

Cluster Server 7.1 Agent for Oracle Installation and Configuration Guide - Linux

Cluster Server Agent for Oracle Installation and Configuration Guide

Last updated: 2016-04-26

Document version: 7.1 Rev 0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the Cluster Server agent for Oracle	9
	About the Cluster Server agent for Oracle	9
	About the agent for Oracle ASM	10
	Supported software for VCS agent for Oracle	10
	Changes introduced in the 7.1 release	10
	How the agent makes Oracle highly available	10
	How the Oracle and Netlsnr agents support intelligent resource monitoring	11
	How the agent makes Oracle ASM instance highly available	12
	How the agent supports container database (CDB) and pluggable database (PDB) in Oracle 12C	13
	About Cluster Server agent functions for Oracle	13
	Oracle agent functions	13
	Netlsnr agent functions	28
	ASMinst agent functions	29
	ASMDG agent functions	32
	Typical Oracle configuration in a VCS cluster	33
	About setting up Oracle in a VCS cluster	34
Chapter 2	Installing and configuring Oracle	36
	About installing Oracle in a VCS environment	36
	Before you install Oracle in a VCS environment	37
	About VCS requirements for installing Oracle	37
	Location of the \$ORACLE_HOME	41
	Failing over Oracle after a VCS node failure during hot backup	42
	About Oracle installation tasks for VCS	42
	Installation tasks for \$ORACLE_HOME on shared disks	43
	Installation tasks for \$ORACLE_HOME on local disks	44
	Installation tasks for Oracle using ASM	44
	Installing ASM binaries for Oracle 11gR2 or 12c in a VCS environment	47
	Configuring VCS to support Oracle Restart function	49
	Configuring Oracle ASM on the first node of the cluster	50

	Enabling the clustering daemon for ASM-managed database	52
	Enabling Oracle Cluster Synchronization Service daemon to start automatically	53
	Configuring and starting up ASM on remaining nodes for 11gR2 or 12c	54
	Installing Oracle binaries on the first node of the cluster	55
	Disabling the clustering daemon for Oracle 10g or later	57
	Configuring the Oracle database	58
	Copying the \$ORACLE_BASE/admin/SID directory	59
	Copying the Oracle ASM initialization parameter file	59
	Verifying access to the Oracle database	60
Chapter 3	Installing and removing the agent for Oracle	62
	Before you install or upgrade the agent for Oracle	62
	Installing the VCS agent for Oracle software	62
	Upgrading the VCS agent for Oracle	64
	Upgrading VCS agent for Oracle to use with VCS 7.1	64
	Disabling the Cluster Server agent for Oracle	65
	Removing the Cluster Server agent for Oracle	66
Chapter 4	Configuring VCS service groups for Oracle	67
	About configuring a service group for Oracle	67
	Configuring Oracle instances in VCS	68
	Configuring a single Oracle instance in VCS	68
	Configuring multiple Oracle instances (single listener) in VCS	70
	Configuring multiple Oracle instances (multiple listeners) in VCS	71
	Configuring an Oracle instance with shared server support in VCS	73
	Before you configure the VCS service group for Oracle	75
	Importing the type definition files for Cluster Server agent for Oracle	76
	Configuring the VCS service group for Oracle	77
	Configuring the VCS service group for Oracle using Cluster Manager (Java console)	78
	Configuring the VCS service group for Oracle using the command-line	80
	Setting up detail monitoring for VCS agents for Oracle	84
	Setting up detail monitoring for Oracle	84
	Setting up detail monitoring for Netlsnr	88

	Enabling and disabling intelligent resource monitoring for agents manually	89
	Administering the AMF kernel driver	91
Chapter 5	Configuring VCS service groups for Oracle using the Veritas High Availability Configuration wizard	93
	Typical VCS cluster configuration in a virtual environment	94
	About configuring application monitoring using the High Availability solution for VMware	95
	Getting ready to configure VCS service groups using the wizard	96
	Before configuring application monitoring	97
	Launching the High Availability Configuration wizard	99
	Configuring the agent to monitor Oracle	100
	Understanding service group configurations	106
	Resource dependency	106
	Service group dependency	106
	Infrastructure service groups	107
	Understanding configuration scenarios	107
	Configuring a single instance/multiple instances in VCS	108
	Configuring multiple Oracle instances in VCS using multiple runs of the wizard	108
	Configuring multiple applications	108
	High Availability Configuration wizard limitations	109
	Troubleshooting	109
	High Availability Configuration wizard displays blank panels	109
	The High Availability Configuration wizard displays the 'hadiscover is not recognized as an internal or external command' error	110
	Running the 'hastop -all' command detaches virtual disks	110
	Log files	110
	Sample configurations	110
	Sample VCS configuration file for single Oracle instance (VxVM)	111
	Sample VCS configuration file for single Oracle instance (LVM)	114
Chapter 6	Administering VCS service groups for Oracle	119
	About administering VCS service groups	119
	Bringing the service group online	120
	Taking the service group offline	120

	Switching the service group	120
	Modifying the service group configuration	121
Chapter 7	Pluggable database (PDB) migration	122
	Migrating Pluggable Databases (PDB) between Container Databases (CDB)	122
Chapter 8	Troubleshooting Cluster Server agent for Oracle	127
	About troubleshooting Cluster Server agent for Oracle	127
	Error messages common to the Oracle and Netlsnr agents	128
	Error messages specific to the Oracle agent	130
	Error messages specific to the Netlsnr agent	134
	Error messages specific to the ASMIInst agent	135
	Error messages specific to the ASMDG agent	136
	Troubleshooting issues specific to Oracle in a VCS environment	137
	Verifying the Oracle health check binaries and intentional offline for an instance of Oracle	137
	Verifying the intentional offline behavior of the VCS Oracle agent	139
	Disabling IMF for a PDB resource	140
Appendix A	Resource type definitions	142
	About the resource type and attribute definitions	142
	Resource type definition for the Oracle agent	142
	Attribute definition for the Oracle agent	143
	About the Sid attribute in a policy managed database	152
	Resource type definition for the Netlsnr agent	153
	Attribute definition for the Netlsnr agent	153
	Resource type definition for the ASMIInst agent	157
	Attribute definition for the ASMIInst agent	157
	Resource type definition for the ASMDG agent	161
	Attribute definition for the ASMDG agent	161
Appendix B	Sample configurations	164
	About the sample configurations for Oracle enterprise agent	164
	Sample single Oracle instance configuration	165
	Sample VCS configuration file for single Oracle instance	165
	Sample multiple Oracle instances (single listener) configuration	168

	Sample VCS configuration file for multiple Oracle instances (single listener)	170
	Sample multiple instance (multiple listeners) configuration	175
	Sample VCS configuration file for multiple Oracle instances (multiple listeners)	175
	Sample Oracle configuration with shared server support	179
	Sample VCS configuration file for Oracle instance configured with shared server support	180
	Sample Oracle ASM configurations	183
	Sample configuration for ASM disks as raw disks	183
	Sample configuration for ASM disks as VxVM volumes	187
	Sample configuration for ASM disks as CVM volumes	191
	Sample configuration of Oracle pluggable database (PDB) resource in main.cf	196
	Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf	199
Appendix C	Best practices	205
	Best practices for multiple Oracle instance configurations in a VCS environment	205
Appendix D	Using the SPFILE in a VCS cluster for Oracle	207
	About the Oracle initialization parameter files	207
	Starting an Oracle instance	207
	Using the SPFILE in a VCS cluster	208
Appendix E	OHASD in a single instance database environment	210
	About OHASD in a single instance database environment	210
	Configuring the Application agent to make OHASD highly available	210
Index		214

Introducing the Cluster Server agent for Oracle

This chapter includes the following topics:

- [About the Cluster Server agent for Oracle](#)
- [Supported software for VCS agent for Oracle](#)
- [Changes introduced in the 7.1 release](#)
- [How the agent makes Oracle highly available](#)
- [About Cluster Server agent functions for Oracle](#)
- [Typical Oracle configuration in a VCS cluster](#)
- [About setting up Oracle in a VCS cluster](#)

About the Cluster Server agent for Oracle

The Cluster Server agent for Oracle provides high availability for Oracle database software.

Veritas high availability agents do the following:

- Monitor specific resources within an enterprise application.
- Determine the status of these resources.
- Start or stop the resources according to external events.

The agents include resource type declarations and agent executables. The agent for Oracle monitors the Oracle and listener processes, brings them online, and takes them offline.

The agent RPM for Oracle contains the following agents that work together to make Oracle highly available:

- The Oracle agent monitors the Oracle database processes.
- The Netlsnr agent monitors the listener process.

About the agent for Oracle ASM

The Cluster Server agent suite for Oracle provides high availability agents for the Oracle 10g R2 or later databases that use Automatic Storage Management (ASM). The agents for Oracle ASM monitor the Oracle ASM instance and ASM disk groups.

The VCS Enterprise Agent package also contains the following agents for Oracle ASM:

- The ASMinst agent monitors the Oracle ASM instance.
- The ASMDG agent monitors the Oracle ASM disk groups.

Supported software for VCS agent for Oracle

Refer to the *Cluster Server Release Notes* for software versions supported by VCS 7.1 agent for Oracle.

Refer to the support matrix at

https://www.veritas.com/support/en_US/article.000107212.

Changes introduced in the 7.1 release

VCS agent for Oracle supports Oracle Flex ASM

VCS agent for Oracle supports the Oracle Flex ASM feature. It detects the Flex ASM in your installation and extends the high availability support to the configuration. The disk group status is reported based on the configured ASM cardinality.

How the agent makes Oracle highly available

The Cluster Server agent for Oracle continuously monitors the Oracle database and listener processes to verify they function properly.

See “[About Cluster Server agent functions for Oracle](#)” on page 13.

The agent provides the following levels of application monitoring:

- Primary or Basic monitoring

This mode has Process check and Health check monitoring options. With the default Process check option, the agent verifies that the Oracle and listener processes are present in the process table. Process check cannot detect whether processes are in a hung or stopped states.

The Oracle agent provides functionality to detect whether the Oracle resource was intentionally taken offline. The agent detects graceful shutdown for Oracle 10g and later. When an administrator brings down Oracle gracefully, the agent does not trigger a resource fault even though Oracle is down. The value of the type-level attribute `IntentionalOffline` and the value of the resource-level attribute `MonitorOption` must be set to 1 to enable Oracle agent's intentional offline functionality.

For example, with the intentional offline functionality, the agent faults the Oracle resource if there is an abnormal termination of the instance. The agent reports the Oracle resource as offline if you gracefully bring down Oracle using commands like `shutdown`, `shutdown immediate`, `shutdown abort`, or `shutdown transactional`.

- Secondary or Detail monitoring
In this mode, the agent runs a perl script that executes commands against the database and listener to verify their status.
- When the Oracle database is in READ WRITE mode, the agent performs an UPDATE query and when it is in READ ONLY mode, the agent runs SELECT query against the database.

The Oracle agent also supports IMF (Intelligent Monitoring Framework) in the process check mode of basic monitoring. IMF enables intelligent resource monitoring. The Oracle agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for resource state change notifications.

See [“How the Oracle and Netlsnr agents support intelligent resource monitoring”](#) on page 11.

The agent detects application failure if the monitoring routine reports an improper function of the Oracle or listener processes. When this application failure occurs, the Oracle service group fails over to another node in the cluster. Thus the agent ensures high availability for the Oracle services and the database.

How the Oracle and Netlsnr agents support intelligent resource monitoring

With intelligent monitoring framework (IMF), VCS supports intelligent resource monitoring in addition to the poll-based monitoring. Poll-based monitoring polls the resources periodically whereas intelligent monitoring performs asynchronous

monitoring. You can enable or disable the intelligent resource monitoring functionality for the Oracle and Netlsnr agents.

See [“Enabling and disabling intelligent resource monitoring for agents manually”](#) on page 89.

When an IMF-enabled agent starts up, the agent initializes the asynchronous monitoring framework (AMF) kernel driver. After the resource is in a steady state, the agent registers the details that are required to monitor the resource with the AMF kernel driver. For example, the Oracle agent registers the PIDs of the Oracle processes with the AMF kernel driver. The agent's *imf_getnotification* function waits for any resource state changes. When the AMF kernel driver module notifies the *imf_getnotification* function about a resource state change, the agent framework runs the monitor agent function to ascertain the state of that resource. The agent notifies the state change to VCS which takes appropriate action.

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

How the agent makes Oracle ASM instance highly available

Oracle 10g and later provides ASM feature to store and manage the Oracle data that includes data files, control files, and log files. These Oracle data files that are stored in ASM disk groups are called ASM files.

For ASM-managed databases, you must start an ASM instance before you start the database instance. The ASM instance mounts ASM disk groups and makes ASM files available to database instances.

ASM requires Cluster Synchronization Services to enable synchronization between an ASM instance and a database instance.

See Oracle documentation.

The Cluster Server agent for Oracle has the following agents to keep the Oracle ASM instance and ASM disk groups highly available:

- ASMInst agent
See [“ASMInst agent functions”](#) on page 29.
- ASMDG agent
See [“ASMDG agent functions”](#) on page 32.

The high availability agent for Oracle continuously monitors the ASM instance and ASM disk groups to verify they function properly.

The Flex ASM functionality introduced in Oracle 12c is supported through ASMDG agent of the VCS agent for Oracle.

For Oracle 10g versions, the VCS agents for Oracle ASM use `sysdba` role to connect to the ASM instances. For Oracle 11g and later, the agents use `sysasm` role to connect to the ASM instances.

How the agent supports container database (CDB) and pluggable database (PDB) in Oracle 12C

The VCS agent for Oracle supports management of container and pluggable databases.

During FAULT of a CDB or critical PDBs, VCS fails over the entire container database from one server to another server along with the configured Pluggable databases. The PDBs, which are part of a CDB, are managed using separate PDB resources.

Note: Oracle agent does not support IMF, health check monitoring and intentional offline.

See [“Disabling IMF for a PDB resource”](#) on page 140.

See [“Migrating Pluggable Databases \(PDB\) between Container Databases \(CDB\)”](#) on page 122.

About Cluster Server agent functions for Oracle

The functions an agent performs are called entry points. Review the functions for the following agents that are part of the Cluster Server agent suite for Oracle:

- Oracle agent functions
See [“Oracle agent functions”](#) on page 13.
- Netlsnr agent functions
See [“Netlsnr agent functions”](#) on page 28.
- ASMInst agent functions
See [“ASMInst agent functions”](#) on page 29.
- ASMDG agent functions
See [“ASMDG agent functions”](#) on page 32.

Oracle agent functions

The Oracle agent monitors the database processes.

[Table 1-1](#) lists the Oracle agent functions.

Table 1-1 Oracle agent functions

Agent operation	Description
Online	<p>Starts the Oracle database by using the following <code>sqlplus</code> command:</p> <pre>startup force pfile=\$PFile</pre> <p>The default Startup option is <code>STARTUP_FORCE</code>. You can also configure the agent to start the database using different Startup options for Oracle.</p> <p>See “Startup and shutdown options for the Oracle agent” on page 15.</p>
Offline	<p>Stops the Oracle database with the specified options by using the following <code>sqlplus</code> command:</p> <pre>shutdown immediate</pre> <p>The default Shutdown option is <code>IMMEDIATE</code>. You can also configure the agent to stop the database using different Shutdown options for Oracle.</p> <p>See “Startup and shutdown options for the Oracle agent” on page 15.</p>
Monitor	<p>Verifies the status of the Oracle processes. The Oracle agent provides two levels of monitoring: basic and detail.</p> <p>See “Monitor options for the Oracle agent in traditional database and container database” on page 20.</p>
oracle_imf_init	<p>Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Oracle agent. This function runs when the agent starts up.</p>
oracle_imf_getnotification	<p>Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.</p>
oracle_imf_register	<p>Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).</p>

Table 1-1 Oracle agent functions (*continued*)

Agent operation	Description
Clean	<p>Forcibly stops the Oracle database by using the following <code>sqlplus</code> command:</p> <pre>shutdown abort</pre> <p>If the process does not respond to the <code>shutdown</code> command, then the agent does the following:</p> <ul style="list-style-type: none"> ■ Scans the process table for the processes that are associated with the configured instance ■ Kills the processes that are associated with the configured instance
Info	<p>Provides the static and dynamic information about the state of the database.</p> <p>Note: This attribute is not applicable PDB resources.</p> <p>See “Info entry point for Cluster Server agent for Oracle” on page 23.</p>
Action	<p>Performs the predefined actions on a resource.</p> <p>See “Action entry point for Cluster Server agent for Oracle” on page 24.</p>

Startup and shutdown options for the Oracle agent

You can specify Startup and Shutdown options for the Oracle instances that are configured.

[Table 1-2](#) lists the startup options that the agent supports for traditional database and container database.

Table 1-2 Startup options for traditional database and container database

Option	Description
STARTUP_FORCE (Default)	<p>Runs the command <code>startup force pfile='location_of_pfile'</code> if the pfile is configured.</p> <p>If the pfile is not configured, the agent runs <code>startup force</code>. It picks up the default parameter files from their default locations.</p>

Table 1-2 Startup options for traditional database and container database
(continued)

Option	Description
STARTUP	<p>Runs the command <code>startup pfile='pfile_name_with_absolute_path'</code> if the pfile is configured.</p> <p>For example:</p> <pre>startup pfile = '</dbs/pfilename>'</pre> <p>If the pfile is not configured, the agent runs <code>startup</code> without <code>pfile</code>.</p>
RESTRICTED	Starts the database in the RESTRICTED mode.
RECOVERDB	Performs a database recovery on instance startup.
CUSTOM	<p>Uses a predefined SQL script (<code>start_custom_\${SID}.sql</code>) and runs custom startup options. The script must be in the <code>/opt/VRTSagents/ha/bin/Oracle</code> directory and must have access to the Oracle Owner OS user. If the file is not present, the agent logs an error message.</p> <p>With a custom script, the agent takes the following action:</p> <pre>sqlplus /nolog <<! connect / as sysdba; @start_custom_\${SID}.sql exit; !</pre>
SRVCTLSTART	<p>Uses the <code>srvctl</code> utility to start an instance of the database.</p> <p>For RAC clusters, you must manually set the default startup option as SRVCTLSTART.</p> <p>Note: This is applicable only when an Oracle grid infrastructure is configured.</p>
SRVCTLSTART_RO	<p>Uses the <code>srvctl</code> utility to start an instance of the database with the <code>-o 'read only'</code> option.</p> <p>Note: This is applicable only when an Oracle grid infrastructure is configured.</p>

Table 1-3 lists the shutdown options that the agent supports for traditional database and container database.

Table 1-3 Shutdown options for traditional database and container database

Option	Description
IMMEDIATE (Default)	Shuts down the Oracle instance by running <code>shutdown immediate</code> .
TRANSACTIONAL	Runs the <code>shutdown transactional</code> command. This option is valid only for the database versions that support this option.
CUSTOM	Uses a predefined SQL script (<code>shut_custom_\$\$SID.sql</code>) and runs custom shutdown options. The script must be in the <code>/opt/VRTSagents/ha/bin/Oracle</code> directory and must have access to the Oracle Owner OS user. If the file is not present, the agent shuts the agent down with the default option.
SRVCTLSTOP	Uses the <code>srvctl</code> utility to stop an instance of the database. For RAC clusters, you must manually set the default option as <code>SRVCTLSTOP</code> . Note: This is applicable only when an Oracle grid infrastructure is configured.
SRVCTLSTOP_TRANSACTIONAL	Uses the <code>srvctl</code> utility to stop an instance of the database with the <code>-o transactional</code> option. Note: This is applicable only when an Oracle grid infrastructure is configured.
SRVCTLSTOP_ABORT	Uses the <code>srvctl</code> utility to stop an instance of the database with the <code>-o abort</code> option. Note: This is applicable only when an Oracle grid infrastructure is configured.
SRVCTLSTOP_IMMEDIATE	Uses the <code>srvctl</code> utility to stop an instance of the database with the <code>-o immediate</code> option.

Startup and shutdown options for the pluggable database (PDB)

You can specify Startup and Shutdown options for the PDB resources that are configured.

[Table 1-4](#) lists the startup options that the agent supports in the PDB.

Table 1-4 Startup options for PDB

Options	Description
STARTUP	Runs the command <code>STARTUP PLUGGABLE DATABASE <pdbname></code> ; if the <code>pdbname</code> is configured.
STARTUP_FORCE	Runs the command <code>STARTUP PLUGGABLE DATABASE <pdbname> FORCE</code> ; if the <code>pdbname</code> is configured.
RESTRICTED	Runs the command <code>STARTUP PLUGGABLE DATABASE <pdbname> RESTRICT</code> ; if the <code>pdbname</code> is configured.
OPEN_RO	Runs the command <code>STARTUP PLUGGABLE DATABASE <pdbname> OPEN READ ONLY</code> ; if the <code>pdbname</code> is configured.

During offline, shutdown of individual PDB will be attempted by running `shutdown immediate`.

See [“Recommended startup modes for pluggable database \(PDB\) based on container database \(CDB\) startup modes”](#) on page 18.

See [“Migrating Pluggable Databases \(PDB\) between Container Databases \(CDB\)”](#) on page 122.

Recommended startup modes for pluggable database (PDB) based on container database (CDB) startup modes

You can specify the startup options for the PDB based on the values configured in the CDB.

[Table 1-5](#) lists the shutdown options that the agent supports.

Table 1-5 CDB Startup options

CDB startup mode	Corresponding PDB startup modes
SRVCTLSTART	<p>When the CDB is set with the startup value as SRVCTLSTART then the PDB can be assigned with one of the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ OPEN_RO
SRVCTLSTART_RO	<p>When the CDB is assigned with the SRVCTLSTART_RO startup value then the PDB should be OPEN_RO.</p>
STARTUP	<p>When the CDB is set with the startup value as STARTUP then the PDB can be assigned with one of the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ OPEN_RO
STARTUP_FORCE	<p>When the CDB is set with the startup value as STARTUP_FORCE then the PDB can be assigned with one of the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ OPEN_RO
RESTRICTED	<p>When the CDB is assigned with the startup value as RESTRICTED then the PDB should be RESTRICTED.</p>
RECOVERDB	<p>When the CDB is set with the startup value as RECOVERDB then the PDB can be assigned with one of the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ OPEN_RO

Table 1-5 CDB Startup options (*continued*)

CDB startup mode	Corresponding PDB startup modes
CUSTOM	The custom script should start the PDB as CUSTOM.

See [“Sample configuration of Oracle pluggable database \(PDB\) resource in main.cf”](#) on page 196.

See [“Migrating Pluggable Databases \(PDB\) between Container Databases \(CDB\)”](#) on page 122.

Monitor options for the Oracle agent in traditional database and container database

The Oracle agent provides two levels of monitoring: basic and detail. By default, the agent does a basic monitoring.

The basic monitoring mode has the following options:

- Process check
- Health check

The MonitorOption attribute of the Oracle resource determines whether the agent must perform basic monitoring in Process check or Health check mode.

[Table 1-6](#) describes the basic monitoring options.

Table 1-6 Basic monitoring options

Option	Description
0 (Default)	<p>Process check</p> <p>The agent scans the process table for the ora_dbw0, ora_smon, ora_pmon, ora_lmon, and ora_lgwr processes to verify that Oracle is running.</p> <p>In this mode, the agent also supports intelligent resource monitoring.</p> <p>See “How the Oracle and Netlsnr agents support intelligent resource monitoring” on page 11.</p>

Table 1-6 Basic monitoring options (*continued*)

Option	Description
1	<p>Health check (supported on Oracle 11g and later)</p> <p>The agent uses the Health Check APIs from Oracle to monitor the SGA and retrieve the information about the instance.</p> <p>If you want to use the Oracle agent's intentional offline functionality, you must enable Health check monitoring.</p> <p>See "How the agent makes Oracle highly available" on page 10.</p> <p>The agent does not support intelligent resource monitoring in this mode.</p> <p>See "How the Oracle agent supports health check monitoring" on page 26.</p>

Review the following considerations if you want to configure basic monitoring:

- Basic monitoring of Oracle processes is user-specific. As a result, an Oracle instance started under the context of another user cannot be detected as online. For example, if an Oracle instance is started under the user "oraVRT" and the agent is configured for a user "oracle", the agent will not detect the instance started by "oraVRT" as online.

This could lead to situations where issuing a command to online a resource on a node might online an already running instance on that node (or any other node).

So, Veritas recommends that instances started outside VCS control be configured with the correct Owner attribute corresponding to the OS user for that instance.

In the detail monitoring mode, the agent performs a transaction on a test table in the database to ensure that Oracle database functions properly. The agent uses this test table for internal purposes. Veritas recommends that you do not perform any other transaction on the test table.

See ["How the agent handles Oracle error codes during detail monitoring"](#) on page 22.

See ["Setting up detail monitoring for VCS agents for Oracle"](#) on page 84.

Monitor for the pluggable database

During monitor, the Oracle agent connects to the CDB and fetches the status of all PDBs that are part of the CDB.

When IntentionalOffline is set for a CDB and PDB resource, and CDB is shut outside VCS then the PDB resource reports as Offline and not FAULTED. To enable this functionality, set the IntentionalOffline attribute to 1 for the PDB resource.

Note: Pluggable database will be failed over along with the Container database. Individual Pluggable database cannot be failed over using VCS.

How the agent handles Oracle error codes during detail monitoring

The Cluster Server agent for Oracle handles Oracle errors during detail monitoring. The agent classifies Oracle errors according to their severity and associates predefined actions with each error code.

The agent includes a reference file oraerror.dat, which lists Oracle errors and the action to be taken when the error is encountered.

The file stores information in the following format:

```
Oracle_error_string:action_to_be_taken
```

For example:

```
01035:WARN
01034:FAILOVER
```

Note: When the oraerror.dat file contains two or more errors, the order in which the Cluster Server agent for Oracle handles the errors is: FAILOVER/NOFAILOVER, UNKNOWN, and IGNORE/WARN.

[Table 1-7](#) lists the predefined actions that the agent takes when an Oracle error is encountered.

Table 1-7 Predefined agent actions for Oracle errors

Action	Description
IGNORE	<p>Ignores the error.</p> <p>When the Veritas agent for Oracle encounters an error, the agent matches the error code in the oraerror.dat file. If the error does not have a matching error code in the file, then the agent ignores the error.</p>
UNKNOWN	<p>Marks the resource state as UNKNOWN and sends a notification if the Notifier resource is configured. See the <i>Cluster Server Administrator's Guide</i> for more information about VCS notification.</p> <p>This action is typically associated with configuration errors or program interface errors.</p>

Table 1-7 Predefined agent actions for Oracle errors (*continued*)

Action	Description
WARN	Marks the resource state as ONLINE and sends a notification if the Notifier resource is configured. This action is typically associated with errors due to exceeded quota limits, session limits/restricted sessions so on.
FAILOVER (Default)	Marks the resource state as OFFLINE. This faults the service group by default, which fails over to the next available system. If the file oraerror.dat is not available, the agent assumes this default behavior for every Oracle error encountered.
NOFAILOVER	Freezes the service group temporarily and marks the resource state as OFFLINE. The agent also sends a notification if the Notifier resource is configured. This action is typically associated with the errors that are not system-specific. For example, if a database does not open from a node due to corrupt Oracle files, failing it over to another node does not help.

Info entry point for Cluster Server agent for Oracle

The Cluster Server agent for Oracle supports the Info entry point, which provides static and dynamic information about the state of the database.

To invoke the Info entry point, type the following command:

```
# hares -value resource ResourceInfo [system] \  
[-clus cluster | -localclus]
```

For the above command to be executed successfully, you must perform the following command:

- Execute the resource refreshinfo using the following command:

```
# hares -refreshinfo <res-name> -sys <sys-name>
```

The entry point retrieves the following static information:

- Version
- InstanceNo
- InstanceName
- DatabaseName
- HostName
- StartupTime
- Parallel
- Thread
- InstanceRole

The entry point retrieves the following dynamic information:

- InstanceStatus
- Logins
- OpenMode
- LogMode
- ShutdownPending
- DatabaseStatus
- Shared Pool Percent free
- Buffer Hits Percent

You can add additional attributes by adding sql statements to the file `/opt/VRTSagents/ha/bin/Oracle/resinfo.sql`. For example:

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

The format of the selected record must be as follows:

```
attribute_type:userkey_name:userkey_value
```

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

Action entry point for Cluster Server agent for Oracle

The Cluster Server agent for Oracle supports the Action entry point, which enables you to perform predefined actions on a resource.

To perform an action on a resource, type the following command:

```
# hares -action res token [-actionargs arg1 ...] \
[-sys system] [-clus cluster]
```

You can also add custom actions for the agent.

For further information, refer to the *Cluster Server Agent Developer's Guide*.

See [Table 1-9](#) on page 25. describes the agent's predefined virtual fire drill actions.

[Table 1-8](#) describes the agent's predefined actions.

Table 1-8 Predefined agent actions

Action	Description
VRTS_GetInstanceName	Retrieves the name of the configured instance. You can use this option for the Oracle and the Netlsnr resources.
VRTS_GetRunningServices	Retrieves the list of processes that the agent monitors. You can use this option for the Oracle and the Netlsnr resources.

Table 1-8 Predefined agent actions (*continued*)

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.
DBTbspBackup	Backs up a tablespace; <code>actionargs</code> contains name of the tablespace to be backed up.

Note: VRTS_GetRunningServices, DBRestrict, DBResume, DBSuspend, DBTbspBackup, and DBUndoRestrict action entry points are not supported for Oracle PDB resources.

[Table 1-9](#) lists the virtual fire drill actions of the Cluster Server agent for Oracle lets you run infrastructure checks and fix specific errors.

Table 1-9 Predefined virtual fire drill actions

Virtual fire drill action	Description
getid (Oracle agent)	Verifies that the Oracle Owner exists on the node.
home.vfd (Oracle agent)	Verifies the following: <ul style="list-style-type: none"> ■ ORACLE_HOME is mounted on the node and corresponding entry is in the fstab. If the ORACLE_HOME is not mounted, the action entry point checks if any other resource has already mounted ORACLE_HOME. ■ Pfile is provided and it exists on the node. ■ Password file from \$ORACLE_HOME/dbs/orapw[SID] is present.
owner.vfd (Oracle agent)	Verifies the uid and gid of the Oracle Owner attribute. Checks if uid and gid of Owner attribute is the same on the node where the Oracle resource is currently ONLINE.

