goes offline and thus the virtual IP becomes unavailable. Before running the Oracle service group configuration wizard, the listener and database processes should be running. To start the listener process in virtual context, create a temporary service group with an IP resource, bring it online, and start the Listener process from services.msc. Once the process is started, delete the temporary service group without taking it offline and proceed with the Oracle service group configuration wizard.

Administering the Oracle service group

This chapter includes the following topics:

- [About administering the Oracle service group](#)
- [About modifying an Oracle service group](#)
- [Deleting an Oracle service group](#)
- [Configuring detail monitoring](#)

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.

About modifying an Oracle service group

The section below describe how to modify the configuration of the service groups using the configuration wizard.

Prerequisites for modifying the Oracle service group

Note the following prerequisites 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.
- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved in the service group configuration. After you modify the service group configuration you can again edit the resource attributes to set the desired value.
- If you have upgraded the product from the previous version to 7.4.1, then do not use the modify mode. Delete the service group and re-configure it using the Oracle configuration wizard.

Modifying an Oracle service group

To modify an Oracle service group

- 1 Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Review the prerequisites and click **Next**.
- 3 On 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 [“About configuring the Oracle service group using the wizard”](#) on page 63.

Deleting an Oracle service group

The section below describes how to delete an Oracle service group using the configuration wizard.

To delete an Oracle service group

- 1 Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Review the prerequisites and click **Next**.
- 3 On the Wizard Options panel, click **Delete service group**, select the service group to delete, and click **Next**.
- 4 On the Service Group Summary panel, click **Next**.
- 5 On the dialog box 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.

Veritas 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 section below describes how to enable detail monitoring using the configuration wizard.

To enable detail monitoring

- 1 Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Review the prerequisites and click **Next**.
- 3 On 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 On the Oracle Configuration panel, select the SIDs to be configured along with their respective listeners and click **Next**.
- 5 On the Detail Monitoring dialog box, specify information to enable detail monitoring, and click **Next**.

Enter the appropriate information in the following fields:

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, <code>check.sql</code> , is located at <code>%VCS_HOME%\bin\Oracle\.</code>

- 6 On 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 [“About configuring the Oracle service group using the wizard”](#) on page 63.

Disabling detail monitoring

The section below describes how to disable detail monitoring using the Oracle Configuration Wizard.

To disable detail monitoring

- 1** Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2** Review the prerequisites and click **Next**.
- 3** On 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** On the Oracle Configuration panel, click **Next**.
- 5** On 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 [“About configuring the Oracle service group using the wizard”](#) on page 63.

Troubleshooting

This chapter includes the following topics:

- [About troubleshooting VCS agents](#)
- [Troubleshooting common problems](#)
- [VCS logging](#)
- [Agent error messages and descriptions](#)

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. The following table 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.	Make sure the Oracle database parameter <code>startmode</code> is set to <code>auto</code> .
	Make sure the path of the initialization file is set correctly.
	Verify access to the shared virtual disk.

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

Error	Description
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 SQLTimeOut. 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 57.</p>
Listener service does not start after configuring the virtual IP address.	<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. Complete the following steps in the order given:</p> <ul style="list-style-type: none"> ■ Stop the listener services on the node. ■ Rename the listener.ora file. ■ Recreate the listeners using the Net Configuration Assistant. ■ Bring up the virtual IP on the node. ■ Configure the virtual IP address for the newly created listeners using the Net Manager. ■ Restart the listener services. ■ 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. The letter A indicates the first log file, B indicates the second, C indicates 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 (YYYY/MM/DD) | Mnemonic | Severity | UMI | Agent Type | Resource Name | Entry Point | Message Text.*

The following table describes the agent log message components and their descriptions.

Table 6-2 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Veritas product logs the message. For Cluster Server, the mnemonic is 'VCS'.
Severity	<p>Denotes the severity of the message. Severity is classified into the following types:</p> <ul style="list-style-type: none"> ■ 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. <p>Among these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.</p>
UMI or Unique Message ID	<p>UMI 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.</p> <p>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.</p>
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:

```
2010/05/10 09:32:44 VCS NOTICE V-16-20046-5 Oracle:
    Oracle-DATA1:online:Service status is Running
```

VCS Cluster Configuration Wizard (VCW) logs

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

Here, `%allusersprofile%` is the file system directory containing application data for all users. A typical path is `C:\ProgramData\`.

The format of the wizard log is of the format *ThreadID | Message Text*.

ThreadID is the ID of the thread initiated by the wizard and Message Text is the actual message generated by the wizard.