Table 1-9 Predefined virtual fire drill actions (*continued*)

Virtual fire drill action	Description
pfile.vfd (Oracle agent)	Checks for the presence of pfile or spfile on the local disk. If both pfile and spfile are not present, the agent function exits. If the Oracle resource is online in the cluster, the agent function logs a message that the spfile must be on the shared storage because the Oracle resource is online.
tnsadmin.vfd (Netlsnr agent)	Checks if listener.ora file is present. If the listener.ora file is not present, it checks if ORACLE_HOME is mounted and displays appropriate messages.

How the Oracle agent supports health check monitoring

The Cluster Server agent for Oracle uses the Oracle health check APIs to determine whether the Oracle instance on a node was shut down gracefully or was it aborted. When an oracle instance is shut down gracefully outside of VCS control the agent acknowledges the operation as intentional offline.

The Oracle health check APIs are supported on all UNIX platforms, Oracle Database 10g Release and later.

You can use the `build_oraapi.sh` script to link the libraries that are required for Oracle health check monitoring. ships the `build_oraapi.sh` script along with the VCS agent for Oracle in the Oracle agent's `/opt/VRTSagents/ha/bin/Oracle` directory.

Generating the executable required for Oracle health check monitoring

Review the considerations before using the `build_oraapi.sh` script.

Considerations to run the `build_oraapi.sh` script

- Run the `build_oraapi.sh` script on each node of the cluster where you want to use Oracle health monitoring.
- You need to run the `build_oraapi.sh` script to relink the Oracle libraries whenever the Oracle database version is upgraded on a node.
- The Oracle health check binaries generated in the Oracle agent directory correspond to the `ORACLE_HOME` directory and the Oracle database version that are specified while executing the `build_oraapi.sh` script. `ORACLE_HOME` can be set as an environment variable or specified as a command line option. The Oracle database version can be specified as a command line option with the `build_oraapi.sh` script.

- Veritas ships the `build_oraapi.sh` script in VCS 6.2 and later releases. Veritas will not ship the pre-built Oracle health check binaries (`oraapi_32`, `oraapi_3211g`, `oraapi_64`, and `oraapi_6411g`).
- To enable health check support in Oracle agent, you must run the `build_oraapi.sh` script to build the Oracle health check binaries based on Oracle version.

Building the Oracle agent executable using the `build_oraapi.sh` script

To build the Oracle agent executable using the `build_oraapi.sh` script:

1 Log in as root user on the node.

2 Set the Oracle home directory path.

```
# export ORACLE_HOME=/u01/oraHome
```

where, `/u01/oraHome` is the Oracle home directory that stores files.

3 Access the Oracle agent directory.

```
# cd /opt/VRTSagents/ha/bin/Oracle
```

4 Run the following command from the Oracle agent directory.

```
# ./build_oraapi.sh -version <version>
```

where *version* is the Oracle database version on the node.

5 Alternatively, if you did not set the `ORACLE_HOME` variable, you can set the `oracle_home` parameter as a command line option.

```
#./build_oraapi.sh -version <version> -oracle_home /u01/oraHome
```

where *version* is the Oracle database version on the node and `/u01/oraHome` is the Oracle home directory path, which is set as the `ORACLE_HOME` variable value.

The script generates binaries under agent directory

`/opt/VRTSagents/ha/bin/Oracle`. Oracle health check binaries generated on a node is based on the system architecture (32 bit or 64 bit) available on that node, and the full Oracle version.

For example, `oraapi_64_121010` is a binary created with Oracle version 12 on a 64-bit system.

See [“Verifying the Oracle health check binaries and intentional offline for an instance of Oracle”](#) on page 137.

Netlsnr agent functions

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.

The Netlsnr agent is IMF-aware.

See [“How the Oracle and Netlsnr agents support intelligent resource monitoring”](#) on page 11.

[Table 1-10](#) lists the Netlsnr agent functions.

Table 1-10 Netlsnr agent functions

Agent operation	Description
Online	Starts the listener process by using the following command: <code>lsnrctl start \$LISTENER</code>
Offline	Stops the listener process by using the following command: <code>lsnrctl stop \$LISTENER</code> If the listener is configured with a password, the agent uses the password to stop the listener.
Monitor	Verifies the status of the listener process. The Netlsnr agent provides two levels of monitoring, basic and detail: <ul style="list-style-type: none"> ■ In the basic monitoring mode, the agent scans the process table for the <code>tnslsnr</code> process to verify that the listener process is running. ■ In the detail monitoring mode, the agent uses the <code>lsnrctl status \$LISTENER</code> command to verify the status of the Listener process. (Default)
netlsnr_imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Netlsnr agent. This function runs when the agent starts up.
netlsnr_imf_getnotification	Gets notification about resource state change. This function runs after the agent registers with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
netlsnr_imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Table 1-10 Netlsnr agent functions (*continued*)

Agent operation	Description
Clean	Scans the process table for <code>tnslsnr \$LISTENER</code> and kills it.
Action	Performs the predefined actions on a resource. See “Action entry point for Cluster Server agent for Oracle” on page 24.

ASMIInst agent functions

The ASMIInst agent monitors the processes of ASM instance.

[Table 1-11](#) lists the ASMIInst agent operations.

Table 1-11 ASMIInst agent operations

Agent operation	Description
Online	Starts the Oracle ASM instance by using the following SQL command: <code>startup nomount</code> The Online operation starts the ASM instance without mounting any disk groups.
Offline	Stops the Oracle ASM instance by using the following SQL command: <code>shutdown immediate</code>

Table 1-11 ASMinst agent operations (*continued*)

Agent operation	Description
Monitor	<p>Verifies the status of the ASM instance.</p> <p>The ASMinst agent can monitor the ASM instance in two ways; process check monitoring and health check monitoring. The default is process check monitoring.</p> <p>In the process check monitoring mode, the agent scans the process table for the <code>asm_pmon</code>, <code>asm_smon</code>, <code>asm_lgwr</code>, <code>asm_rbal</code>, and <code>asm_dbw0</code> processes to verify that the ASM instance is running.</p> <p>The agent also checks if the <code>ocssd.bin</code> process is running. The agent returns offline for the following conditions:</p> <ul style="list-style-type: none"> ■ The process is not running. ■ The process is restarted. <p>Note: Make sure that the OCSSD process is running. The ASMinst agent only monitors the OCSSD process. The agent does not start or stop the process.</p>
Clean	<p>Forcibly stops the Oracle ASM instance by using the following SQL command:</p> <pre>shutdown abort</pre> <p>If the process does not respond to the <code>shutdown</code> command, the agent kills the process using the <code>SIGTERM</code> or <code>SIGKILL</code> commands.</p>

ASMinst agent Startup options

You can specify Startup options for the Oracle ASM instances that are configured.

[ASMinst agent Startup options](#) lists the startup options that the agent supports.

Table 1-12 Startup options

Option	Description
STARTUP	<p>The default startup option for ASM instance resource. Uses <code>sqlplus</code> to start the ASM instance with the <code>startup nomount</code> startup option. If this option is used, none of the ASM disk groups get mounted.</p>
STARTUP_MOUNT	<p>Uses <code>sqlplus</code> to start the ASM instance with the <code>startup mount</code> startup option. If this option is used, all the ASM disk groups get mounted.</p>

Table 1-12 Startup options (*continued*)

Option	Description
STARTUP_OPEN	Uses <code>sqlplus</code> to start the ASM instance with the <code>startup open</code> startup option. If this option is used, all the ASM disk groups are mounted and volumes are enabled.
SRVCTLSTART	Uses <code>srvctl</code> to start the ASM instance with the <code>nomount</code> startup option. If this option is used, none of the ASM disk groups get mounted.
SRVCTLSTART_MOUNT	Uses <code>srvctl</code> to start the ASM instance with the <code>-o mount</code> startup option. If this option is used, all the disk groups get mounted
SRVCTLSTART_OPEN	Uses <code>srvctl</code> to start the ASM instance with the <code>-o open</code> startup option. If this option is used, all the ASM disk groups are mounted and volumes are enabled.

Note: recommends to use STARTUP or SRVCTLSTART option for ASMinst resource.

ASMinst agent Shutdown options

You can specify the Shutdown option for the Oracle ASM instances that are configured.

Table 1-13 Shutdown options

Option	Description
SRVCTLSTOP	Uses the <code>srvctl</code> utility to stop an ASM instance.

By default the agent uses the `sqlplus` command to stop the Oracle ASM instance. To stop the ASM instance using `srvctl` utility set the option to `SRVCTLSTOP`.

See [“Attribute definition for the ASMinst agent”](#) on page 157.

Monitor options for the ASMinst agent

The ASMinst agent's monitoring mode has the following options:

- Process check
- Health check

The MonitorOption attribute of the ASMIInst resource determines whether the agent performs basic monitoring in Process check or Health check mode.

[Table 1-14](#) describes the ASMIInst monitoring options.

Table 1-14 ASMIInst monitoring options

Option	Description
0 (Default)	Process check The agent scans the process table for the ora_dbw0, ora_smon, ora_pmon, and ora_lgwr processes to verify that ASM instance is running.
1	Health check (supported on Oracle 10g and later) The agent uses the Health Check APIs from Oracle to monitor the SGA and retrieve the information about the instance. If you want to use the Oracle agent's intentional offline functionality, you must enable Health check monitoring.

ASMDG agent functions

The ASMDG agent mounts the ASM disk groups that the Oracle databases use, monitors the status, unmounts the ASM disk groups.

You must have specified the disk group names in the DiskGroups attribute of the ASMDG agent.

[Table 1-15](#) lists the ASMDG agent operations.

Table 1-15 ASMDG agent operations

Agent operation	Description
Online	Mounts the specified Oracle ASM disk groups to an ASM instance by using the following SQL command: <code>alter diskgroup dg_name1, dg_name2 mount</code>
Offline	Unmounts the specified Oracle ASM disk groups from an ASM instance by using the following SQL command: <code>alter diskgroup dg_name1, dg_name2 dismount</code> Note: The following Oracle message appears in the VCS log when an ASM instance with no ASM disk groups mounted is shut down: ORA-15100: invalid or missing diskgroup name

Table 1-15 ASMDG agent operations (*continued*)

Agent operation	Description
Monitor	<p>Verifies the status of the specified ASM disk groups.</p> <p>The disk groups can be in one of the following states:</p> <ul style="list-style-type: none"> ■ mounted ■ dismounted ■ unknown ■ broken ■ connected <p>If multiple ASM disk groups are configured for a resource, then the ASMDG agent returns the resource state considering the status of all the specified ASM disk groups.</p> <p>Note: In case of Flex ASM, the ASMDG agent honors the cardinality configured for the ASM Instances. If the ASM instance is down on a particular node it is reported as offline if the number of ASM instances equal to the cardinality value and it is reported as Faulted if the number of ASM instances falls below the cardinality value.</p>
Clean	<p>Forcibly unmounts the Oracle ASM disk groups by using the following SQL command:</p> <pre>alter diskgroup dg_name1, dg_name2 dismount force</pre>

Typical Oracle configuration in a VCS cluster

A typical Oracle configuration in a VCS cluster has the following characteristics:

- VCS is configured in a two-node cluster.
- The Oracle data is installed on shared storage.
- The Oracle binaries are installed locally on both nodes or on shared disks.

If you use Oracle ASM feature, then the characteristics are as follows:

- The Oracle binaries are installed locally on both the nodes.
- The database files are installed on the ASM disk groups.

Review the typical cluster setup illustration for ASM-managed database.

See [Figure 1-2](#) on page 34.

- The Cluster Server agent for Oracle is installed on both nodes.
 For ASM, ASMinst and ASMDG agents are installed on both nodes.

Figure 1-1 depicts a configuration where Oracle binaries and data are installed completely on shared disks.

Figure 1-1 Oracle binaries and data on shared disks

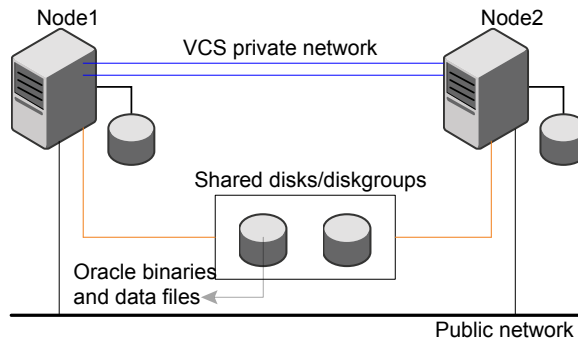
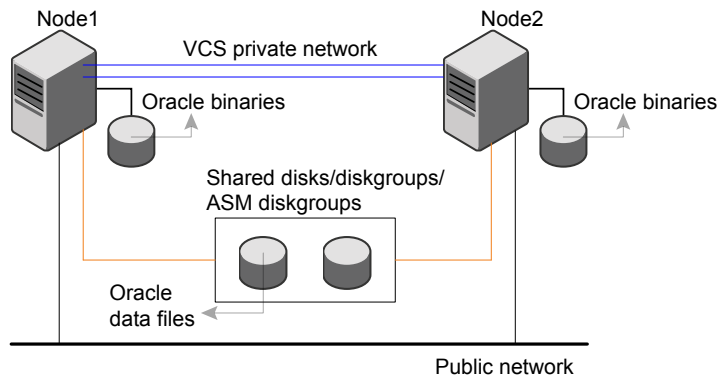


Figure 1-2 depicts a configuration where Oracle binaries are installed locally on each node and Oracle data is installed on shared disks.

Figure 1-2 Oracle binaries on local disk and Oracle data on shared disk



About setting up Oracle in a VCS cluster

Tasks involved in setting up Oracle in a VCS environment include:

- Setting up a VCS cluster
Refer to *Cluster Server Installation Guide* for more information on installing and configuring VCS.
- Installing and configuring Oracle
See [“About installing Oracle in a VCS environment”](#) on page 36.

- Installing the Cluster Server agent for Oracle
See [“Before you install or upgrade the agent for Oracle”](#) on page 62.
- Configuring VCS service groups for Oracle
See [“About configuring a service group for Oracle”](#) on page 67.

Installing and configuring Oracle

This chapter includes the following topics:

- [About installing Oracle in a VCS environment](#)
- [Before you install Oracle in a VCS environment](#)
- [About VCS requirements for installing Oracle](#)
- [About Oracle installation tasks for VCS](#)
- [Installing ASM binaries for Oracle 11gR2 or 12c in a VCS environment](#)
- [Configuring Oracle ASM on the first node of the cluster](#)
- [Configuring and starting up ASM on remaining nodes for 11gR2 or 12c](#)
- [Installing Oracle binaries on the first node of the cluster](#)
- [Configuring the Oracle database](#)
- [Copying the \\$ORACLE_BASE/admin/SID directory](#)
- [Copying the Oracle ASM initialization parameter file](#)
- [Verifying access to the Oracle database](#)

About installing Oracle in a VCS environment

The strategy for installing Oracle into a VCS cluster is aimed at ensuring that installations on all nodes in the cluster are uniform.

See the Oracle documentation on Linux.

You can install Oracle in the following ways in a VCS environment:

\$ORACLE_HOME on the shared disk	Oracle binaries and Oracle data are installed on shared disks.
\$ORACLE_HOME on the local disk	Oracle binaries are installed locally on each node and Oracle data is installed on shared disks.

Note: To use ASM for Oracle database storage, you must install \$ORACLE_HOME on the local disks of all nodes in the cluster.

See [“About Oracle installation tasks for VCS”](#) on page 42.

Note that Oracle data includes the datafiles, control files, redo log files, and archive log files.

When installing Oracle, ensure that the login_id, id_name, group_id, and group_name for the Oracle owner is the same on all the nodes. The user oracle and the group dba must be local and not Network Information Service (NIS and NIS+) users.

Before you install Oracle in a VCS environment

Make sure you meet the following prerequisites:

- Verify that VCS is installed on all nodes in the cluster.
- Verify that all nodes in the cluster have adequate resources to run Oracle and VCS.
- Verify that the network supports the TCP/IP protocol.
- Make sure that you meet the VCS requirements to install Oracle.

About VCS requirements for installing Oracle

Make sure you meet the following requirements to install Oracle in a VCS cluster:

Kernel parameter configuration	<p>Each node on which you want to install Oracle must meet the following Oracle configuration requirements:</p> <ul style="list-style-type: none">■ Disk partitions■ Shared memory■ Swap size■ Semaphores■ File handles <p>See Oracle documentation for the corresponding operating system for specific requirement details.</p>
Linux package group	<p>Veritas recommends you to select the Software Development package option when you install Linux. This option is required for relinking Oracle at install time and to take advantage of some Advanced Server features.</p>
Location of the \$ORACLE_HOME	<p>Depending on your environment, you can place the Oracle home directory (\$ORACLE_HOME) in one of the following ways:</p> <ul style="list-style-type: none">■ Locally on each server's disk■ On the shared storage. <p>If you want to use Oracle ASM, then you must place the Oracle home directory only on the local disks of each node.</p> <p>Review the advantages of each approach to make a decision.</p> <p>See "Location of the \$ORACLE_HOME" on page 41.</p>
Configurations with multiple Oracle instances (SIDs)	<p>You can have multiple Oracle instances that are defined in a single cluster. In such cases, the parameter file for each instance must be accessible on all the nodes in the service group's SystemList attribute.</p> <p>Note: If you installed multiple versions of Oracle on the same system, make sure that the SIDs are unique.</p>

Location of Oracle database tablespaces If you plan to create the tablespaces using regular (ext3 or VxFS) files, the file systems that contain these files must be located on shared disks. Create the same file system mount points on each node.

If you use raw devices on shared disks for Oracle tablespaces, you must meet the following requirements:

- The ownership must be Oracle dba user.
- The permissions or access mode must be 660 on the raw devices that store the Oracle data.

For example, if you use Veritas Volume Manager, type:

```
# vxedit -g diskgroup_name set group=dba \  
user=oracle mode=660 volume_name
```

Note: The user `oracle` and the group `dba` must be local and not Network Information Service (NIS and NIS+) users

Transparent listener failover You can enable Oracle Server clients to reconnect after a node switch without reconfiguring. For such reconnections you must include at least one IP resource in the service group for the Oracle resource. The hostname mapping the IP address of this resource must be used for the Host field in the file `$TNS_ADMIN/listener.ora`.

If you use the TCP/IP protocol for Oracle client/server communication, verify that the file `/etc/services` contains the service name of the Oracle Net Service. You must verify this file on each node that is defined in the service group's SystemList attribute.

Listener authentication in VCS environment The Netlsnr agent supports OS authentication as well as password authentication for the listener process. If you use Oracle 10g or later, recommends you to configure OS authentication. If you want to configure a listener password, make sure that you configure the password correctly. A misconfigured password can cause the listener to fault.

See [“Encrypting Oracle database user and listener passwords”](#) on page 82.

Refer to the Oracle documentation for details on configuring the listener authentication.

Hot backup of Oracle database in VCS environment

The hot backup of Oracle database is enabled by default in VCS environment.

A node can fail during a hot backup of an Oracle database. During such failures, VCS can fail over to another node only if the following requirements are met:

- The AutoEndBkup attribute value must be set to 1, which is the default.
- The Startup option of Oracle agent must be STARTUP, STARTUP_FORCE, or CUSTOM.

See [“Startup and shutdown options for the Oracle agent”](#) on page 15.

If you do not meet VCS requirements, you must manually end the hot backup and then fail over Oracle to another node.

Note: If a node fails during a hot backup of container database or pluggable database for Oracle 12C, you must set AutoEndBkup attribute of the corresponding CDB resource to 1. When the AutoEndBkup is set to 1 for the CDB, it also ends the backup of both CDB and PDB during Online.

See [“Failing over Oracle after a VCS node failure during hot backup”](#) on page 42.

Note: If you set the AutoEndBkup attribute value to 0, then to avoid unexpected VCS behavior you must enable detail monitoring.

See [“Setting up detail monitoring for VCS agents for Oracle”](#) on page 84.

Storage devices for Oracle ASM configurations in VCS

You can choose one of the following storage devices for Oracle ASM:

- ASM disks as raw disks
If you use raw disks, then make sure that the disks have the persistent permissions that are specified for ASM \$ORACLE_HOME.
- ASM disks as Veritas Volume Manager volumes
If you use VxVM volumes, then make sure that the disks have persistent permissions across reboots. The permissions must be the same as that of ASM \$ORACLE_HOME.
See Veritas Volume Manager documentation.
- ASM disks as Cluster Volume Manager volumes
If you use CVM volumes, then make sure that the disks have persistent permissions across reboots. The permissions must be the same as that of ASM \$ORACLE_HOME.
See Veritas Volume Manager documentation.

If you want to configure mirroring for ASM disks that use VxVM or CVM volumes, then you must configure VxVM mirroring and not configure ASM mirroring.

See [“Sample Oracle ASM configurations”](#) on page 183.

From Oracle 11g R2 or 12c, the ASMInst agent does not support pfile or spfile for ASM instances on the ASM disk groups. recommends that you copy this file from ASM disk group to the local file system.

Location of the \$ORACLE_HOME

You can place the Oracle home directory (\$ORACLE_HOME), which contains the Oracle binaries and configuration files, locally on each server's disk. Alternatively, you can place the Oracle home directory on the shared storage. The correct location for Oracle binaries depends on your environment. The following points discuss the advantages of each approach.

\$ORACLE_HOME directory
on shared disks

You can install the Oracle Database Server (\$ORACLE_HOME) on shared disks. Each node in the cluster must have the same mount point directory for the shared file system. Placing the Oracle binaries on shared storage simplifies setting up a given node in a cluster to run an instance. Each database service group is self-contained. An instance can be moved to a new node in the cluster that shares the storage.

For example, in a cluster with four nodes, you can have three database instances or service groups, each at a different Oracle version. If the Oracle binaries are placed on shared storage, three copies of Oracle, that is, one per version are required on shared storage. By contrast, placing the Oracle binaries on local storage, would require as many as 12 copies of the binaries (three versions on four nodes).

The disadvantage of this approach is that a rolling upgrade of Oracle binaries on shared storage is not possible.

\$ORACLE_HOME directory
on local disks

You can install the Oracle Database Server (\$ORACLE_HOME) on the local disk. The advantage is that you can upgrade the Oracle database binaries on an offline node while the database server runs on another node. The database server can later be switched to the upgraded node (provided the database is compatible), permitting a minimum amount of downtime.

The disadvantage of this approach is that with a large number of nodes, it becomes difficult to maintain the various Oracle installations.

If you want to use Oracle ASM, then you must place the Oracle home directory only on the local disks of each node. You can install Oracle ASM on the same Oracle home as Oracle Database, or you can install Oracle ASM on a different Oracle home directory.

See Oracle documentation.

Failing over Oracle after a VCS node failure during hot backup

If a node fails during a hot backup, VCS can succeed with a node failover only if you meet VCS requirements for Oracle hot backup.

See [“About VCS requirements for installing Oracle”](#) on page 37.

If you do not meet VCS requirements, the agent cannot open the database in the backup mode on the failover node. Therefore, VCS cannot online the Oracle group on the failover node. You may encounter the following Oracle errors in this situation:

```
$ ORA-1110 "data file %s: '%s'"
```

or

```
$ ORA-1113 "file %s needs media recovery"
```

In such cases, you must manually end the hot backup and then fail over Oracle to another node.

To manually fail over Oracle after a node failure that occurs during hot backup

- 1 Take the datafiles in the database out of the backup mode.

Refer to the Oracle documentation for instructions on how to change the state of the database files.

- 2 Shut down the database.
- 3 Bring the Oracle service group online on the failover node.

The agent can now start the Oracle instance, mount the database, and open the datafiles.

About Oracle installation tasks for VCS

Tasks to complete Oracle installation in a VCS cluster depend on whether you want the \$ORACLE_HOME on shared disk or local disk.


See [“Installation tasks for \\$ORACLE_HOME on shared disks”](#) on page 43.

See [“Installation tasks for \\$ORACLE_HOME on local disks”](#) on page 44.

For ASM, you must install \$ORACLE_HOME on local disks and configure the Oracle ASM.

[Figure 2-1](#) illustrates the flow of the installation when the \$ORACLE_HOME is on shared disk and on local disk.

Figure 2-1 Comparison of the installation flow



Installation tasks for \$ORACLE_HOME on shared disks

The methodology is to install the Oracle binaries and to create the Oracle database on shared disks during the first installation. Mount the Oracle binaries directory and verify that the database can be started from all nodes.

[Table 2-1](#) lists the tasks to install Oracle such that \$ORACLE_HOME is on shared disks.

Table 2-1 Installation tasks for \$ORACLE_HOME on shared disks

Task	Reference
From any node in the cluster, install Oracle binaries on shared disks.	See “Installing Oracle binaries on the first node of the cluster” on page 55.
Disable the Oracle clustering daemon for Oracle 10g and later.	See “Disabling the clustering daemon for Oracle 10g or later” on page 57.
From the node where you installed Oracle, create the database on shared disks.	See “Configuring the Oracle database” on page 58.

Table 2-1 Installation tasks for \$ORACLE_HOME on shared disks
(continued)

Task	Reference
From each node that would be a part of the Oracle cluster, verify access to the database on shared disks.	See “Verifying access to the Oracle database” on page 60.

Installation tasks for \$ORACLE_HOME on local disks

The methodology is to install the Oracle binaries on the local disk and to create the Oracle database on shared disks during the first installation. Then, install the Oracle binaries on local disks of other nodes. This methodology ensures that all Oracle installations are identical and access the database from the same location on the shared disk.

[Table 2-2](#) lists the tasks to install Oracle such that \$ORACLE_HOME is on local disks.

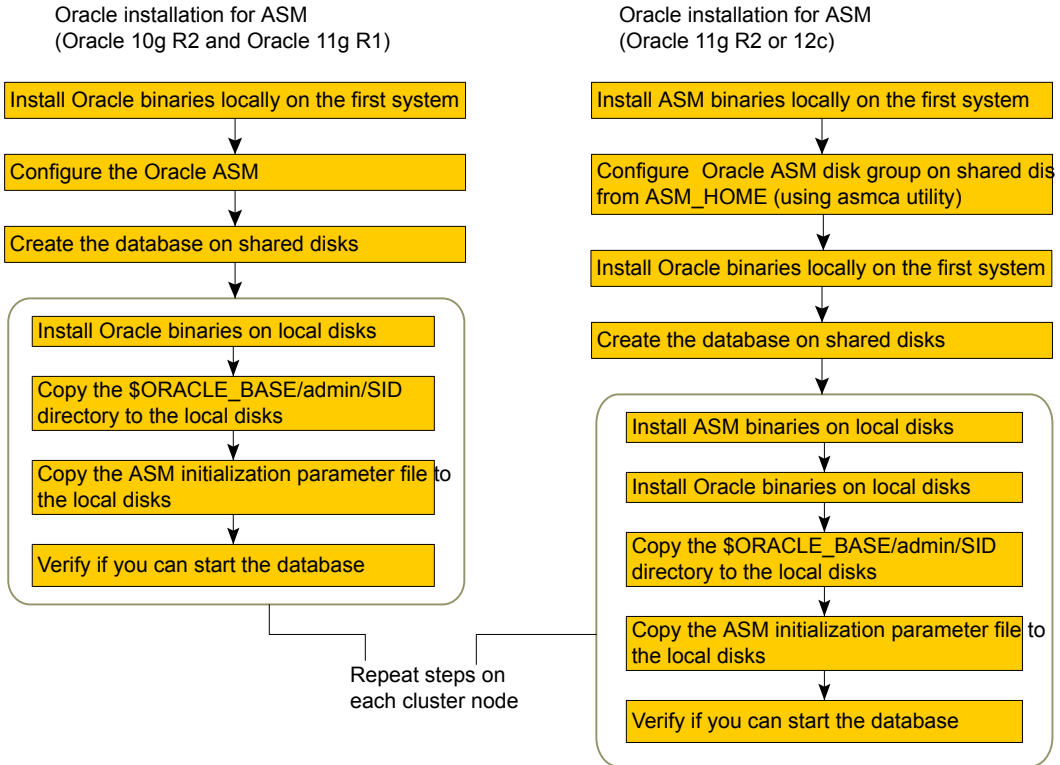
Table 2-2 Installation tasks for \$ORACLE_HOME on local disks

Task	Reference
On the first node in the cluster, install Oracle binaries on the local disk.	See “Installing Oracle binaries on the first node of the cluster” on page 55.
From the first node where you installed Oracle, create the database on shared disks.	See “Configuring the Oracle database” on page 58.
On remaining nodes in the cluster, perform the following tasks: <ul style="list-style-type: none"> ■ Install Oracle binaries on the local disk. ■ Copy the \$ORACLE_BASE/admin/SID directory from the first node to the remaining nodes. Copy \$ORACLE_HOME/dbs/*SID* to the local disk. ■ Verify if you can start the database. 	<ul style="list-style-type: none"> ■ See “Installing Oracle binaries on the first node of the cluster” on page 55. ■ See “Copying the \$ORACLE_BASE/admin/SID directory” on page 59. ■ See “Verifying access to the Oracle database” on page 60.

Installation tasks for Oracle using ASM

[Figure 2-2](#) illustrates the flow of the installation when you want to use ASM for the Oracle database. The installation flow differs if you use Oracle 11g R2.

Figure 2-2 Installation flow for Oracle ASM



The methodology is to perform the following:

- Install the Oracle and the ASM binaries on the local disk, configure the ASM, and to create the Oracle database on ASM disk groups during the first installation.
- Then, install the Oracle and the ASM binaries on local disks of other nodes.

This methodology ensures that all Oracle installations are identical and access the database from the same location.

Table 2-3 lists the tasks to install Oracle and ASM (for Oracle 10g R2 or Oracle 11g R1) such that \$ORACLE_HOME is on local disks.

Table 2-3 Installation tasks for \$ORACLE_HOME on local disks (for Oracle 10g R2 or Oracle 11g R1)

Task	Reference
On the first node in the cluster, install Oracle binaries on the local disk.	See "Installing Oracle binaries on the first node of the cluster" on page 55.
Enable the Oracle clustering daemon for Oracle 10g and later.	See "Enabling the clustering daemon for ASM-managed database" on page 52.
From the first node where you installed Oracle, configure the Oracle ASM.	See "Configuring Oracle ASM on the first node of the cluster" on page 50.
From the first node where you installed Oracle, create the database on ASM disk groups.	See "Configuring the Oracle database" on page 58.
On each node in the cluster, perform the following tasks: <ul style="list-style-type: none"> ■ Install Oracle binaries on the local disk. ■ Copy the \$ORACLE_BASE/admin/SID directory to the local disk. ■ Copy the ASM initialization parameter file to the local disk. ■ Verify if you can start the database. 	<p>See "Installing Oracle binaries on the first node of the cluster" on page 55.</p> <p>See "Copying the \$ORACLE_BASE/admin/SID directory" on page 59.</p> <p>See "Copying the Oracle ASM initialization parameter file" on page 59.</p> <p>See "Verifying access to the Oracle database" on page 60.</p>

[Table 2-4](#) lists the tasks to install Oracle and ASM (for Oracle 11g R2 or 12c) such that \$ORACLE_HOME is on local disks.

Table 2-4 Installation tasks for \$ORACLE_HOME on local disks (for Oracle 11gR2 or 12c)

Task	Reference
On the first node in the cluster, install ASM binaries on the local disk.	See "Installing ASM binaries for Oracle 11gR2 or 12c in a VCS environment" on page 47.
From the first node where you installed ASM, configure Oracle ASM disk groups on shared disks (using <code>asmca</code> utility).	See "Configuring Oracle ASM on the first node of the cluster" on page 50.

Table 2-4 Installation tasks for \$ORACLE_HOME on local disks (for Oracle 11gR2 or 12c) *(continued)*

Task	Reference
On the first node in the cluster, install Oracle binaries on the local disk.	See “Installing Oracle binaries on the first node of the cluster” on page 55.
Enable the Oracle clustering daemon for Oracle 10g and later.	See “Enabling the clustering daemon for ASM-managed database” on page 52.
Enable automatic start up of Oracle Cluster Synchronization Service daemon for Oracle 11g R2	See “Enabling Oracle Cluster Synchronization Service daemon to start automatically” on page 53.
From the first node where you installed Oracle, create the database on ASM disk groups.	See “Configuring the Oracle database” on page 58.
On each node in the cluster, perform the following tasks: <ul style="list-style-type: none"> ■ Install ASM binaries on the local disk. ■ Install Oracle binaries on the local disk. ■ Copy the \$ORACLE_BASE/admin/SID directory to the local disk. ■ Copy the ASM initialization parameter file to the local disk. ■ Verify if you can start the database. 	See “Installing Oracle binaries on the first node of the cluster” on page 55. See “Copying the \$ORACLE_BASE/admin/SID directory” on page 59. See “Copying the Oracle ASM initialization parameter file” on page 59. See “Verifying access to the Oracle database” on page 60.

Installing ASM binaries for Oracle 11gR2 or 12c in a VCS environment

Perform the following procedure if want to configure Oracle ASM instances for high availability for Oracle 11g R2 or 12c.

Note: If you want to use Oracle ASM feature, do not disable the Cluster Synchronization Services.

To install ASM binaries for Oracle 11gR2 or 12c for VCS

- 1 Insert the Oracle CD.
- 2 Set the *DISPLAY* variable, mount the CD drive, and run the Oracle installer as Oracle User.

```
# /mnt/cdrom/grid/runInstaller
```

- 3 In the Select Installation Option dialog box, choose the **Install Grid Infrastructure Software Only** option and click **Next**.
- 4 Follow the wizard instructions and select other options
- 5 In the Specify Installation Location dialog box, specify a location on the local disk for both Oracle base and Software location.
- 6 Follow the wizard instructions and select other options
- 7 In the Summary dialog box, review your selections. Click **Back** to make changes. Otherwise click **Finish**.
- 8 The installer prompts whether you want to instantiate a standalone version of the software or a cluster version of the software. Choose the version that you want to install.

- **Instantiate a standalone version**

Note: If you choose this option, you must execute `roothas.pl`. This option creates a configuration named **Oracle Restart**, which is used on standalone servers and not on RAC clusters. Therefore, you must perform this installation on all nodes in the VCS cluster that are used as failover destinations for Oracle Grid Infrastructure and Oracle ASM.

This option requires you to install the Oracle Grid Infrastructure software locally on every node. You must not use the same set of binaries on more than one node because the same CSSD daemon (used in a particular configuration) cannot be run on nodes with different host names, requiring you to reconfigure the daemon each time it fails over or is used on another node.

In the Oracle Restart configuration the newly created database instance and netlistner resources automatically get registered to the Oracle Grid Infrastructure. In VCS, these resources are configured as part of failover service group.

If the database and netlistner resources accidentally go offline, Oracle Grid Infrastructure tries to restart the resources. However, VCS tries to online

these resources on another node within the failover service group. This behavior causes concurrency violation.

To avoid concurrency violation perform one of the following:

- In the Oracle Grid Infrastructure, set the `auto_start` attribute of each database and `netlistner` resource to 0.
 - In VCS, set the `RestartLimit` attribute value of each database and `netlistner` resource to the corresponding value set for these resources in the `RESTART_ATTEMPTS` attribute in Oracle Grid Infrastructure. For detailed description on configuring VCS to support the Oracle restart function refer to See [“Configuring VCS to support Oracle Restart function”](#) on page 49.
 - **Don't instantiate a standalone version**
Choose this option to instantiate a cluster version of the software. This is required for a RAC database.
- 9 Click **Next** after the installation is complete.
 - 10 In the End of Installation dialog box, click **Exit**.
 - 11 Proceed to configure ASM disk group on shared disks using the `asmca` utility. See [“Configuring Oracle ASM on the first node of the cluster”](#) on page 50.

Configuring VCS to support Oracle Restart function

To ensure that VCS supports the Oracle Restart feature, the user needs to tune the VCS parameters so that it does not attempt a failover action when the Oracle grid infrastructure attempts to restart the resources on the same node. For VCS to take correct decision, the `ToleranceLimit` resource attribute needs to be tuned accordingly.

Oracle Restart feature uses the `CHECK_INTERVAL` and `RESTART_ATTEMPTS` attributes to determine the monitoring interval and restart attempts if the resources are unavailable.

For example, when ASM instance becomes unavailable on a node, the grid infrastructure tries to restart the ASM resource, however VCS detects that the application is unavailable and reports the resource as faulted and tries to failover the service group to another cluster node. To avoid such scenarios, we need to ensure that the VCS policies do not take effect until Oracle completes performing its functions.

You need to modify the VCS resource `ToleranceLimit` attribute such that $ToleranceLimit * MonitorInterval$ is greater than $(CHECK_INTERVAL * RESTART_ATTEMPTS) + MonitorInterval$.

If Oracle's attribute values for ASM resource is CHECK_INTERVAL = 1 (in seconds) and RESTART_ATTEMPTS = 5 (no of restart attempts), and the VCS configuration values for the ASMInst resource has MonitorInterval set to 60, then set the ToleranceLimit to 2 based on the below calculations,

$$2 \text{ (ToleranceLimit)} * 60 \text{ (MonitorInterval)} = 120$$

$$1 \text{ (CHECK_INTERVAL)} * 5 \text{ (RESTART_ATTEMPTS)} + 60 \text{ (MonitorInterval)} = 65$$

The following commands can be used to set the Attribute ToleranceLimit :

- # hares -override <resource name> ToleranceLimit
- # hares -modify <resource name> ToleranceLimit 2

Note: Tune the ToleranceLimit value for Oracle applications such as ASM, ASM Diskgroup, Oracle Database, and Oracle netlistner, that are configured with both Oracle Restart feature and VCS.

Configuring Oracle ASM on the first node of the cluster

If you want to use ASM feature of Oracle 10g or Oracle 11g R1, then configure the Oracle ASM using the Database Configuration Assistant.

You need to configure the ASM only once, from the node on which you installed Oracle first.

Review the procedure to configure ASM for Oracle database.

To configure ASM for the Oracle 10g or Oracle 11g R1 database

- 1 Set the *DISPLAY* variable and start the Oracle Database Configuration Assistant as Oracle User.

```
# dbca
```

- 2 Read the Welcome screen, click **Next**.
- 3 In the Operations dialog box, select the **Configure Automatic Storage Management** option and click **Next**.
- 4 In the Database Templates dialog box, select a template to create the database and click **Next**.
- 5 In the Database Identification dialog box, enter or select a unique name and SID for the global database and click **Next**.
- 6 Follow the wizard instructions and select other options.

- 7 In the Storage Options dialog box, select the **Automatic Storage Option** and click **Next**.
- 8 Review the warning message and click **OK**. Make sure that the Oracle clustering daemon is running.

See [“Enabling the clustering daemon for ASM-managed database”](#) on page 52.
- 9 In the Create ASM Instance dialog box, specify the ASM SYS user password and click **Next**.
- 10 Click **OK** in the message box that appears.
- 11 In the ASM Disk Groups dialog box, choose to create a new disk group or add disks to an existing disk group.
- 12 In the Change Disk Discovery Path dialog box, enter the path name.

Depending on the disk device you want to use for ASM, enter the disk discovery path.
- 13 In the Create Disk Group dialog box, do the following:
 - Choose one of the Redundancy options.
 - In the Select Member Disks area, choose the **Show Candidates** option.
 - Select the disks against the **Disk Path** from the list to create the ASM disk group.
 - Click **OK**.
- 14 Follow the wizard instructions and complete configuring the Oracle ASM.

To configure ASM for the Oracle 11gR2 or 12c database

Note: Create disk groups using Oracle ASM Configuration Assistant only on one server in the VCS cluster. These disk groups will be available on the remaining servers in the VCS cluster because it hosts the database of the disk groups.

- 1 Set the *DISPLAY* variable and start the Oracle ASM Configuration Assistant as Oracle User.


```
# asmca
```
- 2 In the Create ASM Instance dialog box, specify the password for ASM administrator and click **Specify Disk Group** button.
- 3 In the Create Disk Group dialog box, choose to create a new disk group or add disks to an existing disk group.

- 4 In the Change Disk Discovery Path dialog box, enter the path name.
Depending on the disk device you want to use for ASM, enter the disk discovery path.
- 5 In the Create Disk Group dialog box, do the following:
 - Choose one of the Redundancy options.
 - In the Select Member Disks area, choose the **Show Candidates** option.
 - Select the disks against the **Disk Path** from the list to create the ASM disk group.
 - Click **OK**.
- 6 Follow the wizard instructions and complete configuring the Oracle ASM.

Note: After you successfully run Oracle ASM Configuration Assistant, run the `$GRID_HOME/bin/crsctl stat res -t -init` command to check if `ora.<DGname>.dg`, a new Oracle resource(s), is created.

- 7 For Oracle 11g R2 or 12c, copy the PFILE or the SPFILE for ASM instances from the ASM disk group to the local file system.
With Oracle 11g R2 or 12c, the VCS agent for ASMInst agent does not support PFILE or SPFILE for ASM instances on the ASM diskgroups.

Enabling the clustering daemon for ASM-managed database

If your database is ASM-managed, then you must make sure that the Oracle clustering daemon is enabled in each node of the cluster. If the CSSD is not enabled, perform the following procedure.

To enable the clustering daemon for ASM-managed database

- 1 Log in as superuser.
- 2 Run the following command to configure and start the clustering daemon.

For Oracle 10gR2 and 11gR1:

```
# $ORACLE_HOME/bin/localconfig add
```

For Oracle 11gR2 or 12c:

```
# $GRID_HOME/bin/crsctl start resource ora.cssd
```

- 3 Make sure the node has only one Oracle CSSD process running.

Enabling Oracle Cluster Synchronization Service daemon to start automatically

For Oracle 11g R2 or 12c, perform the following procedure on each node of the cluster.

The Oracle ASM instance resource faults if the required Oracle Cluster Synchronization Services process (occssd.bin) is not active. To prevent the resource from faulting, modify the AUTO_START attribute of the CSSD resource (ora.cssd). This configuration change enables the Oracle High Availability Services daemon (ohasd) to start the occssd.bin process as soon as the ohasd daemon starts on each node in the cluster.

Thus, when the occssd.bin process starts, it automatically starts up for the ASM instance. However, this could cause concurrency violation issues if the ASMInst resource is configured as a part of failover service group in the following scenario. Suppose a node goes down or reboots for some reason, all the resources on this node comes online on the second node. After the reboot, the ohasd process automatically starts the occssd.bin process, which automatically starts the ASM instance on node1 which leads to concurrency violation. To resolve this issue, recommends you to disable automatic startup of asm by running the following command:

```
# $GRID_HOME/bin/crsctl modify resource ora.asm -attr AUTO_START=never
```

Note: You must run the command to disable automatic startup of asm on all nodes in the VCS cluster that have Oracle Grid Infrastructure installed.

To enable Oracle Cluster Synchronization Service daemon to start automatically

- ◆ Set the AUTO_START attribute of the ora.cssd resource as follows:

```
# $GRID_HOME/bin/crsctl modify  
resource ora.cssd -attr AUTO_START=always
```

Note: You must perform this step on all nodes in the VCS cluster that have Oracle Grid Infrastructure installed.

Configuring and starting up ASM on remaining nodes for 11gR2 or 12c

This procedure is applicable when Oracle grid infrastructure is installed on all nodes and ASM is configured on the first node. See “[Configuring Oracle ASM on the first node of the cluster](#)” on page 50.

Configure the remaining nodes by following the procedure given below.

To create ASM and start ASM on remaining nodes:

- 1 Copy the `spfile` for ASM from the ASM Disks that is to be used on the remote node.

For example:

```
ASMCMD> spget +DATA/asm/asmparameterfile/<registry_file>

ASMCMD> spcopy +DATA/asm/asmparameterfile/<registry_file>
/u01/oraHome/dbs/spfileASM.ora
```

- 2 Stop the database and the ASM Disk group on the first node.
- 3 Copy the `spfile` from the first node to the remote node.
- 4 Copy `$ORACLE_BASE/admin/SID*` from first node to the remote node.
- 5 Add an asm instance on the remote node.

For example: `$/u01/product/11.2.0/grid/bin/srvctl add asm -p /u01/oraHome/dbs/spfileASM.ora`

- 6 Start the asm instance using the `srvctl` command.
- 7 Import the VxVM disk group on the remote node.

To configure database on remaining nodes that use ASM disk-groups (11gR2 or 12c):

- 1 Run the `srvctl add database` command to register the Oracle Database(s) that is running on the nodes.

Note: Use the credentials of the Oracle software owner to register the database. For more information on the complete list of parameters, refer to the Oracle documentation.

- 2 Login to the Oracle ASM instance running on the remote node.
- 3 Run the `SQL> ALTER disk group <DGname> mount` command to mount the Oracle ASM disk groups.

- 4 Repeat Step 3 to mount all the required disk groups.
- 5 Run the `$GRID_HOME/bin/crsctl stat res -t -init` command to check if the disk groups are auto-registered to OHASD.

The output displays `ora.<DGname>.dg` for the registered disk groups.
- 6 Run the `$GRID_HOME/bin/srvctl modify database -a <diskgroup_list>` command to add the Oracle ASM disk groups as a dependency to the database(s).

Installing Oracle binaries on the first node of the cluster

Install a version of Oracle that the Cluster Server agent for Oracle supports.

Note: If you want to use Oracle ASM feature, do not disable the Cluster Synchronization Services.

To install Oracle 10g for VCS

- 1 Insert the Oracle CD.
- 2 Set the `DISPLAY` variable, mount the CD drive, and run the Oracle installer as Oracle User.


```
# /mnt/cdrom/runInstaller
```
- 3 Read the Welcome screen and click **Next**.
- 4 If you perform the Oracle installation for the first time, the Inventory Location dialog box is displayed. Specify the location of your base directory and click **OK**. Do one of the following:
 - If you plan the `$ORACLE_HOME` to be on local disks in your installation, specify a location on the local disk.
 - If you plan the `$ORACLE_HOME` to be on shared disks in your installation, specify a location on the shared disk.
- 5 In the File Locations dialog box, enter or select a name and a path for the Oracle Home and click **Next**.
- 6 In the Installation Types dialog box, select your installation type and click **Next**.
- 7 In the Select Database Configuration dialog box, select the **Do not create a starter database** option and click **Next**.

- 8 In the Summary dialog box, review your selections. Click **Back** to make changes. Otherwise click **Install**.
- 9 During the installation, the Oracle installer displays a message asking you to run the script `$ORACLE_BASE/oraInventory/orainstRoot.sh`. Run the script and click **Continue** to return to the Oracle installer.
- 10 Click **Next** after the installation is complete.
- 11 In the End of Installation dialog box, click **Exit**.
- 12 Disable the Oracle clustering daemon.
See [“Disabling the clustering daemon for Oracle 10g or later”](#) on page 57.
If you want to use ASM, skip this step and proceed to configuring the Oracle ASM.
See [“Configuring Oracle ASM on the first node of the cluster”](#) on page 50.
- 13 Proceed to configure the database on shared disks.
See [“Configuring the Oracle database”](#) on page 58.

To install Oracle 11g for VCS

- 1 Insert the Oracle CD.
- 2 Set the `DISPLAY` variable, mount the CD drive, and run the Oracle installer as Oracle User.

```
# /mnt/cdrom/runInstaller
```
- 3 In the Select Installation Method dialog box, choose the **Advanced Installation** method and click **Next**.
- 4 If you perform the Oracle installation for the first time, the Specify Inventory directory and credentials dialog box is displayed. Specify the location of your inventory directory and click **Next**. Do one of the following:
 - If you plan the `$ORACLE_HOME` to be on local disks in your installation, specify a location on the local disk.
 - If you plan the `$ORACLE_HOME` to be on shared disks in your installation, specify a location on the shared disk.
- 5 Follow the wizard instructions and select other options
- 6 In the Create Database dialog box, choose **Install database Software only**.
- 7 In the Summary dialog box, review your selections. Click **Back** to make changes. Otherwise click **Install**.

- 8 During the installation, the Oracle installer displays a message asking you to run some configuration scripts. Run the scripts and click **OK** to return to the Oracle installer.
- 9 Click **Next** after the installation is complete.
- 10 In the End of Installation dialog box, click **Exit**.
- 11 Disable the Oracle clustering daemon.
See [“Disabling the clustering daemon for Oracle 10g or later”](#) on page 57.
If you want to use ASM, skip this step and proceed to configuring the Oracle ASM.
See [“Configuring Oracle ASM on the first node of the cluster”](#) on page 50.
- 12 Proceed to configure the database on shared disks.
See [“Configuring the Oracle database”](#) on page 58.

Disabling the clustering daemon for Oracle 10g or later

If you installed Oracle binaries for Oracle versions 10g or later on shared disks, you must disable the Oracle clustering daemon.

Warning: If you want to use ASM feature, then do not disable the Oracle clustering daemon.

Oracle versions 10g and later provide a clustering daemon called Oracle Cluster Synchronization Service Daemon (CSSD). If the Oracle binary files are on a shared storage, the `init` command to start the daemon may result in an error. Because a VCS cluster for Oracle does not require this daemon, Veritas recommends you to disable the daemon.

To disable the daemon on Linux

- ◆ Remove the following line from the `/etc/inittab` file on the node from where you ran the Oracle installation wizard:

```
h1:23:respawn:/etc/init.d/init.cssd run >/dev/null 2>&1 >  
</dev/null
```

Configuring the Oracle database

Configure an Oracle database on shared disks using the Database Configuration Assistant. You need to configure the database only once, from the node on which you installed Oracle first.

If you configured ASM for Oracle database storage, configure the Oracle database on ASM disk groups.

Note: When creating a database in Oracle 12c, uncheck the **Create as a container database** checkbox. Currently VCS 6.1 supports Oracle 12c features that were supported in previous versions.

To configure the Oracle database

- 1 Set the *DISPLAY* variable and start the Oracle Database Configuration Assistant as Oracle User.

```
# dbca
```

- 2 Read the Welcome screen, click **Next**.
- 3 In the Operations dialog box, select the **Create a database** option and click **Next**.
- 4 In the Database Templates dialog box, select a template to create the database and click **Next**.
- 5 In the Database Identification dialog box, enter or select a unique name and SID for the global database and click **Next**.
- 6 Follow the wizard instructions and select other options.
- 7 In the Storage Options dialog box, select the **File System** or **Raw Devices** option and click **Next**.

If you configured ASM, select **Automatic Storage Management (ASM)** and click **Next**.

- 8 If you configured ASM, in the ASM Disk Groups dialog box, select the ASM disk group for the Oracle database storage and click **Next**.
- 9 In the Database File Locations dialog box, specify a location on the shared disk for the database files to be created and click **Next**.

If you configured ASM, choose **Use Oracle-Managed Files**, specify the database area, and click **Next**.

- 10 In the Recovery Configuration dialog box, specify the **Flash Recovery Area** on the shared disk and click **Next**.
If you configured ASM, specify the appropriate **Flash Recovery Area** location and click **Next**.
- 11 If you use Oracle 11g, in the Security Settings dialog box, choose the recommended security setting. Choose the **Keep the enhanced 11g default security settings** option and click **Next**.
- 12 Follow the wizard instructions and complete creating the database.

Copying the \$ORACLE_BASE/admin/SID directory

Follow the instruction on each node in the cluster, only if the \$ORACLE_HOME is on local disks.

To copy the \$ORACLE_BASE/admin/SID directory

- ◆ Remote copy \$ORACLE_BASE/admin/SID on the node where you first configured the database instance to the local disk (\$ORACLE_BASE/admin) on other nodes in the cluster.

The variable *SID* represents the database instance.

Refer to Oracle's documentation on OFA Structure for more information.

Copying the Oracle ASM initialization parameter file

If your Oracle is ASM-managed, then perform the procedure on each node in the cluster.

To copy the ASM initialization parameter file

- 1 For Oracle 10g R2 and 11g R1:
Copy the *initSID.ora* file from the \$ORACLE_HOME/dbs directory on the node where you first installed Oracle ASM on each node of the cluster.
- 2 For Oracle 11g R2 or 12c, copy the *initSID.ora* file from the \$GRID_HOME/dbs directory on the node where you first installed Oracle ASM on each node of the cluster.
where *SID* is the ASM instance identifier.

Verifying access to the Oracle database

You can verify access to the database by running an SQL query. Depending on the location of your `$ORACLE_HOME`, perform one of the following procedures.

Upon completion, the Oracle database is started from the node in the cluster that you last performed this procedure.

To verify access to the Oracle database

- ◆ Depending on the location of your `$ORACLE_HOME`, perform one of the following procedures:

For `$ORACLE_HOME` on shared disks, start the Oracle database.

Verify that you can access the shared database from each node. Perform the steps from each node that would be a part of the Oracle cluster.

To start the Oracle database (for `$ORACLE_HOME` on shared disks)

- 1 If the Oracle binaries are mounted on any other node, unmount it.
- 2 If the data files on shared disk are mounted on any node, unmount it.
- 3 Mount the Oracle binaries and data files.
- 4 Start the database.

For `$ORACLE_HOME` on local disks, start the Oracle database.

Verify that you can access the database individually on each node that would be a part of the Oracle cluster.

To start the Oracle database (for `$ORACLE_HOME` on local disks)

- 1 Make sure you have installed Oracle binaries on the node.
- 2 If the data files on shared disk are mounted on any node, unmount it.
- 3 Mount the data files.
- 4 Start the database.

For ASM, the \$ORACLE_HOME is on local disks and the database is configured on ASM disk groups.

Start the ASM-managed Oracle database.

Verify that you can access the database individually on each node that would be a part of the Oracle cluster.

To start the Oracle database (for \$ORACLE_HOME on local disks)

- 1** Make sure you have installed Oracle binaries on the node.
- 2** If you use VxVM disks as ASM disks, do the following:
 - If the ASM disk groups are mounted on any node, do the following:
 - Unmount the ASM disk groups from the ASM instance.
 - Stop the ASM instance.
 - Stop the VxVM volumes.
 - Deport the VxVM disk groups.
 - Import the VxVM disk groups on this node and start the volumes.
 - Start the ASM instance.
 - Mount the ASM disk groups.
- 3** If you use CVM volumes for ASM, do the following:
 - Unmount the ASM disk group from the ASM instance on the node where the disk group is mounted.
 - Mount the ASM disk groups on the other node.
- 4** If you use raw disks as ASM disks, do the following:
 - If the ASM disk groups are mounted on any node, unmount it.
 - If the ASM instance is not running on this node, then start the ASM instance.
 - Mount the ASM disk groups.
- 5** Start the database.

Installing and removing the agent for Oracle

This chapter includes the following topics:

- [Before you install or upgrade the agent for Oracle](#)
- [Installing the VCS agent for Oracle software](#)
- [Upgrading the VCS agent for Oracle](#)
- [Disabling the Cluster Server agent for Oracle](#)
- [Removing the Cluster Server agent for Oracle](#)

Before you install or upgrade the agent for Oracle

Meet the following prerequisites to install or upgrade the Cluster Server agent for Oracle:

- Make sure that VCS is installed in the cluster.
- Verify that the Oracle Server for Linux and the add-on Listener are installed and configured.
See Oracle documentation.
See [“About installing Oracle in a VCS environment”](#) on page 36.

Installing the VCS agent for Oracle software

The agent binaries for Cluster Server agent for Oracle are part of VRTSvcsea RPM. The VRTSvcsea RPM is already installed if you chose to install all RPMs or recommended RPMs during VCS installation. You must manually install the VRTSvcsea RPM if you installed minimal RPMs during VCS installation.

You can install the Cluster Server agent for Oracle from the product disc. You must install the Oracle enterprise agent on all nodes that will host the Oracle service group.

To install the agent on a Linux node

- 1 Log in as superuser.
- 2 Make sure that the agent is not already installed. Navigate to `/opt/VRTSagents/ha/bin` and list the contents of the directory.
If the command returns the VCS agents for Oracle, you can skip this procedure.
- 3 Insert the software disc into a drive that is connected to the node.
The software automatically mounts the disc as `/mnt/cdrom`.
- 4 If the disc does not automatically mount, then you must mount the disc manually. For example:

```
# mount -t iso9660 -o ro /dev/cdrom /mnt/cdrom
```

- 5 Install the agent software.
 - For RHEL5, RHEL6, OEL5, or OEL6 on x86_64:

```
# cd /mnt/cdrom/dist_x86_64/rpms
```

where *dist* is rhel5 or rhel6.

```
# rpm -i VRTSvcsea-7.1.0.000-GA_dist.x86_64.rpm
```

- For SLES11 on x86_64:

```
# cd /mnt/cdrom/dist_x86_64/rpms
```

```
# rpm -i VRTSvcsea-7.1.0.000-GA_dist.x86_64.rpm
```

where *dist* is SLES11.

The VRTSvcsea RPM includes the agent binaries for Oracle, Netlsnr, ASMInst, and ASMDG agents. The RPM also installs the VCS agents for DB2 and Sybase.

- 6 Repeat step 1 through step 5 on each node that will become part of the Oracle service group.

Upgrading the VCS agent for Oracle

The agent binaries for Cluster Server agent for Oracle are part of VRTSvcsea RPM. The VRTSvcsea RPM is already upgraded if you chose to upgrade to VCS 7.1 on a cluster which has a previous version of Cluster Server agent for Oracle installed.

Before you upgrade the agent, make sure you meet the prerequisites to upgrade the Cluster Server agent for Oracle.

See [“Before you install or upgrade the agent for Oracle”](#) on page 62.

Warning: The agent RPM VRTSvcsea includes the VCS agents for Oracle, Sybase, and DB2. So, the procedure to upgrade the VCS agent for Oracle upgrades all the other agents also. Make sure that you have also met the upgrade prerequisites for the other agents that are part of the VRTSvcsea RPM.

If you had installed Cluster Server agent for Oracle that supported VCS 6.0, 6.0RP1, 6.0PR1, 6.0.1, 6.0.2, 6.0.3, 6.0.4, 6.0.5, 6.1, 6.1.1, 6.2, 6.2.1, and 7.0.1 on your VCS cluster, you can now upgrade the agents for use with VCS 7.1.

Upgrading VCS agent for Oracle to use with VCS 7.1

Perform the following steps on each node of the VCS cluster.

To upgrade VCS agent for Oracle to use with VCS 7.1

- 1 Stop VCS locally.

```
# hastop -local -force
```

- 2 Remove the earlier version of the VCS agent for Oracle.

```
# rpm -e VRTSvcsea
```

- 3 Install the VCS 7.1 agent for Oracle.

See [“Installing the VCS agent for Oracle software”](#) on page 62.

- 4 Make sure that you have copied the OracleTypes.cf file from the /etc/VRTSagents/ha/conf/Oracle directory to /etc/VRTSvcs/conf/config directory.

Make sure to update the newly copied OracleTypes.cf file with all the type-level changes that you had made to the older OracleTypes.cf file.

For example, if you had changed the value of the MonitorInterval attribute from the default 60 to 120 seconds, the OracleTypes.cf file gets updated. You must apply these type-level changes manually to the newly copied OracleTypes.cf file.
- 5 If you use ASM for database storage, then make sure that you have copied the file OracleASMTypes.cf from the directory /etc/VRTSagents/ha/conf/OracleASM to /etc/VRTSvcs/conf/config directory.

Make sure to update the newly copied OracleASMTypes.cf file with all the type-level changes that you had made to the older OracleASMTypes.cf file.
- 6 If you had enabled detail monitoring for the Oracle agent in the previous version, set the value of the LevelTwoMonitorFreq attribute to the same value of that of the DetailMonitor attribute.

Note: The DetailMonitor attribute is deprecated in VCS 6.0 and later.

- 7 Restart the VCS engine.

```
# hastart
```

Disabling the Cluster Server agent for Oracle

To disable the Cluster Server agent for Oracle, 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 the agent

- 1 To remove a system from the service group's SystemList, check if the service group is online:

```
# hagrps -state service_group -sys system_name
```

- 2 If the service group is online, take it offline. Use one of the following commands:
 - To take the service group offline on one node and online it on another node, you can use the -switch option:

```
# hagrps -switch service_group -to system_name
```

- To take the service group offline without bringing it online on any other node in the cluster, enter:

```
# hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the node:

```
# haagent -stop Oracle -sys system_name
```

- 4 When you get the message "Please look for messages in the log file," check the file `/var/VRTSvcsllog/engine_A.log` for a message confirming the agent has stopped.

You can also use the `ps` command to confirm the agent is stopped.

- 5 You can now remove the service group, the resource type, or both from the VCS configuration after disabling the agent on all nodes.

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

Removing the Cluster Server agent for Oracle

Make sure you disabled the agents on all nodes before you remove the service group, the resource type, or both from the VCS configuration.

Warning: The agent RPM `VRTSvcsea` includes the VCS agents for Oracle, Sybase, and DB2. So, the following procedure to remove the VCS agent for Oracle removes all the other agents also.

Perform the procedure to remove the agents on each node in the cluster.

To remove the agent on a node

- 1 Stop all the resources on all the nodes.

```
# hastop -local
```

- 2 Remove the `VRTSvcsea` RPM. Answer prompts accordingly.