A typical wizard log resembles the following:

```
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 the following:

```
00005540-00000064: 5540: STARTING - Discovering NICs on the
selected machines...
00009956-00000064: 9956: STARTING - Generating private network
related files...
00009956-00000048: 9956: COMPLETED - Generating LLT host
files...
00009956-00000048: 9956: COMPLETED - Generating GAB tab files...
```



```
00009956-000000048: 9956: COMPLETED - Generating main.cf file...
00009956-000000064: 9956: STARTING - Configuring LLT on all the
nodes.
00009956-000000048: 9956: COMPLETED - Configuring LLT on all the
nodes.
```

Agent error messages and descriptions

Each ERROR or WARNING message has a description and a recommended solution.

Oracle agent

The following table lists the description of error messages and recommended solution for the Oracle agent.

Table 6-3 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>

Table 6-3 Oracle 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. 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>
Wait for mutex failed.	<p>The wait function for mutex failed. This message is followed by an error code. Check the error code for more information.</p>
Wait for mutex abandoned.	<p>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.</p> <p>Solution: Verify that the SQL script executed correctly.</p>
Cannot wait for mutex object.	<p>The agent could not acquire mutex for the resource. This may occur if two or more Oracle resources configured.</p> <p>Solution: Verify that the SQL script executed correctly.</p>
Failed to decrypt password.	<p>The agent could not decrypt the domain password.</p> <p>Solution: Verify that the password was encrypted using the VCS Encrypt utility.</p>
No domain specified in argument list.	<p>The attribute Domain is not configured for the resource.</p> <p>Solution: Verify that the correct Domain is defined in the configuration file.</p>
No username specified in argument list.	<p>The attribute UserName is not configured for the resource.</p> <p>Solution: Verify that the correct UserName is defined in the configuration file.</p>
No password specified in argument list.	<p>The attribute Password is not configured for the resource.</p> <p>Solution: Verify that the password is encrypted using the VCS Encrypt utility. Verify that the correct encrypted password is defined in the configuration file.</p>

Table 6-3 Oracle agent error messages (*continued*)

Message	Description
No SID specified in argument list.	<p>The configuration definition of the agent does not contain a value for the attribute SID.</p> <p>Solution: Verify that the attribute SID is correctly defined in the configuration file.</p>
SQL*Plus process exited with error error_code.	<p>The agent could not execute the SQL*Plus process correctly. See the associated error code for more information.</p>
Invalid filename in argument list file_name.	<p>The SQL script for the resource is incorrectly defined in the filename.</p> <p>Solution: Verify that the attribute SQLFile is correctly defined in the configuration file. Verify the validity of the file name and path.</p>
An Oracle error occurred.	<p>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.</p>

Netlsnr agent

The following table lists the description of error messages and recommended solutions for the Netlsnr agent.

Table 6-4 Netlsnr 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 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>

Table 6-4 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

The following table lists the description of error messages and the recommended solutions for the NetApp agents.

Table 6-5 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>

Table 6-5 NetApp agents error messages (*continued*)

Message	Description
Volume %s on filer %s is not a SnapMirror replicated volume	Verify replication is set up on the specified volume.
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	There should be only one destination per source volume.
Initialize VLibNetAppHost:Initialize() failed. (error_type: %s, error_code: 0x%s)	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.
Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s. error_message: %s)	This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration: <ul style="list-style-type: none"> ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP Verify the configuration definition of the resource. Make sure each attribute is defined correctly.
Unable to create/delete online lock file %s. Error code %s,	Make sure you have write permissions on the specified directory.

Resource type definitions

This appendix includes the following topics:

- [About resource type definitions](#)
- [Oracle agent](#)
- [Netlsnr agent](#)

About resource type definitions

This appendix lists the resource type definitions and attribute definitions of the 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

The following section describes the resource type definition of Oracle agent:

```
type Oracle (
    static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }
    static i18nstr IMFRegList[] = { ServiceName }
    static keylist SupportedActions = { DBRestrict, DBUndoRestrict,
    DBResume, DBSuspend, DBQuiesce, DBUnQuiesce }
    static i18nstr ArgList[] = { ServiceName, DelayAfterOnline,
```