```
# rpm -e VRTSvcsea
```

Configuring VCS service groups for Oracle

This chapter includes the following topics:

- [About configuring a service group for Oracle](#)
- [Configuring Oracle instances in VCS](#)
- [Before you configure the VCS service group for Oracle](#)
- [Configuring the VCS service group for Oracle](#)
- [Setting up detail monitoring for VCS agents for Oracle](#)
- [Enabling and disabling intelligent resource monitoring for agents manually](#)

About configuring a service group for Oracle

Configuring the Oracle service group involves creating the Oracle service group, its resources, and defining attribute values for the configured resources. You must have administrator privileges to create and configure a service group.

You can configure a VCS service group for Oracle using one of the following:

- Veritas InfoScale Operations Manager
- The Cluster Manager (Java console)
- The command-line

Note: If you use Oracle ASM feature, you must also configure ASMinst resources and ASMDG resources to keep the ASM instance highly available. You must use the command-line to configure these resources.

Configuring Oracle instances in VCS

You can set up Oracle in different ways in a VCS environment. Configuring Oracle for VCS involves configuring the Oracle files `listener.ora` and `tnsnames.ora` as per VCS requirements.

Table 4-1 Oracle configurations in VCS

Configuration type	Reference
Single Oracle instance configuration	See “Configuring a single Oracle instance in VCS” on page 68.
Multiple Oracle instances (single listener) configuration	See “Configuring multiple Oracle instances (single listener) in VCS” on page 70.
Multiple Oracle instances (multiple listeners) configuration	See “Configuring multiple Oracle instances (multiple listeners) in VCS” on page 71.
Configuration with shared server support	See “Configuring an Oracle instance with shared server support in VCS” on page 73.

See [“Best practices for multiple Oracle instance configurations in a VCS environment”](#) on page 205.

Note: Make sure that the Oracle files `$TNS_ADMIN/listener.ora` and `$TNS_ADMIN/tnsnames.ora` are identical across the cluster nodes. If the files are not identical, remote copy the files from a local node to the other nodes in the cluster.

Configuring a single Oracle instance in VCS

Review the resource dependency graph and sample configuration of service groups involving a single Oracle instance.

See [“Sample single Oracle instance configuration”](#) on page 165.

To configure an Oracle instance in VCS

- 1 Review the Oracle and Netlsnr resource types and their attribute definitions. See [“About the resource type and attribute definitions”](#) on page 142.
- 2 Configure the Oracle file tnsnames.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

For clients to connect to the failover instance, in the file tnsnames.ora located at \$TNS_ADMIN, change the host name for all TCP protocol address databases to the virtual IP address for the service group.

The following example assumes that the host name for the database is set to oraprod, which represents the virtual IP address for the service group.

```
prod =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = prod)
    )
  )
```

- 3 Configure the Oracle file listener.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

In the file listener.ora located at \$TNS_ADMIN, edit the "Host=" line in the ADDRESS_LIST section and add the name of the high availability address for the service group, in this case, oraprod.

```
LISTENER_PROD =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
      )
    )
  )
```

- 4 Create the VCS service groups for Oracle. See [“Configuring the VCS service group for Oracle”](#) on page 77.
- 5 Bring the Oracle service group online. See [“Bringing the service group online”](#) on page 120.

Configuring multiple Oracle instances (single listener) in VCS

Review the resource dependency graph and sample configuration of a service group involving multiple Oracle instances.

See [“Sample multiple Oracle instances \(single listener\) configuration”](#) on page 168.

To configure multiple Oracle instances (single listener)

- 1 Review the Oracle and Netlsnr resource types and their attribute definitions. See [“About the resource type and attribute definitions”](#) on page 142.
- 2 Configure the Oracle file tnsnames.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

For clients to connect to the failover instance, in the file tnsnames.ora located at \$TNS_ADMIN, change the host name for all TCP protocol address databases to the virtual IP address for the service group.

The following example assumes that the host name for the database is set to oraprod, which represents the virtual IP address for the service group.

```
prod =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = prod)
    )
  )
)
mktg =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oramktg) (PORT = 1522))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = mktg)
    )
  )
)
```

- 3 Configure the Oracle file listener.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

In the file listener.ora located at \$TNS_ADMIN, edit the "Host=" line in the ADDRESS_LIST section and add the name of the high availability address for the service group, in this case, oraprod.

```
LISTENER_ORACLE =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS_LIST =  
        (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))  
        (ADDRESS = (PROTOCOL = TCP) (HOST = oramktg) (PORT = 1522))  
      )  
    )  
  )
```

- 4 Create the VCS service groups for Oracle.
See ["Configuring the VCS service group for Oracle"](#) on page 77.
- 5 Bring the Oracle service group online.
See ["Bringing the service group online"](#) on page 120.

Configuring multiple Oracle instances (multiple listeners) in VCS

Review the resource dependency graph and sample configuration of a service group involving multiple Oracle instance.

See ["Sample multiple instance \(multiple listeners\) configuration"](#) on page 175.

To configure multiple Oracle instances (multiple listeners)

- 1 Review the Oracle and Netlsnr resource types and their attribute definitions. See [“About the resource type and attribute definitions”](#) on page 142.
- 2 Configure the Oracle file tnsnames.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

For clients to connect to the failover instance, in the file tnsnames.ora located at \$TNS_ADMIN, change the host name for all TCP protocol address databases to the virtual IP address for the service group.

```
prod =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = prod)
    )
  )
)
mktg =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oramktg) (PORT = 1522))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = mktg)
    )
  )
)
```


- 3 Configure the Oracle file listener.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

In the file listener.ora, create independent configurations for each listener.

```
LISTENER_PROD =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
      )
    )
  )

LISTENER_MKTG =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = oramktg) (PORT = 1522))
      )
    )
  )
```

- 4 Create the VCS service groups for Oracle.
See [“Configuring the VCS service group for Oracle”](#) on page 77.
- 5 Bring the Oracle service group online.
See [“Bringing the service group online”](#) on page 120.

Configuring an Oracle instance with shared server support in VCS

Review the resource dependency graph and sample configuration.

See [“Sample Oracle configuration with shared server support”](#) on page 179.

To configure Oracle with shared server support

- 1 Review the Oracle and Netlsnr resource types and their attribute definitions. See [“About the resource type and attribute definitions”](#) on page 142.
- 2 Configure the Oracle file tnsnames.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

For clients to connect to the failover instance, in the file tnsnames.ora located at \$TNS_ADMIN, change the host name for all TCP protocol address databases to the virtual IP address for the service group.

The following example assumes that the host name for the database is set to oraprod, which represents the virtual IP address for the service group.

```
prod =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = prod)
    )
  )
```

- 3 Configure the Oracle file listener.ora as per VCS requirements. The changes required in the file depends on your Oracle configuration.

In the file listener.ora located at \$TNS_ADMIN, edit the "Host=" line in the ADDRESS_LIST section and add the name of the high availability address for the service group, in this case, oraprod.

```
LISTENER_PROD =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP) (HOST = oraprod) (PORT = 1522))
      )
    )
  )
```

- 4 In the initialization parameter file, for the dispatchers parameter, set the host name for all TCP protocol address dispatchers to the virtual IP address for the service group.

In the following example, the host name for the dispatcher is set to oraprod.

```
dispatchers =
  "(ADDRESS = (PROTOCOL = TCP) (HOST = oraprod))
  (MUL=ON) (TICK=15) (POOL=true)"
```

- 5 In the initialization parameter file, set the LOCAL_LISTENER attribute to a listener alias, so that the Oracle PMON process registers information with the listener.

```
local_listener=listener_alias
```

This listener alias gets appended by the default domain set in the file sqlnet.ora.

- 6 In the file tnsnames.ora, create an entry for the listener alias to resolve the address information. In the address parameter, set the host name to the virtual IP address for the service group, without the CONNECT_DATA portion of the connect descriptor.

In the following example, the listener *listener_alias* uses TCP/IP port 1521 and the virtual IP address for the service group is oraprod.

```
listener_alias=
  (address = (PROTOCOL = TCP) (HOST = oraprod))
  (PORT= 1521))
```

- 7 Create the VCS service groups for Oracle.
See [“Configuring the VCS service group for Oracle”](#) on page 77.
- 8 Bring the Oracle service group online.
See [“Bringing the service group online”](#) on page 120.

Before you configure the VCS service group for Oracle

Before you configure the Oracle service group, you must:

- Verify that VCS is installed and configured on all nodes in the cluster where you will configure the service group.
Refer to the *Cluster Server Installation Guide* for more information.

- Verify that Oracle is installed and configured identically on all nodes in the cluster. See [“About installing Oracle in a VCS environment”](#) on page 36.
- Verify that the Cluster Server agent for Oracle is installed on all nodes in the cluster. See [“Before you install or upgrade the agent for Oracle”](#) on page 62.
- Verify that the type definition for Cluster Server agent for Oracle is imported into the VCS engine. See [“Importing the type definition files for Cluster Server agent for Oracle”](#) on page 76.
- If the Oracle database is ASM-managed, make sure that the Oracle clustering daemon is enabled. Else, the OCSSD daemon must be disabled.

Importing the type definition files for Cluster Server agent for Oracle

Perform this procedure if you installed the VRTSvcsea RPM manually. The installer program imports the type definition files if you chose all RPMs or recommended RPMs during VCS installation.

Before configuring the Cluster Server agent for Oracle, you must import the OracleTypes.cf file to the VCS engine.

If you want to configure the ASM agents, you must import the OracleASMTypes.cf file. Import the OracleASMTypes.cf file using the command-line.

To import using the Cluster Manager

- 1 Start Cluster Manager and log on to the cluster.
- 2 From the Cluster Explorer’s **File** menu, choose **Import Types**.
- 3 In the Import Types dialog box, select the `/etc/VRTSagents/ha/conf/Oracle/OracleTypes.cf` file.
- 4 Click **Import**.
- 5 Save the configuration.

To import using the command line

- 1 Log in to the system as superuser.
- 2 Ensure that all changes to the existing configuration have been saved and that further changes are prevented while you modify `main.cf`:

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

- 3 To ensure that VCS is not running while you edit main.cf, stop the VCS engine on all nodes and leave the resources available:

```
# hastop -all -force
```

- 4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.orig
```

- 5 Copy the Oracle types configuration file into place:

```
# cp /etc/VRTSagents/ha/conf/Oracle/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf
```

If you want to configure the ASM agents, do the following:

```
# cp /etc/VRTSagents/ha/conf/OracleASM/OracleASMTypes.cf \
/etc/VRTSvcs/conf/config/OracleASMTypes.cf
```

- 6 Edit the main.cf file to include the OracleTypes.cf file:

```
include "OracleTypes.cf"
```

If you want to configure the ASM agents, you must also include the OracleASMTypes.cf file:

```
include "OracleASMTypes.cf"
```

- 7 Verify the configuration.

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

- 8 Start the VCS engine.

```
# hastart
```

Configuring the VCS service group for Oracle

You can configure Oracle in a VCS environment in one of the ways that VCS supports.

See [“Configuring Oracle instances in VCS”](#) on page 68.

Table 4-2 lists the interfaces available to configure the VCS agent for Oracle after you have configured Oracle and installed the agent.

Table 4-2 Interfaces to configure VCS agent for Oracle

Configuration method	Reference
The Cluster Manager (Java console)	See “Configuring the VCS service group for Oracle using Cluster Manager (Java console)” on page 78.
The command-line	See “Configuring the VCS service group for Oracle using the command-line” on page 80.
Veritas InfoScale Operations Manager	Refer to the Veritas InfoScale Operations Manager documentation for more details.

Review the following to configure the service group:

- Sample configuration files and resource dependency graphs of the Oracle service group.
 See [“About the sample configurations for Oracle enterprise agent”](#) on page 164.
- Resource type and the attribute definitions of the Oracle and Netlshr agents.
 For ASM, resource type and the attribute definitions of the ASMinst and ASMDG agents.
 See [“About the resource type and attribute definitions”](#) on page 142.

Configuring the VCS service group for Oracle using Cluster Manager (Java console)

A template for the Oracle resource groups is automatically installed when you install the Oracle enterprise agent. Using the VCS Cluster Manager, you can view the template, which displays the Oracle service group, its resources and their attributes. You can dynamically modify the attribute values as necessary for your configuration.

Configuration tasks involve:

- Creating a new service group
- Editing the resource attributes

To create a new service group

- 1 Make sure that the Oracle type definition file OracleTypes.cf is imported in your configuration.
- 2 Launch the Cluster Configuration wizard. Do one of the following:
 - From the Cluster Explorer menu, select **Tools > Configuration Wizard**.

- If no service groups exist on the system, **Cluster Explorer** prompts you to launch the **Cluster Configuration wizard**. Click **Yes** when prompted.

The Loading Templates Information window appears, and launches the wizard.

- 3 Read the information on the Welcome screen and click **Next**.
- 4 Specify the name of the service group and the target systems on which it is to be configured.
 - Enter the name of the service group.
 - On the **Available Systems** box, select the systems on which to configure the service group.
 - Click the right arrow to move the selected systems to the **Systems for Service Group** box. To remove a system from the box, select the system and click the left arrow.
 - Indicate the order in which the systems will fail over (priority). System priority is numbered sequentially, with 1 denoting the system that starts first following a failover.
 - Click **Next**.
- 5 In the Would you like to use a template to configure the service group? dialog box, click **Next**.
- 6 In the Select a template to configure the service group dialog box, select the OracleGroup template on which to base the service group.

If applicable, a window opens notifying you that names of some resources within the new service group are already in use. This window also includes a list of alternative names. You can select an alternative name, or modify the name. Click **Next** after resolving the name clashes.
- 7 Click **Next** to create the service group based on the selected template.

A window opens indicating that commands are being sent to add the group, its resources, and the attributes and dependencies specified in the template.

A progress indicator displays the percentage of the commands fired.

The actual commands are displayed at the top of the indicator.
- 8 Click **Next** when prompted that the service group has been successfully created.

A window opens listing the service group's resources and their associated attributes.
- 9 Click **Finish** to accept the default values and complete the configuration.

You can modify the default values of the resources according to your specific configuration requirements, as instructed in the following section.

To edit the resource attributes

- 1 Select the resource from the list on the left pane.
The resource's attributes appear in the right pane.
- 2 Select the attribute to be modified.
See [“About the resource type and attribute definitions”](#) on page 142.
- 3 Click the edit icon in the **Edit** column.
- 4 In the Edit Attribute dialog box, enter the attribute values. To modify the scope of the attribute, click the option buttons for **Global** or **Local**.
- 5 Click **OK**.
- 6 Repeat the procedure for each resource and its attributes and click **Finish**.
Edit the attributes for all the resources according to your configuration.
- 7 Follow the wizard instructions to complete the configuration.
- 8 Click **Finish** to quit the wizard.

Configuring the VCS service group for Oracle using the command-line

The Cluster Server agent for Oracle comes with a sample configuration file that can be used as reference to directly modify your present configuration file. When you use this method, you must stop and restart VCS to implement the configuration.

The Cluster Server agent for Oracle also includes sample configuration files for ASM configurations in a VCS environment.

See [“About the sample configurations for Oracle enterprise agent”](#) on page 164.

To configure the service group using the sample configuration file

- 1 Log in as superuser on the system where oracle is running.
- 2 Ensure that all changes to the existing configuration have been saved and that further changes are prevented while you modify main.cf:

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

- 3 To ensure that VCS is not running while you edit main.cf, issue the following command to stop the VCS engine on all systems and leave the resources available:

```
# hastop -all -force
```


4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.orig
```

5 Edit the main.cf file.

You can use the file `/etc/VRTSagents/ha/conf/Oracle/sample_main.cf` for reference.

Do the following:

- Create the Oracle and Netlsnr resources.
For ASM instances and disk groups, create the ASMinst and ASMDG resources.
- Edit the default attributes to match the parameters in your configuration.
See [“About the resource type and attribute definitions”](#) on page 142.
- Assign dependencies to the newly created resources.
See [“About the sample configurations for Oracle enterprise agent”](#) on page 164.
See the *Cluster Server Administrator’s Guide* for more information on assigning dependencies.

6 Save and close the file.**7** Verify the syntax of the file `/etc/VRTSvcs/conf/config/main.cf`:

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

8 Start the VCS engine:

```
# hastart
```

9 Verify that all Oracle service group resources are brought online:

```
# hagrps -state
```

10 Take the service group offline and verify that all resources are stopped:

```
# hagrps -offline service_group -sys system_name
# hagrps -state
```

- 11** Bring the service group online again and verify that all resources are available:

```
# hagrps -online service_group -sys system_name
# hagrps -state
```

- 12** Start the VCS engine on all remaining nodes in the cluster:

```
# hstart
```

- 13** (Optional) Switch the Oracle service group to one of the cluster nodes to verify that it has been configured correctly for failover:

```
# hagrps -switch service_group -to any
```

Verify that all Oracle service group resources are brought online on the system:

```
# hagrps -state
```

- 14** On all the nodes, look at the following log files for any errors or status:

```
/var/VRTSvcs/log/engine_A.log
/var/VRTSvcs/log/Oracle_A.log
/var/VRTSvcs/log/Netlsnr_A.log

/var/VRTSvcs/log/ASMinst_A.log
/var/VRTSvcs/log/ASMDG_A.log
```

Encrypting Oracle database user and listener passwords

VCS provides a utility to encrypt database user passwords and listener passwords. You must encrypt the Pword attribute in the Oracle agent and the LsnrPwd attribute in the Netlsnr agent before you configure these attributes.

See the `vcscrypt` manual page.

The `vcscrypt` utility also allows you to encrypt the agent passwords using a security key. The security key supports AES (Advanced Encryption Standard) encryption which creates a more secure password for the agent.

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

Note: You need not encrypt passwords if you use the VCS Cluster Manager (Java Console) to configure attributes.

The user passwords that are used for detail monitoring of the Oracle database are encrypted. The listener password that is considered for querying the status of the listener and stopping the listener is also encrypted.

Oracle provides the option of storing the listener password in the listener.ora file, in both clear text and encrypted formats. Irrespective of the format in which the password is stored in Oracle, you must encrypt the password using the `vcseencrypt` utility before you configure the `LsnrPwd` attribute.

If you encrypted the listener password using the Oracle `lsnrctl` utility, do the following depending on the Oracle version you use:

- Pass the clear text password to the `vcseencrypt` utility.
For example, if the listener password is “oracle” pass this clear text password to the `vcseencrypt` utility as follows:

```
# vcseencrypt -agent oracle
```

The `vcseencrypt` utility displays the encrypted password. For example:
bnqPdnCneNnnG

To encrypt passwords

- 1 From the path `$VCS_HOME/bin/`, run the `vcseencrypt` utility.

- Type the following command:

```
# vcseencrypt -agent
```

- Enter the password and confirm it by entering it again. Press Enter.

```
# Enter New Password:
```

```
# Enter Again:
```

- 2 Review as the utility encrypts the password and displays the encrypted password.
- 3 Enter this encrypted password as the value for the attribute.
- 4 Copy the encrypted password for future reference.

Setting up detail monitoring for VCS agents for Oracle

The Cluster Server agent for Oracle provides two levels of application monitoring: primary (basic monitoring) and secondary (detail monitoring).

- In the basic monitoring mode, the agent monitors the Oracle processes to verify that they are continuously active.

Note: The agent for Oracle ASM does only basic monitoring for ASM instances.

- In the detail monitoring mode, the agent executes the script defined in the attribute `MonScript` of the Oracle and the Netlsnr resources. If the script successfully executes, the agent considers the resource available. You can customize the default script according to your configuration.

See [“Monitor options for the Oracle agent in traditional database and container database”](#) on page 20.

You can use the agent’s detail monitoring capability to monitor the status of a database and listener and increase the confidence in their availability. Before setting up detail monitoring, you must have the agent running satisfactorily at the basic level of monitoring.

Note: Disable detail monitoring before undertaking any database maintenance that involves disabling database access to external users.

Setting up detail monitoring for Oracle

Detail monitoring for an Oracle resource verifies whether a database is ready for transactions by performing an update transaction against a table within the database. The update action is taken by the two scripts, `SqlTest.pl` and `SimpleTest.pl`, provided with the Cluster Server agent for Oracle. The scripts are available under the directory `/opt/VRTSagents/ha/bin/Oracle/`. Both scripts update the timestamp to monitor the database.

The `SqlTest.pl` script checks whether the database is open before updating the timestamp. If the database is found to be in restricted mode, quiesced mode, or suspended mode, the monitor returns success. In such a case, only basic monitoring occurs. The `SimpleTest.pl` script does not perform database checks but only issues update statements against the table.

Before enabling detail monitoring for Oracle, you must create a test table (with a timestamp) in the Oracle database. The agent uses this test table for internal purposes. recommends that you do not perform any other transaction on the test table. The detail monitor script, `MonScript`, must exist and have execute permission for root. You can use a custom monitor script, or the scripts provided with the agent. In the monitor script, the return code 100 indicates failure. Return codes from 101 to 110 indicate success.

The example to set up detail monitoring, based on the use of the supplied script, shows how to create and test a table for use by detail monitoring, and how to enable detail monitoring.

To set up detail monitoring for Oracle

- 1 Make the VCS configuration writable:

```
# haconf -makerw
```

- 2 Freeze the service group to avoid automated actions by VCS caused by an incomplete reconfiguration:

```
# hagrps -freeze service_group
```

- 3 Log on as an Oracle user.

```
# su - <Owner>
```

- 4 Set the environment variables for `ORACLE_HOME` and `ORACLE_SID`.

```
# export ORACLE_HOME=<Home>
```

```
# export ORACLE_SID=<Sid>
```

- 5 Start the `sqlplus` utility to set up a database table:

```
$ORACLE_HOME/bin/sqlplus /nolog
```

- 6 As the database administrator, issue the following statements at the `sqlplus` prompt to create the test table.

For traditional or container databases:

```
connect / as sysdba

create user <User>
identified by <Pword>
default tablespace USERS
temporary tablespace TEMP
quota 100K on USERS;

grant create session to <User>;

create table <User>.<Table> ( tstamp date );
insert into <User>.<Table> (tstamp) values (SYSDATE);
```

Note: For a CDB resource, the user name should be prefixed with C##.

To enable detail monitoring for PDB, issue the following statements at the `sqlplus` prompt to create the test table:

```
$ORACLE_HOME/bin/sqlplus /nolog
connect / as sysdba
```

Create common user for CDB if not already created.

```
create user <User>
identified by <Pword>
default tablespace USERS
temporary tablespace TEMP
quota 100K on USERS;
grant connect to <User> container=all;
alter session set container=PDBName
create table <User>.<Table> ( tstamp date );
insert into <User>.<Table> (tstamp) values (SYSDATE);
```

Note: The VCS Agent for Oracle requires a common user for both CDB and PDB detail monitoring, else detail monitoring will not work for PDB.

- 7 To test the database table for use, do the following for CDB or traditional databases:

```
$ORACLE_HOME/bin/sqlplus /nolog
connect / as sysdba
connect <User>/<Pword>
update <User>.<Table> set (tstamp)=SYSDATE;
select TO_CHAR(tstamp, 'MON DD, YYYY HH:MI:SS AM')
from <User>.<Table>;
```

For PDB

```
alter session set container=PDBName;
update <User>.<Table> set ( tstamp ) = SYSDATE;
select TO_CHAR(tstamp, 'MON DD, YYYY HH:MI:SS AM')
from <User>.<Table>;
exit
```

- 8 Enable the detail monitoring for the Oracle resource using the following VCS commands:

```
# hares -modify OracleResource User <User>
# hares -modify OracleResource Pword <Pword>
# hares -modify OracleResource Table <Table>
# hares -modify OracleResource MonScript "./bin/Oracle/SqlTest.pl"

# hares -override OracleResource LevelTwoMonitorFreq
# hares -modify OracleResource LevelTwoMonitorFreq 1

# haconf -dump -makero

# hagrps -unfreeze service_group
```

You can also use Cluster Manager (Java Console) to set these attributes.

Enabling and disabling detail monitoring for Oracle resource

Review the instructions to enable or disable detail monitoring.

To enable detail monitoring

- ◆ Set the value of the LevelTwoMonitorFreq attribute.

```
# hares -override OracleResource LevelTwoMonitorFreq
# hares -modify OracleResource LevelTwoMonitorFreq 1
```

To disable detail monitoring

- ◆ Set the value of the `LevelTwoMonitorFreq` attribute to 0.

```
# hares -modify OracleResource LevelTwoMonitorFreq 0
```

Setting up detail monitoring for Netlsnr

For Netlsnr agent, the detail monitoring is not enabled by default to monitor the listener process.

You can enable detail monitoring for Netlsnr by specifying a value for the `MonScript` attribute. The example to set up detail monitoring uses the supplied monitor script for Netlsnr, `/opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl`. The detail monitoring script for the Netlsnr resource uses the Listener command `lsnrctl status $Listener` to test the Listener process.

To disable detail monitoring for Netlsnr

- ◆ Disable detail monitoring by setting the `MonScript` attribute to an empty string:

```
# haconf -makerw
# hagr -freeze service_group
# hares -modify LsnrResource MonScript ""

# hares -override LsnrResource LevelTwoMonitorFreq
# hares -modify LsnrResource LevelTwoMonitorFreq 0

# haconf -dump -makero
# hagr -unfreeze
```


To set up detail monitoring for Netlsnr

- 1 Make the VCS configuration writable:

```
# haconf -makerw
```

- 2 Freeze the service group to avoid automated actions by VCS caused by an incomplete reconfiguration:

```
# hagrps -freeze service_group
```

- 3 Enable detail monitoring by entering the following commands:

```
# hares -modify LsnrResource MonScript "./bin/Netlsnr/LsnrTest.pl"
```

```
# hares -override LsnrResource LevelTwoMonitorFreq
```

```
# hares -modify LsnrResource LevelTwoMonitorFreq 1
```

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

```
# hagrps -unfreeze service_group
```

Enabling and disabling intelligent resource monitoring for agents manually

Review the following procedures to enable or disable intelligent resource monitoring manually. The intelligent resource monitoring feature is enabled by default. The IMF resource type attribute determines whether an IMF-aware agent must perform intelligent resource monitoring.

See [“How the Oracle and Netlsnr agents support intelligent resource monitoring”](#) on page 11.

To enable intelligent resource monitoring

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```

- 2 Run the following command to enable intelligent resource monitoring.

- To enable intelligent monitoring of offline resources:

```
# hatype -modify resource_type IMF -update Mode 1
```

- To enable intelligent monitoring of online resources:

```
# hatype -modify resource_type IMF -update Mode 2
```

- To enable intelligent monitoring of both online and offline resources:

```
# hatype -modify resource_type IMF -update Mode 3
```

- 3 If required, change the values of the MonitorFreq key and the RegisterRetryLimit key of the IMF attribute.

See [“Attribute definition for the Oracle agent”](#) on page 143.

See [“Attribute definition for the Netlsnr agent”](#) on page 153.

- 4 Save the VCS configuration.

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

- 5 Make sure that the AMF kernel driver is configured on all nodes in the cluster.

```
/etc/init.d/amf status
```

If the AMF kernel driver is configured, the output resembles:

```
AMF: Module loaded and configured
```

Configure the AMF driver if the command output returns that the AMF driver is not loaded or not configured.

See [“Administering the AMF kernel driver”](#) on page 91.

- 6 Restart the agent. Run the following commands on each node.

```
# haagent -stop agent_name -force -sys sys_name
```

```
# haagent -start agent_name -sys sys_name
```

To disable intelligent resource monitoring

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```

- 2 To disable intelligent resource monitoring for all the resources of a certain type, run the following command:

```
# hatype -modify resource_type IMF -update Mode 0
```

- 3 To disable intelligent resource monitoring for a specific resource, run the following command:

```
# hares -override resource_name IMF
# hares -modify resource_name IMF -update Mode 0
```

- 4 Save the VCS configuration.

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

Note: VCS provides haimfconfig script to enable or disable the IMF functionality for agents. You can use the script with VCS in running or stopped state. Use the script to enable or disable IMF for the IMF-aware bundled agents, enterprise agents, and custom agents.

Administering the AMF kernel driver

Review the following procedures to start, stop, or unload the AMF kernel driver.

To start the AMF kernel driver

- 1 Set the value of the AMF_START variable to 1 in the following file, if the value is not already 1:

```
# /etc/sysconfig/amf
```

- 2 Start the AMF kernel driver. Run the following command:

```
# /etc/init.d/amf start
```

To stop the AMF kernel driver

- 1 Set the value of the AMF_STOP variable to 1 in the following file, if the value is not already 1:

```
# /etc/sysconfig/amf
```

- 2 Stop the AMF kernel driver. Run the following command:

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

To unload the AMF kernel driver

- 1 If agent downtime is not a concern, use the following steps to unload the AMF kernel driver:

- Stop the agents that are registered with the AMF kernel driver.
The `amfstat` command output lists the agents that are registered with AMF under the Registered Reapers section.
See the `amfstat` manual page.
 - Stop the AMF kernel driver.
See [“To stop the AMF kernel driver”](#) on page 91.
 - Start the agents.
- 2** If you want minimum downtime of the agents, use the following steps to unload the AMF kernel driver:
- Run the following command to disable the AMF driver even if agents are still registered with it.


```
# amfconfig -Uof
```
 - Stop the AMF kernel driver.
See [“To stop the AMF kernel driver”](#) on page 91.

Configuring VCS service groups for Oracle using the Veritas High Availability Configuration wizard

This chapter includes the following topics:

- [Typical VCS cluster configuration in a virtual environment](#)
- [About configuring application monitoring using the High Availability solution for VMware](#)
- [Getting ready to configure VCS service groups using the wizard](#)
- [Before configuring application monitoring](#)
- [Launching the High Availability Configuration wizard](#)
- [Configuring the agent to monitor Oracle](#)
- [Understanding service group configurations](#)
- [Understanding configuration scenarios](#)
- [High Availability Configuration wizard limitations](#)
- [Troubleshooting](#)
- [Sample configurations](#)

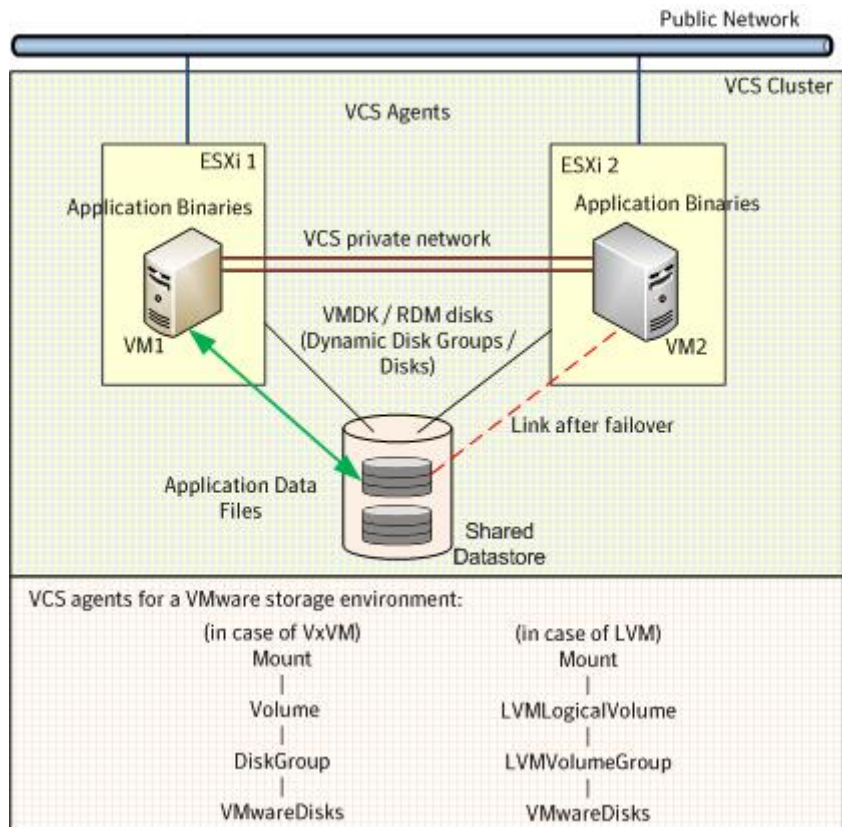
Typical VCS cluster configuration in a virtual environment

A typical VCS cluster configuration for Oracle in a VMware virtual environment involves two or more virtual machines. The virtual machine on which the application is active, accesses a non-shared VMware VMDK or RDM disk that resides on a VMware datastore.

The virtual machines involved in the VCS cluster configuration may belong to a single ESX/ESXi host or could reside on separate ESX/ESXi hosts. If the virtual machines reside on separate ESX/ESXi hosts, the datastore on which the VMware VMDK or RDM disks (on which the application data is stored) reside must be accessible to each of these ESX/ESXi hosts.

The application binaries are installed on the virtual machines and the data files are installed on the VMware disk drive. The VCS agents monitor the application components and services, and the storage and network components that the application uses.

Figure 5-1 Typical VCS cluster configuration in a VMware virtual environment



During a failover, the VCS storage agents move the VMware disks to the new system. The VCS network agents bring the network components online, and the application specific agents then start application services on the new system.

About configuring application monitoring using the High Availability solution for VMware

Consider the following before you proceed:

- You can configure application monitoring on a virtual machine using the High Availability Configuration wizard for VMware. The wizard is launched when you click **Configure application for high availability** on the High Availability tab in VMware vSphere Client.

- Apart from the High Availability Configuration wizard, you can also configure application monitoring using the Cluster Server (VCS) commands. For more information, refer to the *Cluster Server Administrator's Guide*.
- recommends that you first configure application monitoring using the wizard before using VCS commands to add additional components or modify the existing configuration.
Apart from configuring application availability, the wizard also sets up the other components required for successful application monitoring.
- For Oracle, the High Availability Configuration wizard:
 - Discovers the Oracle database instance and Oracle listeners running on the system.
 - Discovers the storage and network dependencies.
 - Configures the appropriate VCS resources and service groups from the discovered information.
- You must not suspend a system if an application is currently online on that machine. If you suspend a system, VCS moves the disks along with the application to another system. Later, when you try to restore the suspended system, VMware does not allow the operation because the disks that were attached before the system was suspended are no longer with the system. To suspend a virtual machine, ensure that the application being monitored is not online on that system.

Note: For details about deploying, configuring, and administering the High Availability solution, refer to the *High Availability Solutions Guide for VMware*.

Getting ready to configure VCS service groups using the wizard

Ensure that you complete the following tasks before configuring application monitoring on a virtual machine:

- Install the VMware vSphere Client.
- Install and enable VMware Tools on the virtual machine, where you want to monitor applications with VCS. Install a version that is compatible with the VMware ESX/ESXi server.
- Install High Availability console on a Windows system in your data center and register the High Availability plug-in with the vCenter server.

- Assign Configure Application Monitoring (Admin) privileges to the logged-on user on the virtual machine where you want to configure application monitoring.
- Install .
- Install the application and the associated components that you want to monitor on the virtual machine.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by the High Availability installer, wizards, and services. Refer to the *High Availability Solutions Guide for VMware* for a list of ports and services used.

Before configuring application monitoring

Note the following points before configuring application monitoring on a virtual machine:

- All the Oracle and Net Listener instances that you want to configure must be running on the system from where the High Availability Configuration wizard is invoked.
- The wizard discovers the disks which are attached and the storage which is currently available. Ensure that the shared storage used by the application is available before you invoke the wizard.
- All the required disks must be attached and all the storage components must be available.
- The Oracle Home directory owner must exist on all the failover nodes.
- The Oracle UID must be the same across all the nodes in the cluster.
- If the Oracle Database is installed on local disks, the Oracle Home directory must exist on all the failover targets.
- If the Oracle Database is installed on shared disks, then the corresponding mount point must be selected when you configure the Oracle instance using the High Availability Configuration wizard.
- You must not restore a snapshot on a virtual machine where an application is currently online, if the snapshot was taken when the application was offline on that virtual machine. Doing this may cause an unwanted failover. This also applies in the reverse scenario; you should not restore a snapshot where the application was online on a virtual machine, where the application is currently offline. This may lead to a misconfiguration where the application is online on multiple systems simultaneously.

- While creating a VCS cluster in a virtual environment, you must configure the cluster communication link over a public network in addition to private adapters. The link using the public adapter should be assigned as a low-priority link. This helps in case the private network adapters fail, leading to a condition where the systems are unable to connect to each other, consider that the other system has faulted, and then try to gain access to the disks, thereby leading to an application fault.
- You must not select teamed network adapters for cluster communication. If your configuration contains teamed network adapters, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed network adapters. A teamed network adapter is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address, due to which you may experience the following issues:
 - SSO configuration failure.
 - The wizard may fail to discover the specified network adapters.
 - The wizard may fail to discover/validate the specified system name.
- Verify that the boot sequence of the virtual machine is such that the boot disk (OS hard disk) is placed before the removable disks. If the sequence places the removable disks before the boot disk, the virtual machine may not reboot after an application failover. The reboot may halt with an "OS not found" error. This issue occurs because during the application failover the removable disks are detached from the current virtual machine and are attached on the failover target system.
- Verify that the disks used by the application that you want to monitor are attached to non-shared controllers so that they can be deported from the system and imported to another system.
- If multiple types of SCSI controllers are attached to the virtual machines, then storage dependencies of the application cannot be determined and configured.
- The term 'shared storage' refers to the removable disks attached to the virtual machine. It does not refer to disks attached to the shared controllers of the virtual machine.
- If you want to configure the storage dependencies of the application through the wizard, the LVM volumes or VxVM volumes used by the application should not be mounted on more than one mount point path.
- The host name of the system must be resolvable through the DNS server or, locally, using `/etc/hosts` file entries.
- By default, the controller ID and port must remain the same on all cluster nodes. If you do not want the resource to have the same controller ID and port, you

should localize the attribute for all cluster nodes. Localization allows all cluster nodes to have different controller IDs and port numbers. For more information about localizing an attribute, refer to the *Cluster Server Administrator's Guide*.

Launching the High Availability Configuration wizard

You can launch the High Availability Configuration wizard from:

- VMware vSphere Client: See [To launch the wizard from the VMware vSphere Client](#).
- A browser window: See [To launch the wizard from a browser window](#).

You must launch the High Availability Configuration wizard from the system where the disk residing on the shared datastore is attached.

To launch the wizard from the VMware vSphere Client

- 1 Launch the VMware vSphere Client and connect to the VMware vCenter Server that hosts the virtual machine.
- 2 From the vSphere Client's Inventory view in the left pane, select the virtual machine where you want to configure application monitoring.
- 3 Skip this step if you have already configured single sign-on during guest installation.

Select the High Availability tab and in the High Availability View page, specify the credentials of a user account that has administrative privileges on the virtual machine and click **Configure**.

The High Availability console sets up a permanent authentication for the user account on that virtual machine.

- 4 Depending on your setup, use one of the following options to launch the wizard:
 - If you have not configured a cluster, click the **Configure application for high availability** link.
 - If you have already configured a cluster, click **Actions > Configure application for high availability** or the **Configure application for high availability** link.
 - If you have already configured a cluster and configured an application for monitoring, click **Actions > Configure application for high availability**.

To launch the wizard from a browser window

- 1 Open a browser window and enter the following URL:

`https://<VMNameorIP>:5634/vcs/admin/application_health.html`

<VMNameorIP> is the virtual machine name or IP address of the system on which you want to configure application monitoring.

- 2 In the Authentication dialog box, enter the username and password of the user who has administrative privileges.
- 3 Depending on your setup, use one of the following options to launch the wizard:
 - If you have not configured a cluster, click the **Configure application for high availability** link.
 - If you have already configured a cluster, click **Actions > Configure application for high availability** or the **Configure application for high availability** link.
 - If you have already configured a cluster and configured an application for monitoring, click **Actions > Configure application for high availability**.

Configuring the agent to monitor Oracle

Perform the following steps to configure monitoring for Oracle on a virtual machine using the High Availability Configuration wizard.

To configure monitoring for Oracle

- 1 Launch the High Availability Configuration wizard. See [“Launching the High Availability Configuration wizard”](#) on page 99.
- 2 Review the information on the Welcome panel and click **Next**.
- 3 On the Application Inputs panel, select **Oracle** from the Supported Applications list and then click **Next**.

You can use the Search box to search for the Oracle application.

- 4 To review and modify listener configurations, in the Oracle Listener Details panel, select a listener from the list that the wizard automatically discovers and displays.

For each listener, you can review the following details:

- Listener
- Home
- Owner

- Env File
 - TnsAdmin
 - Password
- The wizard automatically populates some of the above fields. Among the displayed fields, only the Home and Owner fields represent mandatory attributes. Apart from Listener, you can edit the values in the above fields.
- 5 If you want to add a listener for monitoring, enter the name of the listener in the Listener text box. Click **Add** and then provide the required details.

If you want to remove a listener from the displayed list, select the listener and click **Remove**.
 - 6 Repeat steps 4 and 5 for other listener instances that you want to monitor, and then click **Next**.
 - 7 To review and modify database configurations (including its associated listeners and mount points), on the Database Selection panel, click a database instance from the list of Database SIDs, that the wizard automatically discovers and displays.
 - 8 For each database instance, you can review the following details:
 - Database SID
 - Oracle Home
 - Database Owner
 - PFILE Path
 - Env File

Among the displayed fields, only Database SID, Oracle Home, and Database Owner represent mandatory attributes. For each database instance that you select, the wizard also automatically displays the associated listeners and mount points.
 - 9 For each database, from the Selected Listeners list and Mount Points list, you can select the listeners and mount points that you want to associate with the database instance, and set up dependencies.
 - 10 Repeat steps 7 to 9 for each database instance that you want to monitor, and then click **Next**.

- 11 On the Configuration Inputs panel, use the Edit icon to specify the user name and password of the systems for VCS cluster operations.

Cluster systems lists the systems included in the cluster configuration.

Application failover targets lists the systems to which the application can fail over. Move the required systems to the Application failover targets list. Use the up and down arrow keys to define the priority order of the failover systems. The local system is selected by default for both, the cluster operations and as a failover target.

- 12 Click **Next**.

- 13 Skip this step if you do not want to add more systems to your cluster.

To add a system to the cluster, click **Add System**. In the Add System dialog box, specify the following details of the system that you want to add to the VCS cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account mentioned.
Use the specified user account on all systems	Select to use the specified user account on all the cluster systems that have the same user name and password.

- 14 If you are configuring a cluster and if you want to modify the security settings for the cluster, click **Advanced Settings**. In the Advanced settings dialog box, specify the following details and click **OK**.

Use Single Sign-on	Select to configure single sign-on using VCS Authentication Service for cluster communication. This option is enabled by default.
Use VCS user privileges	Select to configure a user with administrative privileges to the cluster. Specify the username and password and click OK .

Note: The **Advanced Settings** link is not visible if the cluster is already created.

- 15** Skip this step if the cluster is already configured. By default, the links are configured over Ethernet.

On the Network Details panel, select the type of network protocol to configure the VCS cluster network links and then specify the adapters for network communication.

The wizard configures the VCS cluster communication links using these adapters. You must select a minimum of two adapters per system.

Select **Use MAC address for cluster communication (LLT over Ethernet)** or **Use IP address for cluster communication (LLT over UDP)**, depending on the IP protocol that you want to use and then specify the required details to configure the VCS cluster communication network links. You must specify these details for each cluster system.

- To configure LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure LLT over UDP, select the type of IP protocol and then specify the required details for each communication link.

Depending on the IP protocol, specify the following:

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The second link is configured as a high-priority link. To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal VCS cluster communication over the link.

16 Skip this step if the application does not use virtual IP address.

Note: If the application does not use a virtual IP address, you must remove the auto-populated subnet mask entry before proceeding to the next step.

In the Virtual Network Details panel, specify the IP protocol and virtual IP address for the application.

Depending on the IP protocol, specify the following:

Virtual IP address	Specify a unique virtual IP address.
Subnet Mask (IPv4)	Specify the subnet mask details.
Prefix (IPv6)	Select the prefix from the drop-down list.
Network Adapter	Select the network adapter that will host the virtual IP.

If you want to add another virtual IP address for your application, click **Add virtual IP address**.

If you want to remove a virtual IP address, click the Remove icon.

If you select multiple instances for the same virtual IP address, those instances are configured in the same service group.

17 Click **Next**.

18 Skip this step if you did not select mount points.

On the Storage HA Inputs panel, specify all the ESX/ESXi hosts to which the virtual machines can fail over. Each ESX/ESXi host must be able to access the required shared datastores that contain visible disks. Enter the administrative user account details for each ESX/ESXi host and click **Next**.

To specify the ESX/ESXi hosts, click **Add ESX/ESXi Host** and in the Add ESX/ESXi Host dialog box, specify the following details:

ESX/ESXi hostname or IP address	Specify the target ESX/ESXi hostname or IP address. The virtual machines can fail over on this ESX/ESXi host during vMotion. All the additional ESX/ESXi hosts should have access to the datastore on which the disks used by the application reside.
User name	Specify a user account for the ESX/ESXi host. The user account must have administrator privileges on the specified ESX/ESXi host.
Password	Specify the password for the user account provided in the User name text box.

The wizard validates the user account and the storage details on the specified ESX/ESXi hosts.

19 On the Summary panel, review the VCS cluster configuration summary and then click **Next** to proceed with the configuration.

If the network contains multiple clusters, the wizard verifies the cluster ID with the IDs assigned to all the accessible clusters in the network. The wizard does not validate the assigned ID with the clusters that are not accessible during the validation. recommends you to validate the uniqueness of the assigned ID in the existing network. If the assigned ID is not unique or if you want to modify the cluster name or cluster ID, click **Edit**. In the Edit Cluster Details dialog box, modify the details as necessary and click **OK**.

- 20** On the Implementation panel, the wizard creates the VCS cluster, configures the application for monitoring, and creates cluster communication links.

The wizard displays the status of each task. After all the tasks are complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to configure application monitoring.

- 21** On the Finish panel, click **Finish** to complete the wizard workflow.

This completes the application monitoring configuration.

If the application status shows as not running, click **Start** to start the configured components on the system.

Understanding service group configurations

One or more Oracle instances can be discovered on a virtual machine. These Oracle instances may or may not share the same listeners, mount points, disks, or disk groups. The Oracle listeners that do not share any of these forms a separate service group.

Resource dependency

The following are the resource dependencies:

- Net Listener resources associated with an Oracle instance depend on Oracle resources.
- Net Listener resources also depend on the configured IP resources.
- Oracle resource depends on mount point resources which were selected for that particular Oracle instance.
- Mount point resources depend on either LVM (logical volume) or VxVM volume.
 - VxVM volume depends on DiskGroup resources.
 - LVM depends on LVM volume group.
- DiskGroup and LVM volume group resources depend on the shared disks which are configured as VMwareDisks resources.

Service group dependency

The High Availability Configuration wizard does not create service group dependency for Oracle.

Infrastructure service groups

As part of configuring the application, the High Availability Configuration wizard:

- Configures application specific service groups and resources.
- Configures the VCS infrastructure service group (VCSInfraSG).

VCSInfraSG includes a resource called VCSNotifySinkRes. The type of this resource is Process. VCSNotifySinkRes configures and administers the notify_sink process on the guest. The notify_sink process sends the details about service groups and its attributes to the High Availability Console. This information is used for reporting purpose and is displayed on the Dashboard.

Note: VCSInfraSG is an internal service group. You must not add or delete resources from this service group.

The following are the VCSInfraSG notes:

- Before you configure the application for monitoring, ensure that SSO is configured between the High Availability Console and the guest. If SSO is not configured, VCSInfraSG fails to come online.
- If VCSInfraSG or VCSNotifySinkRes faults, ensure that SSO is configured between the High Availability Console and the guest. Clear the faults and bring the resource online again.
- VCSInfraSG or VCSNotifySinkRes must not be taken offline because it affects the information displayed on the Dashboard.

Understanding configuration scenarios

You can configure Oracle instances in different ways using the High Availability Configuration wizard.

Table 5-1 Oracle configurations

Configuration Type	Reference
Configuring a single instance/multiple instances in VCS	See "Configuring a single instance/multiple instances in VCS" on page 108.
Configuring multiple Oracle instances in VCS using multiple runs of the wizard	See "Configuring multiple Oracle instances in VCS using multiple runs of the wizard" on page 108.

Table 5-1 Oracle configurations (*continued*)

Configuration Type	Reference
Configuring multiple applications	See “Configuring multiple applications” on page 108.

Configuring a single instance/multiple instances in VCS

The High Availability Configuration wizard can be used to configure one or more Oracle instances in a single run.

Select the instances from the Oracle instance selection panel of the wizard. The related listeners and mount points are displayed and can be selected for each instance.

Configuring multiple Oracle instances in VCS using multiple runs of the wizard

If you are configuring the first Oracle instance on a machine where more than one instance is running, you can configure it by following the steps in the [Configuring a single instance/multiple instances in VCS](#) section.

The High Availability Configuration wizard will not allow you to configure the next Oracle instance if any of the mandatory dependent resources such as mount point, disk group, disk, and listener are already configured in VCS.

- If existing resources are part of the Oracle service group, unconfigure the existing service group and then reconfigure the new instance along with the old instances/listeners which were part of the pre-existing service group.
- If existing resources are part of an application service group other than Oracle, the wizard does not support configuring multiple applications. You can configure these applications through CLI or InfoScale Operations Manager.

Configuring multiple applications

If you run the High Availability Configuration wizard multiple times, you can configure multiple applications of different types.

If you are configuring the first application on a machine where more than one application is running, you can configure it by following the steps in the [Configuring a single instance/multiple instances in VCS](#) section.

The High Availability Configuration wizard will not allow you to configure the next application if any of the mandatory dependent resources such as mount point, disk group, disk, and listener are already configured in VCS.

High Availability Configuration wizard limitations

The following are the High Availability Configuration wizard limitations:

- If the Oracle database instance or listener instance is already configured, the Oracle instance is not shown and it will not be available for configuration. In such a scenario, no error message will be displayed.
For example, if there are two Oracle instances running on the system with the same Oracle home or listener and assuming that the first instance is configured, if you run the wizard to configure the next instance, the following error message is displayed:

```
The wizard has failed to discover Oracle on the system.
```

However, if the Oracle home or listener do not share any resources (storage and network), you can configure the Oracle instance.

- The wizard supports discovery of only LVM or VxVM type of storage.
- You cannot use the wizard to discover the instances that use ASM or raw disks.
- The wizard will not discover the disks used by the application if the controllers attached to the virtual machine are of different type.
To correctly discover and identify the association of mount points to the virtual disks, all the controllers attached to the virtual machine must be of same type.
- The wizard will not discover disks which are attached to the virtual machine in shared mode.

Troubleshooting

This section lists common troubleshooting scenarios that you may encounter while or after configuring application monitoring.

High Availability Configuration wizard displays blank panels

The High Availability Configuration wizard may fail to display the wizard panels. The window may appear blank.

Workaround:

Verify that the ApplicationHA Service is running on the High Availability Console host and then launch the wizard again.

The High Availability Configuration wizard displays the 'hadiscover is not recognized as an internal or external command' error

While configuring application monitoring the High Availability Configuration wizard may display the "hadiscover is not recognized as an internal or external command" error, after you click Next on the Application Selection panel.

This issue occurs if you launch the wizard from a system where you have reinstalled the High Availability guest components.

Workaround:

Close the wizard, restart the Veritas Storage Foundation Messaging Service, and then re-run the wizard.

Running the 'hastop -all' command detaches virtual disks

The `hastop -all` command takes offline all the components and component groups of a configured application, and then stops the VCS cluster. In the process, the command detaches the virtual disks from the VCS cluster nodes.

Workaround:

If you want to stop the VCS cluster (and not the applications running on cluster nodes), instead of the `hastop -all` command, use the following command:

```
hastop -all -force
```

This command stops the cluster without affecting the virtual disks attached to the VCS cluster nodes.

Log files

The log files are stored in the virtual machine on which you configured application monitoring.

The `healthview_A.log` file contains the steps performed by the back-end to configure the application. To check the file, you must access:

```
/var/VRTSvcS/log/healthview_A.log
```

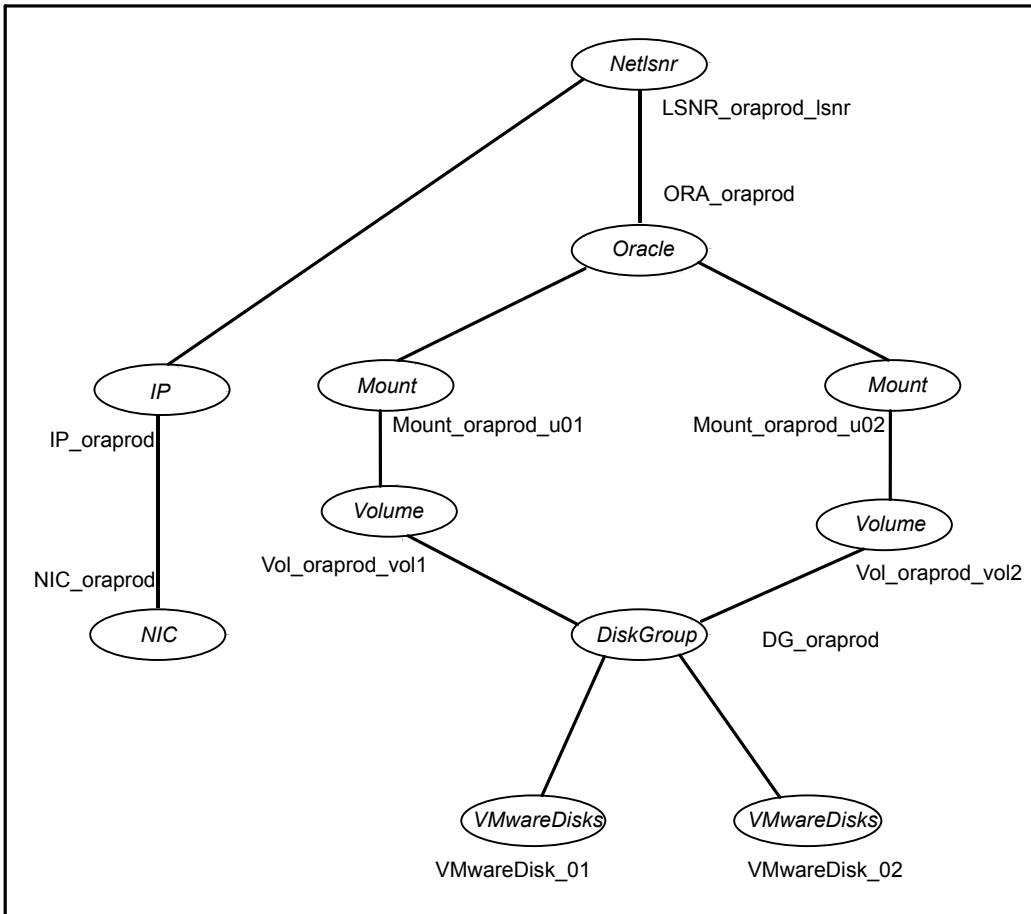
Sample configurations

The sample configurations includes description for typical service groups that are configured using the High Availability Configuration wizard.

Sample VCS configuration file for single Oracle instance (VxVM)

Figure 5-2 shows a typical service group configured to monitor the state of an Oracle instance. In this example, the Oracle instance uses VxVM volume as storage.

Figure 5-2 Dependency graph for single Oracle instance (VxVM)



Review the sample configuration with a resource of type Oracle that is configured as follows in main.cf file.

```

include "types.cf"
include "OracleTypes.cf"

cluster vcs (
)

```

```

system sys1 (
)

system sys2 (
)

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_oraprod (
    DiskGroup = ora_prod_dg
    StartVolumes = 0
    StopVolumes = 0
)

IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.22"
    Netmask = "255.255.255.0"
)

Mount Mount_oraprod_u01 (
    MountPoint = "/prod/u01"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol1"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oraprod_u02 (
    MountPoint = "/prod/u02"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol1"
    FSType = vxfs
    FsckOpt = "-n"
)

NIC NIC_oraprod (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)

```



```

Netlsnr LSNR_oraprod_lsnr (
    Owner = oraprod
    Home = "/orahome/Oracle"
    TnsAdmin = "/orahome/Oracle/network/admin"
    Listener = LISTENER_PROD
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
    LsnrPwd = cqfOdoOolOo
)

Oracle ORA_oraprod (
    Sid = PROD
    Owner = oraprod
    Home = "/orahome/Oracle"
    EnvFile = "/tmp/env.sh"
    MonScript = "./bin/Oracle/SqlTest.pl"
    User = thor
    Pword = hvlTptWvj
    DBAUser = "orauser"
    DBAPword = "DPSrFFeRiRGPiRF"
    Table = thor
    MonitorOption = 0
)

Volume Vol_oraprod_vol1 (
    Volume = u01-vol
    DiskGroup = ora_prod_dg
)

Volume Vol_oraprod_vol2 (
    Volume = u02-vol
    DiskGroup = ora_prod_dg
)

VMwareDisks VMwareDisk_01 (
    ESXDetails = {
        "ESX1.domain.com" = "root=Encrypted Password",
        "ESX2.domain.com" = "root=Encrypted Password" }
    DiskPaths = {
        "6000C29d-5235-4480-823e-1234abcd1234:
        [datastore1] disk1.vmdk" = "0:1" }
)

VMwareDisks VMwareDisk_02 (
    ESXDetails = {

```

```

"ESX1.domain.com" = "root=Encrypted Password",
"ESX2.domain.com" = "root=Encrypted Password" }
DiskPaths = {
    "RDM:6000C290-2462-33ca-7019-1234abcd1234:
    [datastore1] disk2.vmdk" = "0:2" }
)

```

```

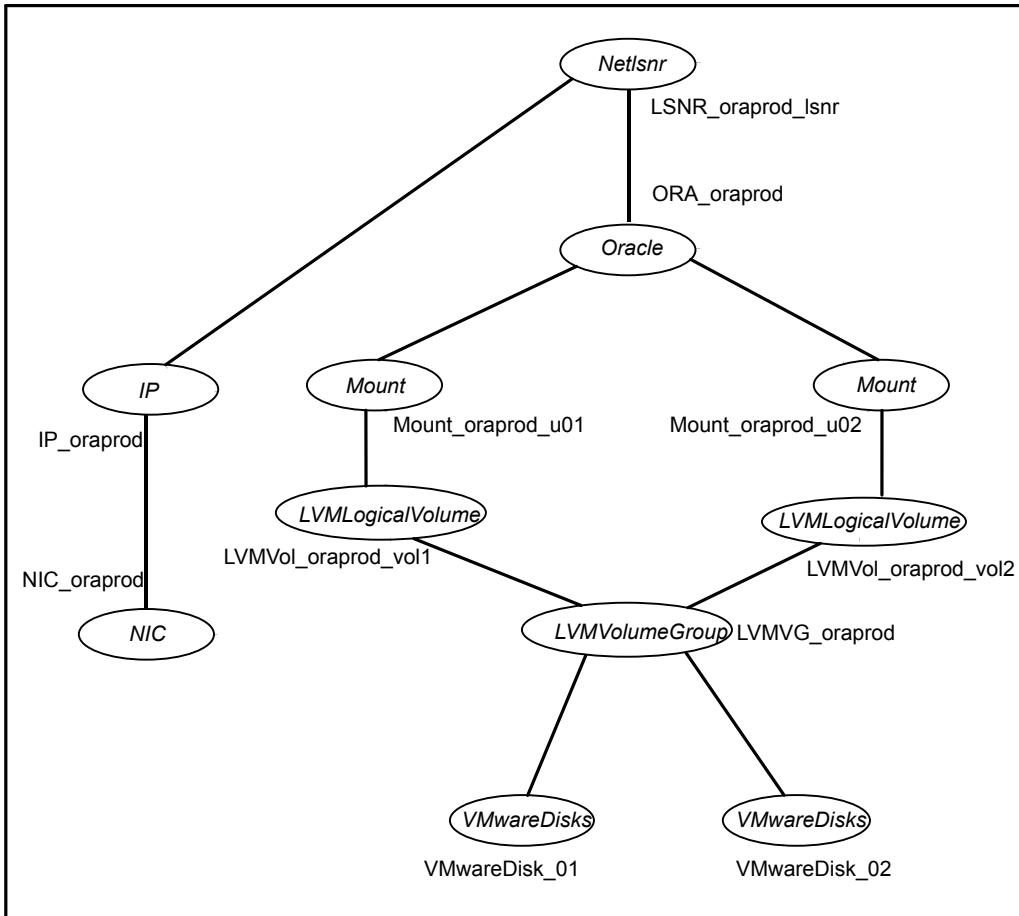
IP_oraprod requires NIC_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
ORA_oraprod requires Mount_oraprod_u01
ORA_oraprod requires Mount_oraprod_u02
Mount_oraprod_u01 requires Vol_oraprod_vol1
Mount_oraprod_u02 requires Vol_oraprod_vol2
Vol_oraprod_vol1 requires DG_oraprod
Vol_oraprod_vol2 requires DG_oraprod
DG_oraprod requires VMwareDisk_01
DG_oraprod requires VMwareDisk_02

```

Sample VCS configuration file for single Oracle instance (LVM)

[Figure 5-3](#) shows a typical service group configured to monitor the state of an Oracle instance. In this example, the Oracle instance uses LVM volume as storage.

Figure 5-3 Dependency graph for single Oracle instance (LVM)



Review the sample configuration with a resource of type Oracle that is configured as follows in main.cf file.