```
DelayAfterOffline, DetailMonitor, IndepthMonitorCount, Domain, SID,  
UserName, EncryptedPasswd, SQLFile, SQLTimeOut }  
str ServiceName  
int DelayAfterOnline = 20  
int DelayAfterOffline = 20  
boolean DetailMonitor = 0  
int IndepthMonitorCount = 5  
str Domain  
str SID  
str UserName  
str EncryptedPasswd  
str SQLFile  
int SQLTimeOut = 20  
)
```

Attribute definitions

The following table describes the definitions of required attribute of Oracle agent.

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 OracleServiceSID, where SID 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.

The following table describes the definitions of optional attributes of Oracle agent.

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
DetailMonitor	boolean-scalar	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. Default = 0 If the DetailMonitor flag is set to 1, the following attributes must be defined: <ul style="list-style-type: none">■ SID■ SQLFile
IndepthMonitorCount	integer-scalar	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. Default = 5 Note: Do not set the value of this attribute to 0.
SQLTimeOut	integer-scalar	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. Default = 20

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 %VCS_HOME%\bin\Oracle\. 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

The following section describes the resource type definition of the Netlsnr agent:

```
type Netlsnr (
    static int IMF{} = { Mode=3, MonitorFreq=5,
    RegisterRetryLimit=3 }
    static i18nstr IMFRegList[] = { ServiceName }
    static i18nstr ArgList[] = { ServiceName, DelayAfterOnline,
    DelayAfterOffline }
    str ServiceName
    int DelayAfterOnline = 20
    int DelayAfterOffline = 20
)
```

Attribute definitions

The following table describes the definition of required attributes of Netlsnr agent.

Table A-3 Netlsnr agent - required attributes

Required Attribute	Type and Dimension	Definition
ServiceName	string-scalar	<p>The name of the configured listener service.</p> <p>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.</p>

The following table describes the definitions of optional attributes of Netlsnr agent.

Table A-4 Netlsnr agent - optional attributes

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

Sample configuration

This appendix includes the following topics:

- [About Oracle sample configurations](#)
- [Sample configuration file](#)

About Oracle 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.

See the *Cluster Server Bundled Agents Reference Guide*.

Sample configuration file

Sample configuration files are provided in the following section for your reference:

Sample configuration file

```
include "types.cf"

cluster cluskj (
    UserNames = { a = annH }
    Administrators = { a }
)

system KRN-EX1 (
)
```

```
group sgkj (
  SystemList = { KRN-EX1 = 0 }
  AutoStartList = { KRN-EX1 }
)

IP sgkj-IP (
  Address = "10.217.57.32"
  SubNetMask = "255.255.252.0"
  MACAddress @KRN-EX1 = 00-50-56-8D-56-AA
)

MountV sgkj-MountV (
  MountPath = "K:"
  VolumeName = OraVol
  VMDGResName = sgkj-VMDg
)

NIC sgkj-NIC (
  MACAddress @KRN-EX1 = 00-50-56-8D-56-AA
)

Netlsnr Netlsnr-LISTENER (
  ServiceName = OracleOraDb11g_home1TNSListener
)

Oracle Oracle-DB1 (
  ServiceName = OracleServiceDB1
  Domain = "vcsnetapp.com"
  SID = DB1
  UserName = administrator
  EncryptedPasswd = FTLrITiNF1GLhLi1J
)

VMDg sgkj-VMDg (
  DiskGroupName = OraDG
  DGGuid = 2c9d16a8-cbf1-452a-a0d9-e05abfa1912e
)

Oracle-DB1 requires Netlsnr-LISTENER
Oracle-DB1 requires sgkj-MountV
Netlsnr-LISTENER requires sgkj-IP
```