```

include "types.cf"
include "OracleTypes.cf"

cluster vcs (
)

system sys1 (
)
  
```

```

system sys2 (
)

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

LVMVolumeGroup LVMVG_oraprod (
    LVMVolumeGroup = ora_prod_vg
)

IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.22"
    Netmask = "255.255.255.0"
)

Mount Mount_oraprod_u01 (
    MountPoint = "/prod/u01"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oraprod_u02 (
    MountPoint = "/prod/u02"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

NIC NIC_oraprod (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)

Netlsnr LSNR_oraprod_lsnr (
    Owner = oraprod
    Home = "/orahome/Oracle"
    TnsAdmin = "/orahome/Oracle/network/admin"
    Listener = LISTENER_PROD
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
)

```

```

        LsnrPwd = cqfOdoOolOo
    )

Oracle ORA_oraprod (
    Sid = PROD
    Owner = oraprod
    Home = "/orahome/Oracle"
    EnvFile = "/tmp/env.sh"
    MonScript = "./bin/Oracle/SqlTest.pl"
    User = thor
    Pword = hvlTptWvj
    DBAUser = "orauser"
    DBAPword = "DPSrFFeRiRGPiRF"
    Table = thor
    MonitorOption = 0
)

LVMLogicalVolume LVMVol_oraprod_vol1 (
    LVMLogicalVolume = ora_oraprod_vol1
    LVMVolumeGroup = ora_prod_vg
)

LVMLogicalVolume LVMVol_oraprod_vol2 (
    LVMLogicalVolume = ora_oraprod_vol2
    LVMVolumeGroup = ora_prod_vg
)

VMwareDisks VMwareDisk_01 (
    ESXDetails = {
        "ESX1.domain.com" = "root=Encrypted Password",
        "ESX2.domain.com" = "root=Encrypted Password" }
    DiskPaths = {
        "6000C29d-5235-4480-823e-1234abcd1234:
        [datastore1] disk1.vmdk" = "0:1" }
)

VMwareDisks VMwareDisk_02 (
    ESXDetails = {
        "ESX1.domain.com" = "root=Encrypted Password",
        "ESX2.domain.com" = "root=Encrypted Password" }
    DiskPaths = {
        "RDM:6000C290-2462-33ca-7019-1234abcd1234:
        [datastore1] disk2.vmdk" = "0:2" }
)

```

```
IP_oraprod requires NIC_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
ORA_oraprod requires Mount_oraprod_u01
ORA_oraprod requires Mount_oraprod_u02
Mount_oraprod_u01 requires LVMVol_oraprod_vol1
Mount_oraprod_u02 requires LVMVol_oraprod_vol2
LVMVol_oraprod_vol1 requires LVMVG_oraprod
LVMVol_oraprod_vol2 requires LVMVG_oraprod
LVMVG_oraprod requires VMwareDisk_01
LVMVG_oraprod requires VMwareDisk_02
```

Administering VCS service groups for Oracle

This chapter includes the following topics:

- [About administering VCS service groups](#)
- [Bringing the service group online](#)
- [Taking the service group offline](#)
- [Switching the service group](#)
- [Modifying the service group configuration](#)

About administering VCS service groups

You can administer service groups in Cluster Server using one of the following interfaces:

- Veritas Operations Manager
See the Veritas InfoScale Operations Manager documentation for more information.
- The Cluster Manager (Java Console)
See [“Bringing the service group online”](#) on page 120.
See [“Taking the service group offline”](#) on page 120.
See [“Switching the service group”](#) on page 120.
See [“Modifying the service group configuration”](#) on page 121.
See the *Cluster Server Administrator’s Guide* for more information.
- The command-line
See the *Cluster Server Administrator’s Guide* for more information.

Bringing the service group online

Perform the following steps to bring the service group online.

To bring a service group online

- 1 In the Cluster Explorer configuration tree, select the newly created service group.
- 2 Right-click the service group name, and select **Enable Resources** to enable all resources in the service group.
- 3 Right-click the service group name, and select the systems on which to enable the service group (Right-click>Enable>*system_name* or Right-click>Enable>All).
- 4 Save your configuration (File>Close Configuration).
- 5 Right-click the service group and select to online the service group on the system (Right-click>Online>*system_name*).

Taking the service group offline

Perform the following steps to take the service group offline.

To take a service group offline

- 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 **Offline**, and choose the appropriate system from the pop-up menu (Right-click>Offline>*system_name*).

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

- 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*).

Modifying the service group configuration

You can dynamically modify the Cluster Server agent for Oracle agent using one of the supported interfaces to administer VCS.

See [“About administering VCS service groups”](#) on page 119.

Pluggable database (PDB) migration

This chapter includes the following topics:

- [Migrating Pluggable Databases \(PDB\) between Container Databases \(CDB\)](#)

Migrating Pluggable Databases (PDB) between Container Databases (CDB)

You can use the `hapdbmigrate` utility to perform a planned end-to-end migration of pluggable databases across containers of the same database version. The destination container database may be on the same node or on another node. The PDB must have its own independent storage.

If the PDB has parent groups, they will be unlinked and frozen by the utility during the migration. After successful migration, they will be relinked and unfrozen by the utility.

The migration is supported on database version 12.1.0.2.

The configuration must meet the following requirements:

- The PDB and CDB databases must be on Veritas Cluster File System (CFS) or Veritas File system (VxFS).
- There must be no parent resources for the PDB resource.
- The PDB must be plugged in to the source CDB. The PDB resources may or may not be offline.
- The source and destination CDB resources for the migration must be different.
- The version of the destination CDB must be the same as the source CDB.

- The PDB to be migrated is mounted on an independent file system (separate mount points for CDB and PDB datafiles).
- The CDB and PDB resources must be configured in the same service group.

The utility performs the following actions during the migration:

- Unlinks and freezes the parent groups, if any, depending on the source CDB group where the PDB resource to be migrated is configured.
- Takes PDB resources offline.
- Unplugs the PDB from the CDB and creates an XML file, `<pdb_res_name>_<dest_cdb_res_name>.xml` in the XML directory provided by the user.
- Drops the PDB from the CDB keeping the datafiles.
- Takes offline all the PDB child resources. Unlinks the PDB resource from the source CDB resource and deletes the PDB resource. Unlinks and deletes all PDB child resources.
- Recreates the PDB resource and all its children with original dependencies in the destination CDB service group.
- Brings online the PDB child resources on all the nodes where the destination CDB service group is online
- Plugs the PDB in the destination CDB.
- Brings the PDB resource online after successful plugging. Unfreezes and links the parent group to the source CDB group.

The `hapdbmigrate` utility performs certain pre-requisite checks before the migration to verify that the cluster is ready for PDB migration. If the utility encounters any issues, you will need to manually fix the issues.

The utility is present in the `$VCSHOME/bin` directory.

The log files of the migration are located at `$VCSLOG/log/hapdbmigrate.log`. The logs are rotated after the file exceeds 5 MB and is saved in `.gz` format `hapdbmigrate.log[1..7].gz`.

Note: Ensure that only one instance of the `hapdbmigrate` utility is running at a time.

To migrate Pluggable Databases (PDB) between Container Databases (CDB)

- 1 Back up the VCS configuration file `/etc/VRTSvcs/conf/config/main.cf`:

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

- 2 Verify that the high availability daemon (had) is running on all the nodes in the cluster.

```
# hastatus -summary
```

- 3 Verify that there are no resources in faulted or unknown state.

```
# hares -state | grep FAULTED
# hares -state | grep UNKNOWN
```

- 4 Verify that the `PDBName` attribute is present for the PDB resource with the correct value in the `main.cf` configuration file.

- 5 Verify that the source and destination CDB resources are online.

```
# hares -state resname
```

- 6 On the destination CDB, verify the following:

- The destination CDB is not in `suspended` mode.
Any instance of the destination CDB is not in `mounted` state.
See the Oracle documentation for more information.
- If any instance of the destination CDB is in `restricted` state, ensure that the PDB resource you want to migrate has the `StartUpOpt` attribute set to `restricted`.

```
# haconf -makerw
# hares -modify pdb1 StartUpOpt \
RESTRICTED
# haconf -dump -makero
```

- 7 Verify that existing dependencies do not conflict with the migration process.

The PDB child resources must not be dependent on the CDB resource or any of its child resources.

The PDB child resources must not have any parent, which is not a part of the PDB child hierarchy.

- 8 Verify that the XML data directory has read and write permissions for the "oracle" user. The XML data directory must be located either on PDB mounts or at a location accessible to both source and destination CDBs.
- 9 Run the `hapdbmigrate` utility as the root user:

Note: If there are parent groups dependent on the source CDB group, specify the `-ignoreparentgrp` option.

```
# $VCS_HOME/bin/hapdbmigrate -pdbres pdb_resname -cdbres cdb_resname \
-XMLdirectory xml_directory [-ignoreparentgrp] [-prechecks]
  -pdbres: Name of the PDB resource, which needs to be migrated
  -cdbres: Name of the CDB resource, where the PDB needs to migrate
  -XMLdirectory: XML directory location for the unplugged PDB
  -ignoreparentgrp: Utility proceeds even
                    if the PDB group has parent groups
  -prechecks: Performs prechecks and validation
  -help|h: Prints usage
```

The migration log file is located at `$VCSLOG/log/hapdbmigrate.log`.

- 10 Verify that the PDB resource is online on the destination CDB.

```
# hares -state pdb_resname
```

- 11 Relink the parent service group of the source CDB group manually to the destination CDB group, if it depends on the migrated PDB.

```
# haconf -makerw
# hagr -dep parent_sg
#Parent          Child          Relationship
parent_sg        source_CDB        online local firm
# hagr -offline parent_sg -any
# hagr -unlink parent_sg source_CDB
# hagr -link parent_sg dest_CDB online local firm
# hagr -dep parent_sg
#Parent          Child          Relationship
parent_sg        dest_CDB         online local firm
# haconf -dump -makero
```

See [“Sample configuration of migratable Oracle pluggable database \(PDB\) resource in main.cf”](#) on page 199.

See [“Resource type definition for the Oracle agent”](#) on page 142.

See [“Resource type definition for the Netsnr agent”](#) on page 153.

Troubleshooting Cluster Server agent for Oracle

This chapter includes the following topics:

- [About troubleshooting Cluster Server agent for Oracle](#)
- [Error messages common to the Oracle and Netlsnr agents](#)
- [Error messages specific to the Oracle agent](#)
- [Error messages specific to the Netlsnr agent](#)
- [Error messages specific to the ASMinst agent](#)
- [Error messages specific to the ASMDG agent](#)
- [Troubleshooting issues specific to Oracle in a VCS environment](#)
- [Verifying the Oracle health check binaries and intentional offline for an instance of Oracle](#)
- [Disabling IMF for a PDB resource](#)

About troubleshooting Cluster Server agent for Oracle

Review the information on the error logs that you must access:

- To check the Oracle installation error log, you must access:

```
$ORACLE_BASE/oraInventory/logs/installActionsdate_time.log
```

This file contains the errors that occurred during installation. It clarifies the nature of the error and at exactly which point it occurred during the installation. If there are any installation problems, you must send this file to Tech Support for debugging the issue.

- To check the log file, you must access:

```
/var/VRTSvcs/log/engine_A.log
/var/VRTSvcs/log/Oracle_A.log
/var/VRTSvcs/log/Netlsnr_A.log

/var/VRTSvcs/log/ASMIInst_A.log
/var/VRTSvcs/log/ASMDG_A.log
```

These files contain all the actions that the VCS engine and other agents for Oracle perform.

Review the description of the error messages for the following agents and the possible solutions:

- Oracle agent
- Netlsnr agent
- ASMIInst agent
- ASMDG agent

Error messages common to the Oracle and Netlsnr agents

[Table 8-1](#) lists the Cluster Server agent for Oracle error messages with the description and a recommended solution, if available.

Table 8-1 Cluster Server agent for Oracle error messages

Message	Description and solution
No ORACLE_HOME specified	The Home attribute in the Oracle or Netlsnr type has not been set. Solution: Set the Home attribute value to the correct full path name of the Oracle home.
Oracle home directory <Oracle_Home> does not exist	The string that is specified for the Home attribute in the Oracle or Netlsnr type is incorrect. Solution: Set the Home attribute value to the correct full path name of the Oracle home.

Table 8-1 Cluster Server agent for Oracle error messages (*continued*)

Message	Description and solution
File <env_file> is not a valid text file	<p>The file that the EnvFile attribute specifies for sourcing the environment variables is not present, not readable, or is not a text file.</p> <p>Solution: Set the EnvFile attribute value to the correct full path name. Ensure that the file format is valid.</p>
VCSAgExec returned failure when trying to execute in-depth test	<p>Internal error.</p> <p>Solution: Contact Technical Support for further assistance.</p>
Unable to open pipe from %s	<p>Internal error.</p> <p>Solution: Contact Technical Support for further assistance.</p>
Process <Process_name> restarted	<p>Warning message to indicate that the PID for the Oracle process that is specified is different than the one registered by the previous monitor cycle.</p>
Monitor procedure <monscript> returned <output>	<p>MonScript failed to execute correctly.</p> <p>Solution: Debug MonScript to assess the exact cause of failure.</p>
Monitor procedure <monscript> did not exit	<p>Internal error while executing MonScript.</p> <p>Solution: Contact Technical Support for further assistance.</p>
No owner for Oracle executables was specified	<p>The Owner attribute in the Oracle type has not been set.</p> <p>Solution: Set the Owner attribute value to the correct owner of the database binaries.</p>
Invalid owner <owner_name> for Oracle executables was specified	<p>The Operating System user that the Owner attribute specifies is invalid.</p> <p>Solution: Set the Owner attribute value to the correct owner of the database binaries.</p>
Access to Monscript <monscript> denied. Detail Monitoring will not be enabled!! Please specify a valid file.	<p>The file that the MonScript attribute specifies is not accessible or not found.</p> <p>Solution: Make sure that the file name indicates a valid and accessible file.</p>

Table 8-1 Cluster Server agent for Oracle error messages (*continued*)

Message	Description and solution
Encountered errors while decrypting password!	The agent cannot decrypt the password you specified. Solution: Use vscencrypt utility to create a new encrypted password and supply the password.

Error messages specific to the Oracle agent

Table 8-2 lists the error messages for the VCS agent for Oracle with the description and a recommended solution, if available.

Table 8-2 Oracle agent error messages

Message	Description and solution
No SID specified	The Sid attribute in the Oracle type has not been set. Solution: Set the Sid attribute value to the correct database instance.
sqlplus not found in <Oracle_Home>/bin	The client utility sqlplus is not found in the \$ORACLE_HOME/bin directory. Solution: Verify that the Oracle home has been correctly specified and that this executable is present.
srvctl not found in <Oracle_Home>/bin	The client utility srvctl is not found in the \$ORACLE_HOME/bin directory. Solution: Verify that the Oracle home has been correctly specified and that this executable is present.

Table 8-2 Oracle agent error messages (*continued*)

Message	Description and solution
Oracle <SID> failed to stop	<p>Warning message to indicate that the following commands were not successful in closing the Oracle instance in the clean or offline entry point:</p> <ul style="list-style-type: none"> ■ Shutdown immediate ■ Shutdown abort
Oracle database <SID> not running	<p>Warning message to indicate that the database instance was not running even before the clean or offline entry points were executed.</p> <p>Solution: No action required.</p>
Oracle (<SID>) kill TERM <PID1, PID2, ...>	<p>Warning message to indicate that the Oracle processes would be signaled with SIGTERM.</p> <p>Solution: No action required.</p>
Oracle (<SID>) kill KILL <PID1, PID2, ...>	<p>Warning message to indicate that the Oracle processes would be signaled with SIGKILL.</p> <p>Solution: No action required.</p>
Database in QUIESCING/QUIESCED mode	<p>Warning message to indicate that database is in QUIESCING or QUIESCED mode.</p>
Database in RESTRICTED mode	<p>Warning message to indicate that database is in RESTRICTED mode.</p>

Table 8-2 Oracle agent error messages (*continued*)

Message	Description and solution
Database in SUSPENDED state	Warning message to indicate that database is in SUSPENDED state.
Resource <Res_name> - monitor procedure did not complete within the expected time.	<p>Refer to Oracle's alert log for more information.</p> <p>When a monitor times out as many times as the value specified, the corresponding resource is brought down by calling the clean entry point. The resource is then marked FAULTED, or it is restarted, depending on the RestartLimit attribute value.</p> <p>Solution: Set the FaultOnMonitorTimeouts attribute value to 0 so that the monitor failures are not considered indicative of a resource fault.</p> <p>Another possible reason could be that automatic archiving was not enabled while setting up the database.</p> <p>Solution: Archive the database manually. If automatic archival is enabled, set the LOG_ARCHIVE_START parameter value in the file init.ora to TRUE.</p>
VCS ERROR Resource <Res_name>: online procedure did not complete within the expected time.	Solution: Increase the Oracle agent's OnlineTimeout interval.

Table 8-2 Oracle agent error messages (*continued*)

Message	Description and solution
<p>Custom script /opt/VRTSagents/ha/bin/Oracle/start_custom_<SID>.sql does not exist. Will not be able to start the database.</p>	<p>The agent could not find the custom script at the specified location to start the database.</p> <p>Solution: Make sure the custom file exists at the specified location and has valid permissions.</p>
<p>Custom script /opt/VRTSagents/ha/bin/Oracle/shut_custom_<SID>.sql does not exist. Using default shutdown option.</p>	<p>The agent could not find the custom script at the specified location to stop the database.</p> <p>Solution: Make sure the custom file exists and the specified location and has valid permissions.</p>
<p>oraerror.dat did not have records that could be parsed</p>	<p>The file oraerror.dat is not present or has records in an unsupported format.</p> <p>Solution: Make sure the file exists and has data in the supported format.</p>
<p>Incorrect Monitor Option</p>	<p>The MonitorOption value is less than 0 or greater than 1.</p> <p>Solution: Set the MonitorOption attribute value to 0 or 1.</p>
<p>MonitorOption value not applicable for this Oracle Version</p>	<p>The health check monitoring option is selected when Oracle version is not Oracle 10g or later.</p> <p>Solution: Set the MonitorOption value to 0 to select the process check monitoring option.</p>

Table 8-2 Oracle agent error messages (*continued*)

Message	Description and solution
VCSAgExec returned failure when trying to execute health check monitor test	Internal error. Solution: Contact Technical Support for further assistance.
VCSAgExec returned failure while trying to find Oracle version	Internal error. Solution: Contact Technical Support for further assistance.
One or more of the attributes User:Pword:Table:MonScript are not set correctly. Detail monitoring will not be enabled!!	Detail Monitoring has been enabled but the necessary attributes for detail monitoring have not been set correctly. Solution: Set the values of the required attributes for detail monitoring correctly.

Error messages specific to the Netlsnr agent

[Table 8-3](#) lists the Netlsnr agent error messages with the description and a recommended solution, if available.

Table 8-3 Netlsnr agent error messages

Message	Description and solution
Cannot open process directory.	The agent could not process the /proc entries in the particular monitor cycle. Solution: No action required.
Listener process <Listener> not running	Warning message to indicate that the Listener process was not running even before the clean or offline entry points were executed. Solution: No action required.

Table 8-3 Netlsnr agent error messages (*continued*)

Message	Description and solution
Listener <Listener> kill TERM <PID1, PID2, ...>	Warning message to indicate that the Listener process would be signaled with SIGTERM. Solution: No action required.
Listener <Listener> kill KILL <PID1, PID2, ...>	Warning message to indicate that the Listener process would be signaled with SIGKILL. Solution: No action required.
Isnrctl not found in <Oracle_Home>/bin	The client utility Isnrctl is not found in the \$ORACLE_HOME/bin directory. Solution: Verify that the Oracle home has been correctly specified and that this executable is present.
Isnrctl operation timed out	The tnslnsr process does not respond. Solution: Verify the underlying network protocol.

Error messages specific to the ASMIInst agent

[Table 8-4](#) lists the ASMIInst agent error messages with the description and a recommended solution, if available.

Table 8-4 ASMIInst agent error messages

Message	Description and solution
VCSAgExec failed while retrieving process name from ps output.	Internal error. Solution: Contact Technical Support for further assistance.
Attribute Home cannot have multiple tokens.	The Home attribute of the ASMIInst agent has multiple tokens. Solution: Correct the value of the Home attribute.

Table 8-4 ASMIInst agent error messages (*continued*)

Message	Description and solution
Cluster Synchronization Service process is not running.	The Oracle CSSD process is not running. Solution: Enable the CSSD process. See “Enabling the clustering daemon for ASM-managed database” on page 52.
Cluster Synchronization Service died or is restarted.	The Oracle CSSD process has died. Solution: Enable the CSSD process. See “Enabling the clustering daemon for ASM-managed database” on page 52.

Error messages specific to the ASMDG agent

[Table 8-5](#) lists the ASMIInst agent error messages with the description and a recommended solution, if available.

Table 8-5 ASMDG agent error messages

Message	Description and solution
No ASM Diskgroup name specified, or is null.	The DiskGroups attribute value for ASMDG agent is not specified. Solution: Specify the value of the Diskgroup attribute.
Agent unable to identify state of the ASMDG resource. The asm_diskstring parameter is not set correctly in the ASM initialization parameter file.	The ASM initialization parameter file does not have the correct value for ASM instance to search the ASM disk groups. Solution: Specify the correct value for the asm_diskstring parameter in the ASM initialization parameter file.
Agent unable to identify state of the resource.	The ASMDG agent cannot identify the state of the resource. Solution: Contact Technical Support for further assistance.

Troubleshooting issues specific to Oracle in a VCS environment

Table 8-6 lists any Oracle issues that you may encounter in a VCS environment.

Table 8-6 Oracle common issue

Message	Description and solution
ORA-15097: Cannot SHUTDOWN ASM instance with connected RDBMS instance.	This message appears for ASM-managed database if you try to offline the Oracle service group within 2-3 minutes after you brought the service group online. Refer to Oracle bug 5045309 for more information.

Verifying the Oracle health check binaries and intentional offline for an instance of Oracle

This section describes the steps to verify the state of Oracle instance, the Oracle health check binaries, and intentional offline behavior for an Oracle agent.

Note: The steps listed in the table below should be executed by the operating system user specified in Owner attribute.

Table 8-7 lists the checks you can verify with Oracle health check binaries.

Table 8-7 Oracle health check options

Verify	Solution
If the ORACLE_HOME variable is set.	Run the following command to verify that the \$ORACLE_HOME variable is correctly set. # echo \$ORACLE_HOME Set the ORACLE_HOME environment variable if it is not already set. For example, run the following command to set the ORACLE_HOME variable: # export ORACLE_HOME=/u01/oraHome where /u01/oraHome is the Oracle home directory path.

Verifying the Oracle health check binaries and intentional offline for an instance of Oracle

Table 8-7 Oracle health check options (*continued*)

Verify	Solution
If the SID for the Oracle instance is set.	<p>Run the following command to verify that the <code>\$ORACLE_SID</code> variable is correctly set.</p> <pre># echo \$ORACLE_SID</pre> <p>If required, run the <code># export ORACLE_SID=db</code> command to set the <code>\$ORACLE_SID</code> variable.</p>
If the library path is set.	<p>Run the following command to verify that the <code>\$LD_LIBRARY_PATH</code> variable is correctly set.</p> <pre># echo \$LD_LIBRARY_PATH</pre> <p>If required, run the <code># export LD_LIBRARY_PATH=\$ORACLE_HOME/lib:\$LD_LIBRARY_PATH</code> command to set the <code>\$LD_LIBRARY_PATH</code> variable.</p>
If the Oracle instance is online.	<p>Run the following command to verify that the Oracle health check binary reports the status correctly:</p> <pre># ./oraapi_<Arch>_<Oracle_FullVersion></pre> <p>where <code><Arch>_<Oracle_FullVersion></code> is the system architecture and the Oracle full version.</p> <p>For example, <code># ./oraapi_64_121010</code>, where <code>oraapi_64_121010</code> is a binary built for Oracle version 12 on a 64-bit system.</p> <p>The system displays the following message:</p> <pre>Instance is online</pre> <p>Run <code># echo \$?</code></p> <p>If the system displays 110 the instance is online and active.</p>

Table 8-7 Oracle health check options (*continued*)

Verify	Solution
<p>If the Oracle instance is offline.</p>	<p>Run the following command to verify that the Oracle health check binary reports the status correctly:</p> <pre># ./oraapi_<Arch>_<Oracle_FullVersion></pre> <p>where <Arch>_<Oracle_FullVersion> is the system architecture and the Oracle full version.</p> <p>For example, # ./oraapi_64_121010, where oraapi_64_121010 is a binary built for Oracle version 12 on a 64-bit system.</p> <p>The system displays the following message:</p> <pre>Failure: Instance Shutdown class: SOFT</pre> <p>Proper Shutdown</p> <pre>Run # echo \$?</pre> <p>If the system displays 100 the instance is offline.</p>
<p>If the Oracle instance is abnormally terminated.</p>	<p>If any of the important Oracle process is killed (for example, ora_pmon_<sid>, ora_smon_<sid>), the Oracle instance is abnormally terminated. In such a case run the Oracle health check binary, and the system displays following message:</p> <p>For example, run # ./oraapi_64_121010</p> <pre>Failure: Abnormal Termination class: SOFT Abnormal termination</pre> <pre>Run # echo \$?</pre> <p>If the system displays 98 the termination of Oracle processes. The VCS resource fails over as Oracle termination was abnormal and not intentional.</p>

Verifying the intentional offline behavior of the VCS Oracle agent

Perform the following steps to verify if VCS Oracle agent determines a graceful shutdown of a resource as intentional offline.

- 1 Bring the resource online.

```
# hares -online ORA_oraprod -sys system
```

Where *system* is the name of the system that has Oracle database installed.

- 2 Verify that `IntentionalOffline` and `MonitorOption` attributes are set to 1.
- 3 Stop the Oracle instance properly outside VCS control.

```
# su - oracle
# bash
# bash-3.00$ sqlplus "/" as sysdba"
```

The system displays the following message:

```
SQL> shutdown immediate
Database closed.
Database dismounted.
ORA-03113: end-of-file on communication channel
```

- 4 You can verify the log files to confirm if the shutdown was graceful.

Sample log message:
VCS INFO V-16-1-13470 Resource ORA_oraprod
(Owner: Unspecified, Group: ORA_PROD_Group) is offline on system.
(Intentional But NOT initiated by VCS)

Oracle agent has identified the Intentional offline for the resource.

Disabling IMF for a PDB resource

To disable IMF for PDB resource perform the following steps:

- 1 Override the IMF attribute for a particular PDB resource:

```
# hares -override <pdb_resname> IMF
```

- 2 Verify if the IMF attribute is overridden using the following command:

```
# hares -value pdb1 IMF
Mode      3      MonitorFreq      5      RegisterRetryLimit      3
```

3 Modify the Mode value to 0

```
# hares -modify pdb1 IMF -update Mode 0
```

4 Verify the Mode value using the following command:

```
# hares -value pdb1 IMF  
Mode      0      MonitorFreq      5      RegisterRetryLimit      3
```

Resource type definitions

This appendix includes the following topics:

- [About the resource type and attribute definitions](#)
- [Resource type definition for the Oracle agent](#)
- [Resource type definition for the Netlsnr agent](#)
- [Resource type definition for the ASMInst agent](#)
- [Resource type definition for the ASMDG agent](#)

About the resource type and attribute definitions

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 explain the attributes associated with the agent. The Required attributes explain the attributes that must be configured for the agent to function properly.

Resource type definition for the Oracle agent

The Oracle agent of the Cluster Server agent for Oracle is represented by the Oracle resource type in VCS.

```
type Oracle (  
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Oracle"  
    static keylist SupportedActions = { VRTS_GetInstanceName,  
        VRTS_GetRunningServices, DBRestrict, DBUndoRestrict,  
        DBResume, DBSuspend, DBTbspBackup, PDBMigrate_PreCheck,  
        PlugPDB, "home.vfd", "owner.vfd", "getid", "pfile.vfd" }  
    static str ArgList[] = { Sid, Owner, Home, Pfile, StartUpOpt,
```

```
    ShutDownOpt, DBAUser, DBAPword, EnvFile, AutoEndBkup,  
    User, Pword, Table, MonScript, Encoding, MonitorOption,  
    DBName, ManagedBy, PDBName }  
static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }  
static str IMFRegList[] = { Home, Owner, Sid, MonitorOption }  
str Sid  
str Owner  
str Home  
str Pfile  
str StartUpOpt = STARTUP_FORCE  
str ShutDownOpt = IMMEDIATE  
str DBName  
str ManagedBy = "ADMIN"  
str DBAUser  
str DBAPword  
str EnvFile  
boolean AutoEndBkup = 1  
str MonScript = "./bin/Oracle/SqlTest.pl"  
str User  
str Pword  
str Table  
str Encoding  
int MonitorOption = 0  
static boolean IntentionalOffline = 0  
static boolean AEPTimeout = 1  
static int OnlineWaitLimit = 5  
str PDBName  
)
```

Attribute definition for the Oracle agent

Review the description of the Oracle agent attributes. The agent attributes are classified as required, optional, and internal.

[Table A-1](#) lists the required attributes. You must assign values to the required attributes.

Table A-1 Required attributes for Oracle agent

Required attributes	Type and dimension	Definition
Sid	string-scalar	The variable \$ORACLE_SID that represents the Oracle instance. The Sid is considered case-sensitive by the Oracle agent and by the Oracle database server. For a policy managed database, the Sid attribute should be set to Sid prefix. See “About the Sid attribute in a policy managed database” on page 152.
Owner	string-scalar	The Oracle user who has privileges to start or stop the database instance. The agent also supports LDAP users as Oracle user.
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path.

[Table A-2](#) lists the optional attributes for Oracle agent. You can configure the optional attributes if necessary.

Table A-2 Optional attributes for Oracle agent

Optional Attributes	Type and Dimension	Definition
DBAUser	string-scalar	The database user who has sysdba privileges to start or stop the database.
DBAPword	string-scalar	Encrypted password for DBAUser. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the VCS Encrypt Utility (/opt/VRTSvcs/bin/vcsencrypt). See “Encrypting Oracle database user and listener passwords” on page 82.

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
StartUpOpt	string-scalar	<p>Startup options for the Oracle instance. This attribute can take the following values for traditional database and container database:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ RECOVERDB ■ SRVCTLSTART ■ CUSTOM ■ SRVCTLSTART_FORCE <p>Default is STARTUP_FORCE.</p> <p>See “Startup and shutdown options for the Oracle agent” on page 15.</p> <p>Startup options for the Oracle instance in pluggable database. This attribute can take the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_FORCE ■ RESTRICTED ■ OPEN_RO <p>See “Startup and shutdown options for the pluggable database (PDB)” on page 18.</p>
ShutDownOpt	string-scalar	<p>Shut down options for the Oracle instance. This attribute can take the following values for traditional database and container database:</p> <ul style="list-style-type: none"> ■ IMMEDIATE ■ TRANSACTIONAL ■ SRVCTLSTOP ■ CUSTOM ■ SRVCTLSTOP_TRANSACT ■ SRVCTLSTOP_ABORT ■ SRVCTLSTOP_IMMEDIATE <p>Default is IMMEDIATE.</p> <p>See “Startup and shutdown options for the Oracle agent” on page 15.</p> <p>Shut down options for the Oracle instance in pluggable database is IMMEDIATE.</p> <p>See “Startup and shutdown options for the pluggable database (PDB)” on page 18.</p>

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
EnvFile	string-scalar	<p>The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on.</p> <p>The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input.</p>
Pfile	string-scalar	<p>The name of the initialization parameter file with the complete path of the startup profile.</p> <p>You can also use the server parameter file. Create a one-line text initialization parameter file that contains only the SPFILE parameter. See the Oracle documentation for more information.</p> <p>Note: This attribute is applicable only for traditional and container databases.</p> <p>See “Using the SPFILE in a VCS cluster” on page 208.</p>
AutoEndBkup	boolean-scalar	<p>Setting the AutoEndBkup attribute to a non-zero value takes the datafiles in the database out of the backup mode, during Online.</p> <p>Default = 1</p> <p>Note: If a node fails during a hot backup of container database or pluggable database for Oracle 12C, you must set AutoEndBkup attribute of the corresponding CDB resource to 1. When the AutoEndBkup is set to 1 for the CDB, it also ends the backup of both CDB and PDB during Online.</p> <p>See “About VCS requirements for installing Oracle” on page 37.</p> <p>See “Failing over Oracle after a VCS node failure during hot backup” on page 42.</p>

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
MonitorOption	integer-scalar	<p>Monitor options for the Oracle instance. This attribute can take values 0 or 1.</p> <p>For traditional and container databases:</p> <ul style="list-style-type: none">■ 0—Process check monitoring (recommended)■ 1—Health check monitoring <p>For pluggable databases:</p> <ul style="list-style-type: none">■ 0—Basic monitoring <p>You must set the value of this attribute as 1 to use the intentional offline functionality of the agent.</p> <p>Default = 0</p> <p>The agent supports intelligent resource monitoring only for traditional and CDBs when the attribute value is set to 0.</p> <p>See "Monitor options for the Oracle agent in traditional database and container database" on page 20.</p> <p>See "Monitor for the pluggable database" on page 21.</p>

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
IMF	integer-association	

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
		<p>This resource-type level attribute determines whether the Oracle agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. <p>Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources <p>Default: 3</p> <ul style="list-style-type: none"> ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. <p>Default: 5</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources <ul style="list-style-type: none"> ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. <p>Default: 3</p> <p>Note: IMF is supported only in traditional and container databases.</p> <p>See “Enabling and disabling intelligent resource monitoring for agents</p>

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
		manually ” on page 89.
MonScript	string-scalar	<p>Pathname to the script provided for detail monitoring. The default (basic monitoring) is to monitor the database PIDs only.</p> <p>Note: Second-level monitoring is disabled if the value of the attribute MonScript is invalid or is set to an empty string.</p> <p>The pathname to the supplied detail monitor script is /opt/VRTSagents/ha/bin/Oracle/SqlTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with “./”, as in the path ./bin/Oracle/SqlTest.pl.</p>
User	string-scalar	Internal database user. Connects to the database for detail monitoring.
LevelTwoMonitorFreq	integer-scalar	<p>Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. You can also override the value of this attribute at resource-level.</p> <p>The value indicates the number of monitor cycles after which the agent will monitor Oracle in detail. For example, the value 5 indicates that the agent will monitor Oracle in detail every five online monitor intervals.</p> <p>If you manually upgraded to the VCS 7.1 agent, and if you had enabled detail monitoring in the previous version, then do the following:</p> <ul style="list-style-type: none"> ■ Set the value of the LevelTwoMonitorFreq attribute to the same value of that of the DetailMonitor attribute. <p>Note: If you set the AutoEndBkup attribute value to 0, then make sure that the LevelTwoMonitorFreq attribute value is 1 for detail monitoring.</p> <p>Default = 0</p>
Pword	string-scalar	<p>Encrypted password for internal database-user authentication.</p> <p>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the VCS Encrypt Utility (/opt/VRTSvcs/bin/vcscrypt).</p> <p>Note: This attribute is not applicable for PDB resources.</p> <p>See “Encrypting Oracle database user and listener passwords” on page 82.</p>
Table	string-scalar	Table for update by User/Pword.

Table A-2 Optional attributes for Oracle agent (*continued*)

Optional Attributes	Type and Dimension	Definition
Encoding	string-scalar	<p>Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output.</p> <p>Default is "".</p> <p>Note: This attribute is not applicable for PDB resources.</p>
IntentionalOffline		<p>This resource-type level attribute defines how VCS reacts when Oracle is intentionally stopped outside of VCS control.</p> <p>If you stop Oracle traditional and container database out of VCS control, the agent behavior is as follows:</p> <ul style="list-style-type: none"> ■ 0—The Oracle agent registers a fault and initiates the failover of the service group. ■ 1—The Oracle agent takes the Oracle resource offline when Health check monitoring is enabled. <p>If Health check monitoring is not enabled, the agent registers a fault and initiates the failover of the service group.</p> <p>Note: If you want to use the intentional offline functionality of the agent, you must set the value of the MonitorOption attribute as 1 to enable Health check monitoring in the CDB.</p> <p>Note: If a CDB resource is configured with IntentionalOffline, and the configured CDB is brought down outside the VCS control, then the PDB resources will also report as offline. To enable this functionality, set the IntentionalOffline attribute to 1 for the PDB resource.</p> <p>See the <i>Cluster Server Administrator's Guide</i>.</p>
DBName	string-scalar	<p>Set this attribute only when the database is a policy managed RAC database. The value of this attribute must be set to the database unique name.</p>
ManagedBy	string-scalar	<p>Default value for this attribute is ADMIN. In a policy managed RAC database this attribute must be set to POLICY.</p>
PDBName	string-scalar	<p>This attribute must be configured for a PDB. And the value should be set for a PDB database name. For traditional and CDB database, do not set this attribute.</p>

Table A-3 lists the internal attribute for Oracle agent. This attribute is for internal use only. Veritas recommends not to modify the value of this attribute.

Table A-3 Internal attributes for Oracle agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	Specifies the location of binaries, scripts, and other files related to the Oracle agent. Default is /opt/VRTSagents/ha/bin/Oracle.

About the Sid attribute in a policy managed database

The SID attribute is a required attribute. This section provides information to define the SID attribute in a policy managed database.

The SID prefix comprises of the first 8 alphanumeric characters of the database unique name. It can be a combination of letters a-z; uppercase and lowercase and numbers 0-9.

The SID prefix cannot have operating system special characters. Therefore, avoid the use of special characters in the first 8 characters of the database unique name. Special characters are omitted if used in the first 8 characters. There is a single SID prefix for every database. The SID prefix for a database must be unique within the cluster.

For an Oracle RAC database, each instance has a unique identifier, ORACLE_SID, which consists of the SID prefix and an instance number. The ORACLE_SID for Oracle RAC database instances is generated differently, depending on how you choose to manage the database. If you select a policy-managed database, then Oracle generates the SID in the format name_#, where name is the first eight alphanumeric characters of DB_UNIQUE_NAME, and # is the instance number. If you select an admin-managed database, then DBCA generates the SID for the instance names in advance, and the SID is in the format name#.

To find the Sid prefix name, run the following command:

```
# ${GRID_HOME}/bin/crsctl status resource ora.${DBName}.db -f | grep
GEN_USR_ORA_INST_NAME@ | tail -1 | sed 's/.*///' | sed 's/_[0-9]$//',
```

where *GRID_HOME* is grid home path and *DBName* is the database unique name.

Note: When a policy managed database is created, the Sid prefix is displayed on the confirmation page of the installation procedure.

See [“Attribute definition for the Oracle agent”](#) on page 143.

Resource type definition for the Netlsnr agent

The Netlsnr agent of the Cluster Server agent for Oracle is represented by the Netlsnr resource type in VCS.

```
type Netlsnr (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Netlsnr"
    static keylist SupportedActions = { VRTS_GetInstanceName,
        VRTS_GetRunningServices, "tnsadmin.vfd" }
    static str ArgList[] = { Owner, Home, TnsAdmin, Listener,
        EnvFile, MonScript, LsnrPwd, Encoding }
    static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }
    static str IMFRegList[] = { Home, Owner, Listener }
    str Owner
    str Home
    str TnsAdmin
    str Listener = "LISTENER"
    str EnvFile
    str MonScript = "./bin/Netlsnr/LsnrTest.pl"
    str LsnrPwd
    str Encoding
    static boolean IntentionalOffline = 0
)
```

Attribute definition for the Netlsnr agent

Review the description of the Netlsnr agent attributes. The agent attributes are classified as required, optional, and internal.

[Table A-4](#) lists the required attributes for Netlsnr agent. You must assign values to the required attributes.

Table A-4 Required attributes for Netlsnr agent

Required attributes	Type and dimension	Definition
Owner	string-scalar	The Oracle user who has privileges to start or stop the listener process. The agent also supports LDAP users as Oracle user.

Table A-4 Required attributes for Netlsnr agent (*continued*)

Required attributes	Type and dimension	Definition
Home	string-scalar	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Do not append a slash (/) at the end of the path.

[Table A-5](#) lists the optional attributes for Netlsnr agent. You can configure the optional attributes if necessary.

Table A-5 Optional attributes for Netlsnr agent

Optional attributes	Type and dimension	Definition
TnsAdmin	string-scalar	The \$TNS_ADMIN path to directory in which the Listener configuration file resides (listener.ora). Default is /var/opt/oracle.
Listener	string-scalar	Name of Listener. The name for Listener is considered case-insensitive by the Netlsnr agent and the Oracle database server. Default is LISTENER.
LsnrPwd	string-scalar	The VCS encrypted password used to stop and monitor the listener. This password is set in the Listener configuration file. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the VCS Encrypt utility. See “Encrypting Oracle database user and listener passwords” on page 82.
EnvFile	string-scalar	Specifies the full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle listener environment such as LD_LIBRARY_PATH and so on. The syntax for the contents of the file depends on the login shell of Owner. This file must be readable by Owner. The file must not contain any prompts for user input.

Table A-5 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
IMF	integer-association	<p>This resource-type level attribute determines whether the Netlsnr agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. <p>Valid values are as follows:</p> <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources <p>Default: 3</p> <ul style="list-style-type: none"> ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. <p>Default: 5</p> <p>You can set this attribute to a non-zero value in some cases where the agent requires to perform poll-based resource monitoring in addition to the intelligent resource monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources <ul style="list-style-type: none"> ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>netlsnr_imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. <p>Default: 3</p> <p>See “Enabling and disabling intelligent resource monitoring for agents manually” on page 89.</p>

Table A-5 Optional attributes for Netlsnr agent (*continued*)

Optional attributes	Type and dimension	Definition
MonScript	string-scalar	<p>Pathname to the script provided for detail monitoring. By default, the detail monitoring is enabled to monitor the listener process.</p> <p>Note: If the value of the attribute MonScript is set to an empty string, the agent disables detail monitoring.</p> <p>The pathname to the supplied detail monitoring script is /opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Netlsnr/LsnrTest.pl.</p>
LevelTwoMonitorFreq	integer-scalar	<p>Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring.</p> <p>If you enabled detail monitoring, then set the value of the LevelTwoMonitorFreq attribute.</p> <p>Default = 0</p>
Encoding	string-scalar	<p>Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output.</p> <p>Default is "".</p>
IntentionalOffline		<p>For future use.</p> <p>Do not change the value of this attribute.</p> <p>Default = 0</p>

[Table A-6](#) lists the internal attribute for Netlsnr agent. This attribute is for internal use only. Veritas recommends not to modify the value of this attribute.

Table A-6 Internal attributes for Netlsnr agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	<p>Specifies the location of binaries, scripts, and other files related to the Netlsnr agent.</p> <p>Default is /opt/VRTSagents/ha/bin/Netlsnr.</p>

Resource type definition for the ASMinst agent

The ASMinst agent of the Cluster Server agent for Oracle is represented by the ASMinst resource type in VCS.

```
type ASMinst (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/ASMinst"
    static str ArgList[] = { Sid, Owner, Home, DBAUser,
        DBAPword, Pfile, StartUpOpt, ShutDownOpt,
        EnvFile, Encoding, MonitorOption }
    str Sid
    str Owner
    str Home
    str DBAUser
    str DBAPword
    str Pfile
    str StartUpOpt
    str ShutDownOpt
    str EnvFile
    str Encoding
    int MonitorOption = 0
)
```

Attribute definition for the ASMinst agent

Review the description of the ASMinst agent attributes. The agent attributes are classified as required, optional, and internal.

[Table A-7](#) lists the required attributes. You must assign values to the required attributes.

Table A-7 Required attributes for ASMinst agent

Required attributes	Type and dimension	Definition
Sid	string-scalar	The variable \$ORACLE_SID that represents the ASM instance. The Sid is considered case-sensitive by the ASMinst agent.
Owner	string-scalar	The Oracle user who has privileges to start or stop ASM instance. The agent also supports LDAP users as Oracle user.

Table A-7 Required attributes for ASMinst agent (*continued*)

Required attributes	Type and dimension	Definition
Home	string-scalar	<p>The \$ORACLE_HOME path to Oracle ASM binaries and configuration files. For example, you could specify the path as /opt/ora_home.</p> <p>Note: Do not append a slash (/) at the end of the path.</p>

[Table A-8](#) lists the optional attributes for ASMinst agent. You can configure the optional attributes if necessary.

Table A-8 Optional attributes for ASMinst agent

Optional Attributes	Type and Dimension	Definition
EnvFile	string-scalar	<p>The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on.</p> <p>The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input.</p>
Pfile	string-scalar	<p>The name of the initialization parameter file of ASM instance with the complete path of the startup profile.</p> <p>You can also use the server parameter file. Create a one-line text initialization parameter file that contains only the SPFILE parameter. See the Oracle documentation for more information.</p> <p>See “Using the SPFILE in a VCS cluster” on page 208.</p>

Table A-8 Optional attributes for ASMinst agent (*continued*)

Optional Attributes	Type and Dimension	Definition
MonitorOption	integer-scalar	<p>Monitor options for the ASM instance. This attribute can take values 0 or 1.</p> <ul style="list-style-type: none"> ■ 0—Process check monitoring (recommended) ■ 1—Health check monitoring <p>You must set the value of this attribute as 1 to use the intentional offline functionality of the agent.</p> <p>Default = 0</p> <p>See “Monitor options for the ASMinst agent” on page 31.</p>
DBAUser	string-scalar	<p>The ASM user who has sysasm privileges to start or stop the ASM instance. You can create ASM users for Oracle 11g R1 and later.</p>
DBAPword	string-scalar	<p>Encrypted password for DBAUser.</p> <p>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the VCS Encrypt utility.</p> <p>See “Encrypting Oracle database user and listener passwords” on page 82.</p>
Encoding	string-scalar	<p>Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".</p>

Table A-8 Optional attributes for ASMIInst agent (*continued*)

Optional Attributes	Type and Dimension	Definition
StartUpOpt	string-scalar	<p>Startup option for the Oracle ASM instance. The agent uses the sqlplus command to start the Oracle ASM instance.</p> <p>This attribute can take the following values:</p> <ul style="list-style-type: none"> ■ STARTUP ■ STARTUP_MOUNT ■ STARTUP_OPEN ■ SRVCTLSTART ■ SRVCTLSTART_MOUNT ■ SRVCTLSTART_OPEN <p>Default is "STARTUP".</p> <p>Note: recommends to use STARTUP or SRVCTLSTART option for ASMIInst resource.</p>
ShutDownOpt	string-scalar	<p>Shut down option for the Oracle ASM instance. By default the agent uses the sqlplus command to stop the Oracle ASM instance. To stop the ASM instance using srvctl utility set the option to SRVCTLSTOP.</p> <p>Default is "".</p>

[Table A-9](#) lists the internal attribute for ASMIInst agent. This attribute is for internal use only. Veritas recommends not to modify the value of this attribute.

Table A-9 Internal attributes for ASMIInst agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	<p>Specifies the location of binaries, scripts, and other files related to the ASMIInst agent.</p> <p>Default is /opt/VRTSagents/ha/bin/ASMIInst.</p>

Resource type definition for the ASMDG agent

The ASMDG agent is represented by the ASMDG resource type in VCS. The following extract shows the type definition of the ASMDG resource in the OracleASMTTypes.cf file.

```
type ASMDG (  
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/ASMDG"  
    static str ArgList[] = { Sid, Owner, Home, DBAUser,  
        DBAPword, DiskGroups, EnvFile, Encoding }  
    static boolean IntentionalOffline = 1  
    str Sid  
    str Owner  
    str Home  
    str DBAUser  
    str DBAPword  
    keylist DiskGroups  
    str EnvFile  
    str Encoding  
)
```

Attribute definition for the ASMDG agent

Review the description of the ASMDG agent attributes. The agent attributes are classified as required, optional, and internal.

[Table A-10](#) lists the required attributes. You must assign values to the required attributes.

Table A-10 Required attributes for ASMDG agent

Required attributes	Type and dimension	Definition
DiskGroups	keylist	The ASM disk groups, where you store the Oracle database files.
Sid	string-scalar	The variable \$ORACLE_SID that represents the ASM instance. The Sid is considered case-sensitive by the ASMinst agent.

Table A-10 Required attributes for ASMDG agent (*continued*)

Required attributes	Type and dimension	Definition
Owner	string-scalar	The Oracle user who has privileges to mount or unmount the ASM disk group. The agent also supports LDAP users as Oracle user.
Home	string-scalar	The \$ORACLE_HOME path to Oracle ASM binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path.

[Table A-11](#) lists the optional attributes for ASMDG agent. You can configure the optional attributes if necessary.

Table A-11 Optional attributes for ASMDG agent

Optional Attributes	Type and Dimension	Definition
DBAUser	string-scalar	The ASM user who has sysasm privileges to start or stop the ASM instance. You can create ASM users for Oracle 11g R1 and later.
DBAPword	string-scalar	Encrypted password for DBAUser. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the VCS Encrypt utility. See “Encrypting Oracle database user and listener passwords” on page 82.

Table A-11 Optional attributes for ASMDG agent (*continued*)

Optional Attributes	Type and Dimension	Definition
EnvFile	string-scalar	The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on. The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input.
Encoding	string-scalar	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".

[Table A-12](#) lists the internal attribute for ASMDG agent. This attribute is for internal use only. Veritas recommends not to modify the value of this attribute.

Table A-12 Internal attributes for ASMDG agent

Optional Attributes	Type and Dimension	Definition
AgentDirectory	static-string	Specifies the location of binaries, scripts, and other files related to the ASMDG agent. Default is /opt/VRTSagents/ha/bin/ASMDG.

Sample configurations

This appendix includes the following topics:

- [About the sample configurations for Oracle enterprise agent](#)
- [Sample single Oracle instance configuration](#)
- [Sample multiple Oracle instances \(single listener\) configuration](#)
- [Sample multiple instance \(multiple listeners\) configuration](#)
- [Sample Oracle configuration with shared server support](#)
- [Sample Oracle ASM configurations](#)
- [Sample configuration of Oracle pluggable database \(PDB\) resource in main.cf](#)
- [Sample configuration of migratable Oracle pluggable database \(PDB\) resource in main.cf](#)

About the sample configurations for Oracle enterprise agent

The sample configuration includes descriptions for typical service groups that are configured to monitor the state of Oracle in a VCS cluster.

See [“Configuring Oracle instances in VCS”](#) on page 68.


The sample dependency graphs depict the resource types, resources, and resource dependencies within the service group. The sample configuration file (main.cf) is also included for your reference.

Review these dependencies carefully before configuring the agent. For more information about VCS resource types, see the *Cluster Server Bundled Agents Reference Guide*.

Sample single Oracle instance configuration

Figure B-1 describes a typical service group configured to monitor the state of an Oracle instance in a VCS cluster.

Figure B-1 Dependency graph for single Oracle instance



The shared disk groups and volumes in the cluster are configured as resources of type DiskGroup and Volume respectively. The volumes are mounted using the Mount agent. The virtual IP address for the service group is configured using the IP and NIC resource types. The Oracle server can be started after each of these resources is brought online.

If your configuration does not use Veritas Volume Manager, use the DiskReservation resource type along with LVMVolumeGroup and LVMLogicalVolume resource types to configure shared storage instead of the DiskGroup and Volume resource types.

Sample VCS configuration file for single Oracle instance

Review the sample configuration with a resource of type Oracle that is configured as follows in main.cf file.

```
include "types.cf"
include "OracleTypes.cf"

cluster vcs (
)

system sys1 (
)

system sys2 (
)

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

    DiskGroup DG_oraprod (
        DiskGroup = ora_prod_dg
        StartVolumes = 0
        StopVolumes = 0
    )

    IP IP_oraprod (
        Device = eth0
        Address = "192.168.1.22"
        Netmask = "255.255.255.0"
    )

    Mount Mount_oraprod_u01 (
        MountPoint = "/prod/u01"
        BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol"
        FSType = vxfs
        FsckOpt = "-n"
    )

    Mount Mount_oraprod_u02 (
        MountPoint = "/prod/u02"
        BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol"
        FSType = vxfs
        FsckOpt = "-n"
    )
)
```

```
NIC NIC_oraprod (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)

Netlsnr LSNR_oraprod_lsnr (
    Owner = oraprod
    Home = "/orahome/Oracle"
    TnsAdmin = "/orahome/Oracle/network/admin"
    Listener = LISTENER_PROD
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
    LsnrPwd = cqfOdoOolOo
)

Oracle ORA_oraprod (
    Sid = PROD
    Owner = oraprod
    Home = "/orahome/Oracle"
    EnvFile = "/tmp/env.sh"
    MonScript = "./bin/Oracle/SqlTest.pl"
    User = thor
    Pword = hvlTptWvj

    DBAUser = "orauser"
    DBAPword = "DPSrFPeRIrGPiRF"

    Table = thor
    MonitorOption = 0
)

Volume Vol_oraprod_vol1 (
    Volume = u01-vol
    DiskGroup = ora_prod_dg
)

Volume Vol_oraprod_vol2 (
    Volume = u02-vol
    DiskGroup = ora_prod_dg
)

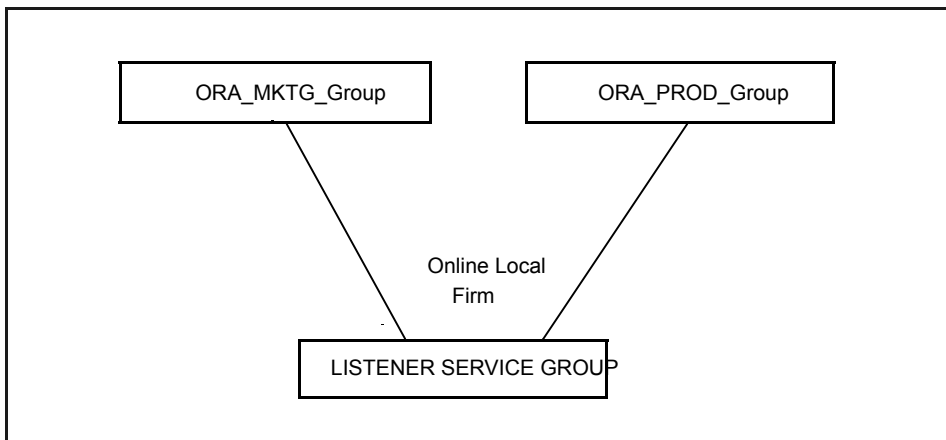
IP_oraprod requires NIC_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
Mount_oraprod_u01 requires Vol_oraprod_vol1
```

```
Mount_oraprod_u02 requires Vol_oraprod_vol2
ORA_oraprod requires Mount_oraprod_u01
ORA_oraprod requires Mount_oraprod_u02
Vol_oraprod_vol1 requires DG_oraprod
Vol_oraprod_vol2 requires DG_oraprod
```

Sample multiple Oracle instances (single listener) configuration

Figure B-2 describes a typical VCS configuration to monitor two Oracle instances sharing a listener. This configuration has a service group for each Oracle instance. The Listener too is configured in a separate service group.

Figure B-2 Two Oracle instances sharing a listener



The Oracle service groups are made dependent on the Listener service group using an Online Local Firm dependency.

Figure B-3 shows the dependency graph for one of the Oracle instances in the VCS configuration. In the Oracle service group, the shared disk groups and volumes in the cluster are configured as resources of type DiskGroup and Volume respectively. The volumes are mounted using the Mount agent.

Figure B-3 Dependency graph for one of the Oracle instances

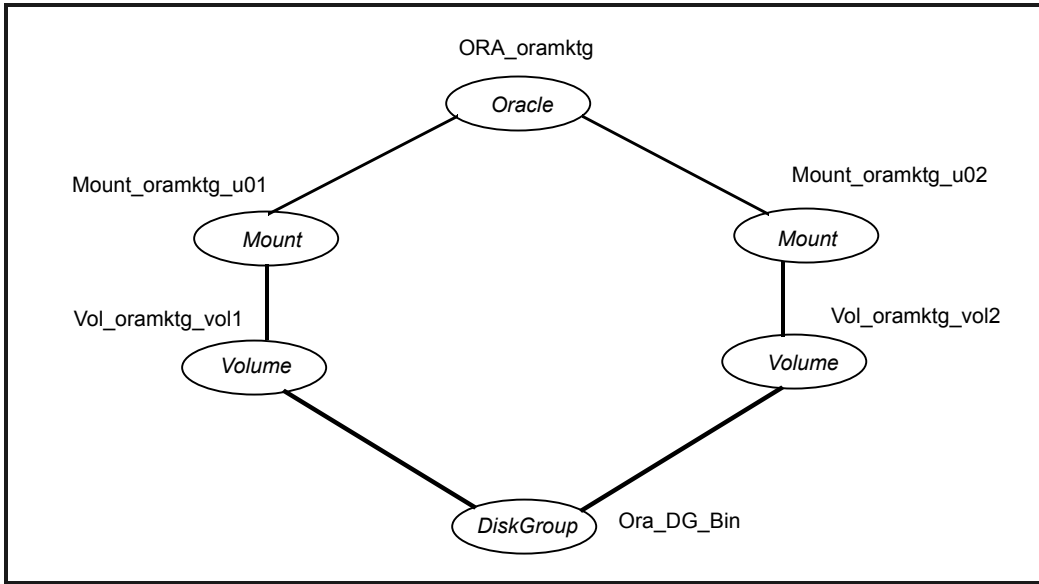
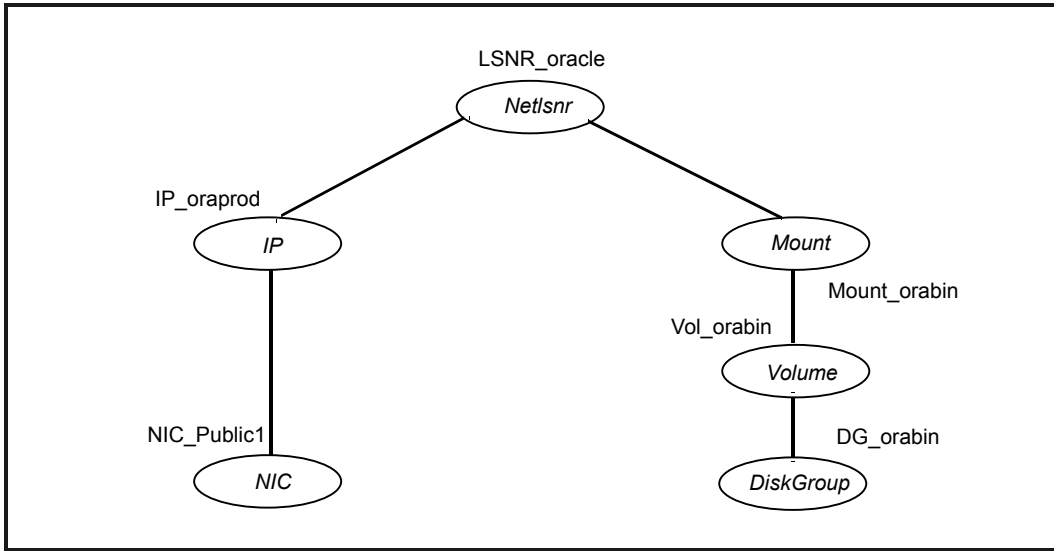


Figure B-4 shows the dependency graph for the listener that the two Oracle instances share in the VCS configuration. In the Listener service group, the virtual IP address is configured using the IP and NIC resource types. The Listener can be started after the IP and NIC resources are brought online.

Figure B-4 Dependency graph for the single listener



The Oracle server can be started after the Listener service group and the resources in the Oracle service group are brought online.

If your configuration does not use Veritas Volume Manager, use the DiskReservation resource type along with LVMVolumeGroup and LVMLogicalVolume resource types to configure shared storage instead of the DiskGroup and Volume resource types.

Note: In this case, make sure you have modified all proper system files, such as /etc/system, /etc/passwd, /etc/group, and /etc/shadow to support multiple databases. Pay particular attention to system requirements like physical memory and shared memory segment availability. Also ensure a single system is capable of sustaining a multiple instance load in the event of a server failure and extended operation on the backup server.

Sample VCS configuration file for multiple Oracle instances (single listener)

Review the sample configuration file for multiple Oracle instances.