```
sgkj-MountV requires sgkj-VMDg
sgkj-IP requires sgkj-NIC
```

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 = {
```

```
"ign.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
Netlsnr-LISTENER requires ora-iscsi-IP
Oracle-MerlinA requires Netlsnr-LISTENER
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 = vcsnetapp1
  StorageIP = "10.255.254.252"
)

NetAppSnapDrive ora-fc-NetAppSnapDrive (
  FilerResName = ora-fc-NetAppFiler
  VolumeName = vol1
  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 appendix includes the following topics:

- [About Oracle Enterprise Manager 10g Grid Control with VCS](#)
- [Installing the VCS agent for Oracle](#)
- [Installing Oracle Enterprise Manager server](#)
- [Installing Oracle](#)
- [Installing the Oracle Management Agent](#)
- [Creating and configuring Oracle database and listener on the first node](#)
- [Configuring Oracle on failover nodes](#)
- [Configuring the Oracle service group](#)
- [Configuring a virtual network name](#)
- [Configuring the listener for virtual network name](#)
- [Configuring the database service for virtual network name](#)
- [Configuring an additional Oracle management agent](#)
- [Adding the database to the OEM 10g Grid Control](#)
- [Making the additional agent highly available](#)

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.

Installing the VCS agent for Oracle

Install the required InfoScale product (InfoScale Availability or InfoScale Enterprise) on all the systems that comprise the cluster. The VCS database agent for Oracle is installed as part of the product installation.

For details, see the *Veritas InfoScale Installation and Upgrade Guide*.

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.

Review the following topics:

- See [“Prerequisites for installing Oracle”](#) on page 37.
 - See [“Managing storage using Windows Logical Disk Manager”](#) on page 40.
See [“Managing storage using NetApp filer”](#) on page 44.
 - See [“Creating a temporary service group ”](#) on page 47.
 - See [“Bringing the temporary service group online”](#) on page 48.
 - See [“Installing Oracle ”](#) on page 37.
- 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 50.
- 4 Configure the listener to use the virtual IP address.
 See [“Configuring the listeners to use the virtual IP address”](#) on page 51.
- 5 Use the Oracle Database Configuration Assistant (DBCA) to create the database.

Complete the following steps to create the database:

- Start the Database Configuration Assistant from **Start > All Programs > Oracle-Oracle_Home > Configuration and Migration Tools > Database Configuration Assistant** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen. 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 49.

- 6 Configure the database on the shared virtual disk. Review the procedure.
See [“Configuring the Oracle database”](#) on page 52.
- 7 Configure the Oracle and listener services.
Follow all the steps given in the procedure with one exception. In step 2, bullet 6, choose **Local System Account** instead of **This Account**.
See [“Configuring the Oracle and listener services”](#) on page 58.

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 [“Configuring Oracle on failover nodes ”](#) on page 59.

Configuring the Oracle service group

Configure the Oracle service group.

See [“About configuring the Oracle service group ”](#) on page 62.

Configuring a virtual network name

Review the information on how to create a resource, bring it online and link a resource.

See the *Cluster Server Administrator’s Guide*.

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 Volume Replicator.
- 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 the steps given below 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 the steps given below 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 =
  (ADDRESS=(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 Oracle management agent that you installed in the earlier section 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 vcsw2h283.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 the steps given below 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.

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 Oracle management agent that you created earlier.
- 3 Specify the attributes DelayAfterOnline and DelayAfterOffline. Set the value as 30 seconds.

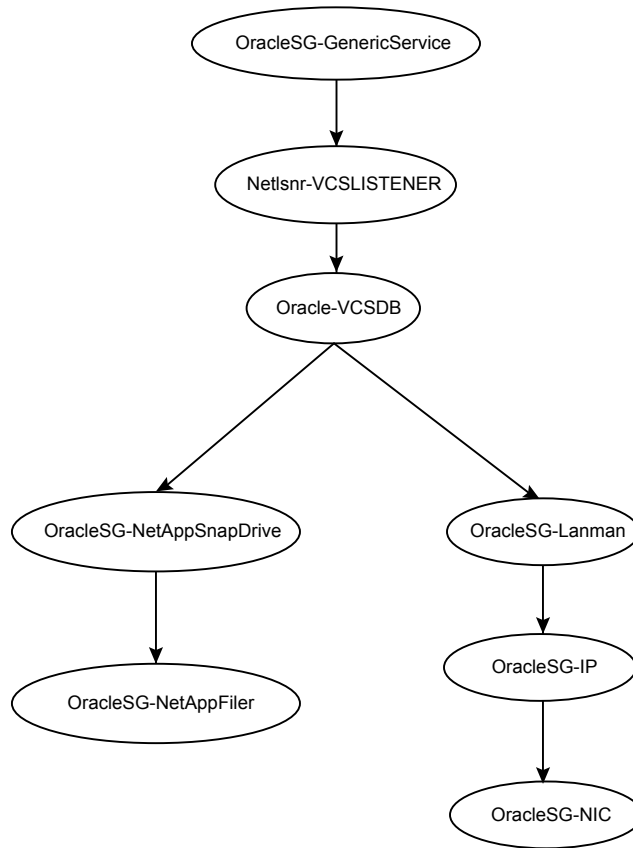
Sometimes the GenericService resource may go into an unknown state. This is because the Oracle agent service takes some time in coming up. The workaround for this is to increase the attributes DelayAfterOnline and DelayAfterOffline, or probe the resource after the Windows service has started.

- 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

The following figure illustrates a typical Oracle service group resource dependency graph.

Figure C-1 Oracle service group resource dependency graph



Make a note of the following:

- 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.