```

include "types.cf"
include "OracleTypes.cf"

cluster vcs (

```

```
)  
system sys1 (  
)  
  
system sys2 (  
)  
  
group ORA_MKTG_Group (  
    SystemList = { sys1 = 0, sys2 = 1 }  
    AutoStartList = { sys1 }  
)  
  
    DiskGroup DG_oramktg (  
        DiskGroup = ora_mktg_dg  
        StartVolumes = 0  
        StopVolumes = 0  
    )  
  
        Mount Mount_oramktg_u01 (  
            MountPoint = "/mktg/u01"  
            BlockDevice = "/dev/vx/dsk/ora_mktg_dg/u01-vol"  
            FSType = vxfs  
            FsckOpt = "-n"  
        )  
  
        Mount Mount_oramktg_u02 (  
            MountPoint = "/mktg/u02"  
            BlockDevice = "/dev/vx/dsk/ora_mktg_dg/u02-vol"  
            FSType = vxfs  
            FsckOpt = "-n"  
        )  
  
    Oracle ORA_oramktg (  
        Sid = MKTG  
        Owner = oramktg  
        Home = "/orahome/Oracle"  
        EnvFile = "/tmp/env.sh"  
        MonScript = "./bin/Oracle/SqlTest.pl"  
        User = thor  
        Pword = hvlTptWvj  
        Table = thor  
        MonitorOption = 0  
    )  
)
```

```
Volume Vol_oramktg_vol1 (
    Volume = u01-vol
    DiskGroup = ora_mktg_dg
)

Volume Vol_oramktg_vol2 (
    Volume = u02-vol
    DiskGroup = ora_mktg_dg
)

requires group Common_Service online local firm
Mount_oramktg_u01 requires Vol_oramktg_vol1
Mount_oramktg_u02 requires Vol_oramktg_vol2
ORA_oramktg requires Mount_oramktg_u01
ORA_oramktg requires Mount_oramktg_u02
Vol_oramktg_vol1 requires DG_oramktg
Vol_oramktg_vol2 requires DG_oramktg

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_oraprod (
    DiskGroup = ora_prod_dg
    StartVolumes = 0
    StopVolumes = 0
)

Mount Mount_oraprod_u01 (
    MountPoint = "/prod/u01"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oraprod_u02 (
    MountPoint = "/prod/u02"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol"
    FSType = vxfs
    FsckOpt = "-n"
```

```
)

Oracle ORA_oraprod (
    Sid = PROD
    Owner = oraprod
    Home = "/orahome/Oracle"
    EnvFile = "/tmp/env.sh"
    MonScript = "./bin/Oracle/SqlTest.pl"
    User = thor
    Pword = hvlTptWvj
    Table = thor
    MonitorOption = 0
)

Volume Vol_oraprod_vol1 (
    Volume = u01-vol
    DiskGroup = ora_prod_dg
)

Volume Vol_oraprod_vol2 (
    Volume = u02-vol
    DiskGroup = ora_prod_dg
)

requires group Common_Service online local firm
Mount_oraprod_u01 requires Vol_oraprod_vol1
Mount_oraprod_u02 requires Vol_oraprod_vol2
ORA_oraprod requires Mount_oraprod_u01
ORA_oraprod requires Mount_oraprod_u02
Vol_oraprod_vol1 requires DG_oraprod
Vol_oraprod_vol2 requires DG_oraprod

group Common_Service (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_orabin (
    DiskGroup = ora_bin_dg
    StartVolumes = 0
    StopVolumes = 0
)
```

```
IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.22"
    Netmask = "255.255.255.0"
)

Mount Mount_orabin (
    MountPoint = "/orahome/Oracle"
    BlockDevice = "/dev/vx/dsk/ora_bin_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

NIC NIC_Public1 (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)

Netlsnr LSNR_oracle (
    Owner = oracle
    Home = "/orahome/Oracle"
    TnsAdmin = "/orahome/Oracle/network/admin"
    Listener = LISTENER_ORACLE
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
)

Volume Vol_orabin (
    Volume = u01-vol
    DiskGroup = ora_bin_dg
)

IP_oraprod requires NIC_Public1
LSNR_oracle requires IP_oraprod
LSNR_oracle requires Mount_orabin
Mount_orabin requires Vol_orabin
Vol_orabin requires DG_orabin
```

Sample multiple instance (multiple listeners) configuration

This configuration has several single-instance configurations. Each Oracle instance is configured in a separate service group. The resource dependencies are similar to the single Oracle instance configuration.

See [“Sample single Oracle instance configuration”](#) on page 165.

Sample VCS configuration file for multiple Oracle instances (multiple listeners)

Review the sample configuration file for the multiple Oracle instances that have multiple listeners.

```
include "types.cf"
include "OracleTypes.cf"

cluster vcs (
)

system sys1 (
)

system sys2 (
)

group ORA_MKTG_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_oramktg (
    DiskGroup = ora_mktg_dg
    StartVolumes = 0
    StopVolumes = 0
)

IP IP_oramktg (
    Device = eth0
    Address = "192.168.1.22"
    Netmask = "255.255.255.0"
)
```

```
Mount Mount_oramktg_u01 (
    MountPoint = "/mktg/u01"
    BlockDevice = "/dev/vx/dsk/ora_mktg_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oramktg_u02 (
    MountPoint = "/mktg/u02"
    BlockDevice = "/dev/vx/dsk/ora_mktg_dg/u02-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Netlsnr LSNR_oramktg_lsnr (
    Owner = oramktg
    Home = "/orahome/Oracle"
    TnsAdmin = "/orahome/Oracle/network/admin"
    Listener = LISTENER_MKTG
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
    LsnrPwd = cqfOdoOolOo
)

Oracle ORA_oramktg (
    Sid = MKTG
    Owner = oramktg
    Home = "/orahome/Oracle"
    EnvFile = "/tmp/env.sh"
    MonScript = "./bin/Oracle/SqlTest.pl"
    User = thor
    Pword = hvlTptWvj
    Table = thor
    MonitorOption = 0
)

Proxy NICProxy_oramktg (
    TargetResName = NIC_Public1
)

Volume Vol_oramktg_voll (
    Volume = u01-vol
    DiskGroup = ora_mktg_dg
```



```
)

Volume Vol_oramktg_vol2 (
    Volume = u02-vol
    DiskGroup = ora_mktg_dg
)

IP_oramktg requires NICProxy_oramktg
LSNR_oramktg_lsnr requires IP_oramktg
LSNR_oramktg_lsnr requires ORA_oramktg
Mount_oramktg_u01 requires Vol_oramktg_vol1
Mount_oramktg_u02 requires Vol_oramktg_vol2
ORA_oramktg requires Mount_oramktg_u01
ORA_oramktg requires Mount_oramktg_u02
Vol_oramktg_vol1 requires DG_oramktg
Vol_oramktg_vol2 requires DG_oramktg

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_oraprod (
    DiskGroup = ora_prod_dg
    StartVolumes = 0
    StopVolumes = 0
)

IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.21"
    Netmask = "255.255.255.0"
)

Mount Mount_oraprod_u01 (
    MountPoint = "/prod/u01"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oraprod_u02 (
    MountPoint = "/prod/u02"
```

```
BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol"
FSType = vxfs
FsckOpt = "-n"
)

Netlsnr LSNR_oraprod_lsnr (
  Owner = oraprod
  Home = "/orahome/Oracle"
  TnsAdmin = "/orahome/Oracle/network/admin"
  Listener = LISTENER_PROD
  MonScript = "./bin/Netlsnr/LsnrTest.pl"
  LsnrPwd = cqfOdoOolOo
)

Oracle ORA_oraprod (
  Sid = PROD
  Owner = oraprod
  Home = "/orahome/Oracle"
  EnvFile = "/tmp/env.sh"
  MonScript = "./bin/Oracle/SqlTest.pl"
  User = thor
  Pword = hvlTptWvj
  Table = thor
  MonitorOption = 0
)

Proxy NICProxy_oraprod (
  TargetResName = NIC_Public1
)

Volume Vol_oraprod_vol1 (
  Volume = u01-vol
  DiskGroup = ora_prod_dg
)

Volume Vol_oraprod_vol2 (
  Volume = u02-vol
  DiskGroup = ora_prod_dg
)

IP_oraprod requires NICProxy_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
```

```
Mount_oraprod_u01 requires Vol_oraprod_vol1
Mount_oraprod_u02 requires Vol_oraprod_vol2
ORA_oraprod requires Mount_oraprod_u01
ORA_oraprod requires Mount_oraprod_u02
Vol_oraprod_vol1 requires DG_oraprod
Vol_oraprod_vol2 requires DG_oraprod


group Parallel_Service (
    SystemList = { sys1 = 0, sys2 = 1 }
    Parallel = 1
    AutoStartList = { sys1 }
)

NIC NIC_Public1 (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)
```

Sample Oracle configuration with shared server support

[Figure B-5](#) describes a typical service group configured to monitor Oracle with shared server support.

Figure B-5 Dependency for Oracle configured with shared server support



The shared disk groups and volumes in the cluster are configured as resources of type `DiskGroup` and `Volume` respectively. The volumes are mounted using the `Mount` agent. The virtual IP address for the service group is configured using the `IP` and `NIC` resource types. The Oracle server can be started after each of these resources is brought online.

If your configuration does not use Veritas Volume Manager, use the `DiskReservation` resource type along with `LVMVolumeGroup` and `LVMLogicalVolume` resource types to configure shared storage instead of the `DiskGroup` and `Volume` resource types.

Sample VCS configuration file for Oracle instance configured with shared server support

Review the configuration file for an Oracle instance that is configured with shared server support.

```

include "types.cf"
include "OracleTypes.cf"

cluster vcs (

```

```
)

system sys1 (
)

system sys2 (
)

group ORA_PROD_Group (
    SystemList = { sys1 = 0, sys2 = 1 }
    AutoStartList = { sys1 }
)

DiskGroup DG_oraprod (
    DiskGroup = ora_prod_dg
    StartVolumes = 0
    StopVolumes = 0
)

IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.22"
    Netmask = "255.255.255.0"
)

Mount Mount_oraprod_u01 (
    MountPoint = "/prod/u01"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u01-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

Mount Mount_oraprod_u02 (
    MountPoint = "/prod/u02"
    BlockDevice = "/dev/vx/dsk/ora_prod_dg/u02-vol"
    FSType = vxfs
    FsckOpt = "-n"
)

NIC NIC_ORAPROD (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)
```

```
Netlsnr LSNR_oraprod_lsnr (  
    Owner = oraprod  
    Home = "/orahome/Oracle"  
    TnsAdmin = "/orahome/Oracle/network/admin"  
    Listener = LISTENER_PROD  
    MonScript = "./bin/Netlsnr/LsnrTest.pl"  
    LsnrPwd = cqfOdoOolOo  
)  
  
Oracle ORA_oraprod (  
    Sid = PROD  
    Owner = oraprod  
    Home = "/orahome/Oracle"  
    EnvFile = "/tmp/env.sh"  
    MonScript = "./bin/Oracle/SqlTest.pl"  
    User = thor  
    Pword = hvlTptWvj  
    Table = thor  
    MonitorOption = 0  
)  
  
Volume Vol_oraprod_vol1 (  
    Volume = u01-vol  
    DiskGroup = ora_prod_dg  
)  
  
Volume Vol_oraprod_vol2 (  
    Volume = u02-vol  
    DiskGroup = ora_prod_dg  
)  
  
IP_oraprod requires NIC_ORAPROD  
LSNR_oraprod_lsnr requires IP_oraprod  
Mount_oraprod_u01 requires Vol_oraprod_vol1  
Mount_oraprod_u02 requires Vol_oraprod_vol2  
ORA_oraprod requires IP_oraprod  
ORA_oraprod requires Mount_oraprod_u01  
ORA_oraprod requires Mount_oraprod_u02  
Vol_oraprod_vol1 requires DG_oraprod  
Vol_oraprod_vol2 requires DG_oraprod
```

Sample Oracle ASM configurations

Review the dependency graphs for the service group configured to monitor the state of an Oracle instance that is ASM-managed in a VCS cluster. You can have the following Oracle ASM configurations in a VCS environment:

- ASM disks as raw disks
- ASM disks as Veritas Volume Manager volumes
- ASM disks as Veritas Cluster Volume Manager volumes

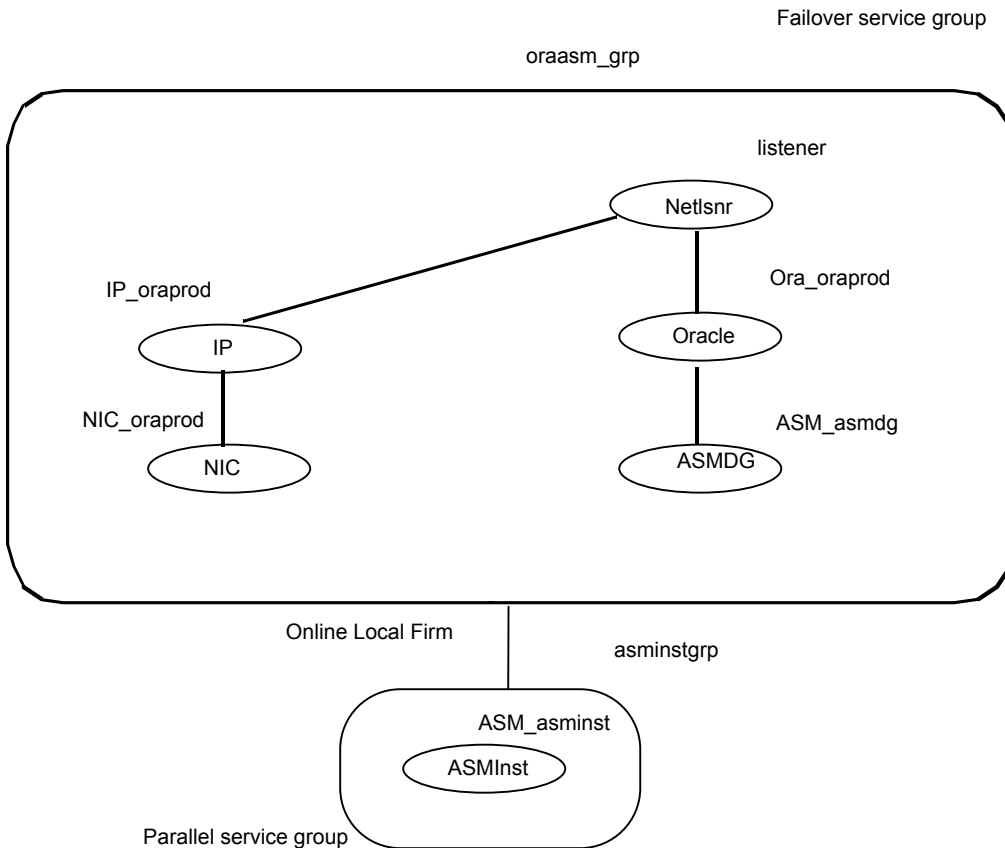
Sample configuration for ASM disks as raw disks

If you use raw devices as ASM disks, you can configure your service group in the following way:

- Oracle and ASMDG resources as parent failover service groups, and ASMInst resource as parallel service group
- The service groups are made dependent using an Online Local Firm dependency.

[Figure B-6](#) describes a typical service group with ASMInst resource as a parallel service group.

Figure B-6 Dependency graph with ASMInst resource as a parallel service group



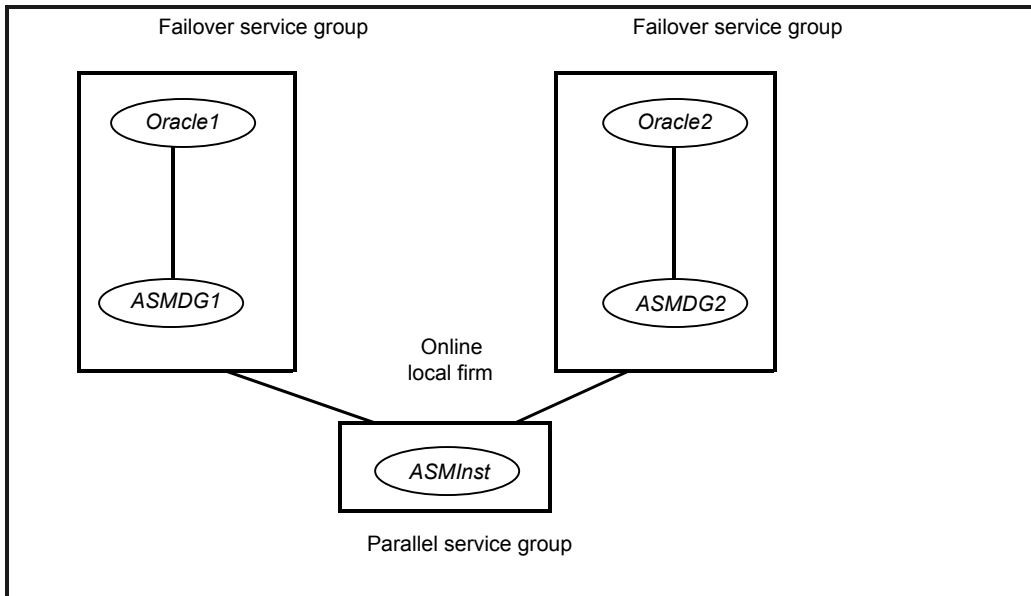
The Oracle ASM instance and ASM disk groups in the cluster are configured as resources of type ASMInst and ASMDG respectively. The ASMInst agent is configured as parallel service group, asminstgrp.

The virtual IP address for the service group is configured using the IP and NIC resource types. The Oracle and ASMDG resources are configured as failover service group, oraasm_grp. The Oracle server can be started after each of these resources is brought online.

The oraasm_grp is made dependent on the asminstgrp using an Online Local Firm dependency.

[Figure B-7](#) describes a typical service group with multiple Oracle instances sharing an ASMInst resource that is configured as a parallel service group.

Figure B-7 Dependency graph for Oracle ASM with multiple Oracle instances on a node



If you have multiple Oracle instances, a failover service group is configured for each Oracle instance. The Oracle service groups share a single ASM instance that is configured as a parallel service group. The Oracle service groups are made dependent on the ASMInst service group using an Online Local Firm dependency. However, each database must use exclusive ASM disk groups, so that the Cluster Server agent for Oracle can fail over the disk group.

Sample VCS configuration file for ASM-managed Oracle instance

Review the sample configuration for an Oracle instance that is ASM-managed. The sample file has the ASMInst resource as part of a parallel service group.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTTypes.cf"

cluster vcs (
)

system symnode01 (
```

```
)

system symnode02 (
)

group asminstgrp (
    SystemList = { symnode01 = 0, symnode02 = 1 }
    Parallel = 1
)

    ASMInst ASM_asminst (
        Sid = "+ASM"
        Owner = "oraprod"
        Home = "/orahome/Oracle"
    )
group oraasm_grp (
    SystemList = { symnode01 = 0, symnode02 = 1 }
    AutoStartList = { symnode01 }
)

    ASMDG ASM-asmdg (
        Sid = "+ASM"
        Owner = "oraprod"
        Home = "/orahome/Oracle"
        DiskGroups = { asmhighdg }
    )

    IP IP_oraprod (

        Device = eth0

        Address = "192.168.1.22"
        NetMask = "255.255.240.0"
    )

    NIC NIC_oraprod (

        Device = eth0

        NetworkHosts = { "192.168.1.1" }
    )

    Netlsnr LSNR_oraprod_lsnr (
```

```
Owner = "oraprod"
Home = "/orahome/Oracle"
TnsAdmin = "/orahome/Oracle/network/admin"
Listener = LISTENER_PROD
MonScript = "./bin/Netlsnr/LsnrTest.pl"
)

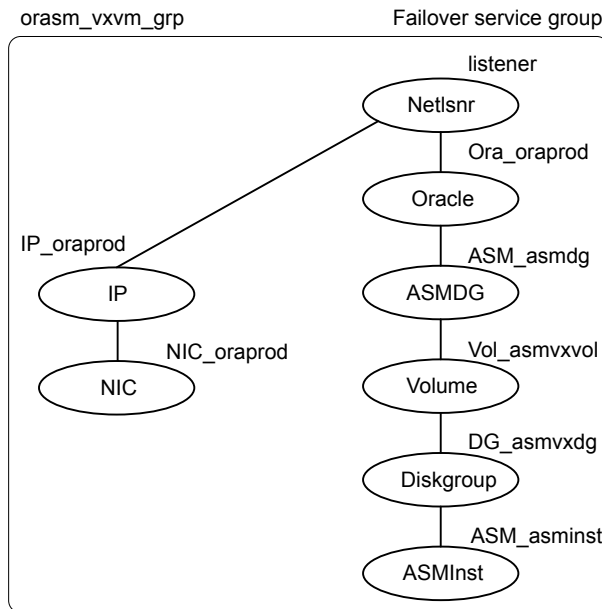
Oracle ORA_oraprod (
  Sid = PROD
  Owner = "oraprod"
  Home = "/orahome/Oracle"
  EnvFile = "/tmp/env.sh"
  MonScript = "./bin/Oracle/SqlTest.pl"
  User = thor
  Pword = hvlTptWvj
  Table = thor
  MonitorOption = 0
)

requires group asminstgrp online local firm
IP_oraprod requires NIC_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
ORA_oraprod requires ASM_asmdg
```

Sample configuration for ASM disks as VxVM volumes

[Figure B-8](#) describes a typical service group with VxVM disks chosen for Oracle ASM.

Figure B-8 Dependency graph for ASM on VxVM disks

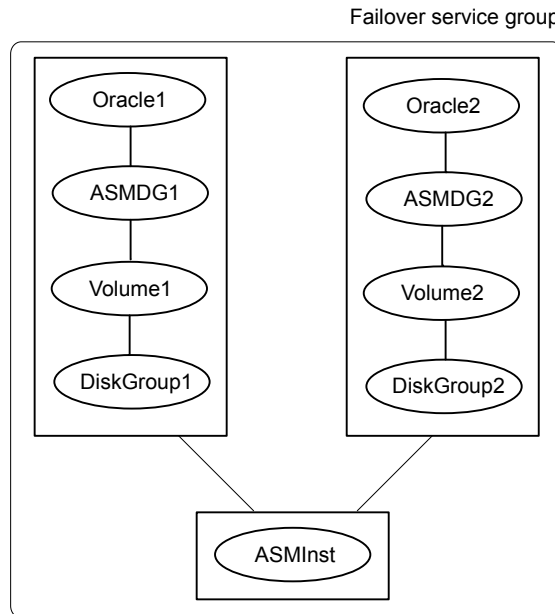


This configuration has a single failover service group. The Oracle ASM instance and ASM disk groups in the cluster are configured as resources of type ASMInst and ASMDG respectively.

The VxVM disk groups and volumes in the cluster are configured as resources of type DiskGroup and Volume respectively. The virtual IP address for the service group is configured using the IP and NIC resource types. The Oracle server can be started after each of these resources is brought online.

[Figure B-9](#) describes a typical service group with multiple Oracle instances sharing the ASMInst resource that is part of a single failover service group.

Figure B-9 Dependency graph for Oracle ASM with multiple Oracle instances on a node



Sample VCS configuration file for ASM disks as VxVM disks

Review the sample configuration for an Oracle instance that is ASM-managed.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"

cluster vcs (
)

system symnode01 (
)

system symnode02 (
)

group orasm_vxvm_grp (
    SystemList = { symnode01 = 0, symnode02 = 1 }
    AutoStartList = { symnode01 }
```

```
)

ASMDG ASM-asmdg (
  Sid = "+ASM"
  Owner = oraprod
  Home = "/orahome/Oracle"
  DiskGroups = { asmhhdg }

  DBAUser = "orauser"
  DBAPword = "DPSrFFeRiRGPiRF"

)

ASMinst ASM_asminst (
  Sid = "+ASM"
  Owner = oraprod
  Home = "/orahome/Oracle"

  DBAUser = "orauser"
  DBAPword = "DPSrFFeRiRGPiRF"

)

DiskGroup DG_asmvxdg (
  DiskGroup = asmvxdg

)

IP IP_oraprod (

  Device = eth0

  Address = "192.168.1.22"
  NetMask = "255.255.240.0"

)

NIC NIC_oraprod (

  Device = eth0

)

NetworkHosts = { "192.168.1.1" }

)

Netlsnr LSNR_oraprod_lsnr (
  Owner = oraprod
```

```
Home = "/orahome/Oracle"
TnsAdmin = "/orahome/Oracle/network/admin"
Listener = LISTENER_PROD
MonScript = "./bin/Netlsnr/LsnrTest.pl"
LsnrPwd = cqfOdoOolOo
)

Oracle ORA_oraprod (
  Sid = PROD
  Owner = oraprod
  Home = "/orahome/Oracle"
  EnvFile = "/tmp/env.sh"
  MonScript = "./bin/Oracle/SqlTest.pl"
  User = thor
  Pword = hv1TptWvj
  Table = thor
  MonitorOption = 0
)

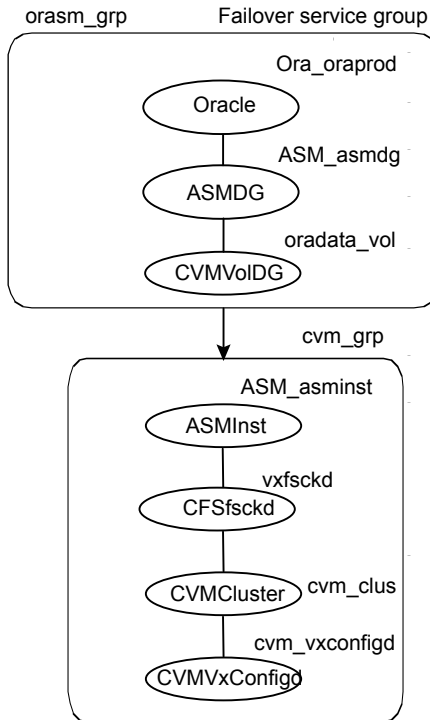
Volume Vol_asmvxvol (
  Volume = asmvxvol
  DiskGroup = asmvxdg
)

IP_oraprod requires NIC_oraprod
LSNR_oraprod_lsnr requires IP_oraprod
LSNR_oraprod_lsnr requires ORA_oraprod
ASM_asmdg requires VOL_asmvxvol
ORA_oraprod requires ASM_asmdg
VOL_asmvxvol requires DG_asmvxdg
DG_asmvxdg requires ASM_asminst
```

Sample configuration for ASM disks as CVM volumes

[Figure B-10](#) describes a typical service group with CVM volumes chosen for Oracle ASM.

Figure B-10 Dependency graph for ASM on CVM volumes

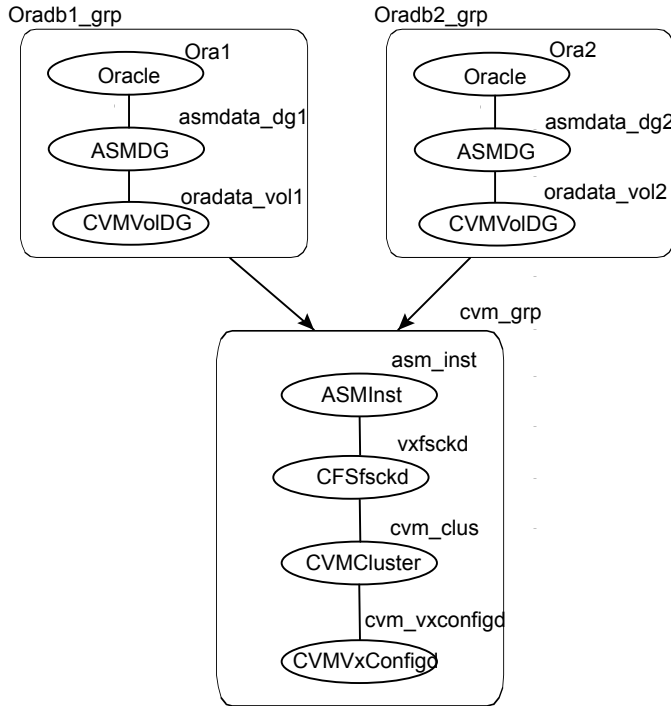


This configuration has two service groups. The Oracle, ASMDG, and CVMVoIDG resources are part of the parent failover service group oraasm_grp. The ASMIInst resource belongs to the CVM service group cvm_grp, which is a parallel service group. The service groups are linked with online local firm dependency.

After the CVM volume where the database resides comes online, the ASMDG agent mounts the ASM disk group that the database requires. The virtual IP address for the service group is configured using the IP and NIC resource types. The Oracle server can be started after each of these resources come online.

Figure B-11 describes a typical service group with multiple Oracle instances which share the ASMIInst resource that is part of a CVM parallel service group.

Figure B-11 Dependency graph for Oracle ASM with multiple Oracle instances on a node



If you have multiple Oracle instances, a failover service group is configured for each Oracle instance. The Oracle service groups share a single ASM instance that is configured as part of the CVM parallel service group. The Oracle service groups are made dependent on the CVM service group using an Online Local Firm dependency. However, each database must use exclusive ASM disk groups, so that the Cluster Server agent for Oracle can fail over the disk group.

Sample VCS configuration file for ASM that uses CVM volumes

Review the sample configuration for an Oracle instance that uses CVM volumes for ASM.

```

include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"
  
```

```
cluster vcscclus_asm (
  UserNames = { admin = abcdef }
  Administrator = { admin }
  UseFence = SCSI3
  HacliUserLevel = COMMANDROOT
)

system symnode01 (
)

system symnode02 (
)

group cvm (
  SystemList = { symnode01 = 0, symnode02 = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { symnode01, symnode02 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
  CVMClustName = vcscclus_asm
  CVMNodeId = { symnode01 = 0, symnode02 = 1 }
  CVMTransport = gab
  CVMTimeout = 200
)

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

ASMInst ASM_asminst (
  Sid = "+ASM"
  Owner = oracle
  Home = "/orahome/Oracle"
)
```

```
ASM_asminst requires vxfscd
vxfscd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```

```
group oraasm_grp (
  SystemList = { symnode01 = 0, symnode02 = 1 }
  AutoFailOver = 1
  AutoStartList = { symnode01, symnode02 }
)
```

```
CVMVolDg oradata_vol (
  CVMDiskGroup = ora_dg
  CVMVolume = { oradatavol }
  CVMActivation = sw
)
```

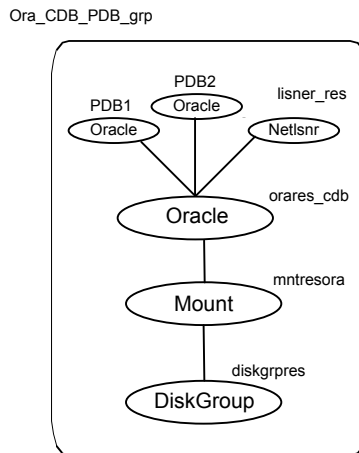
```
ASMDG ASM_asmdg (
  Sid = "+ASM"
  Owner = oracle
  Home = "/orahome/Oracle"
  DiskGroups = { ASM_DG }
)
```

```
Oracle Ora_oraprod (
  Owner = oracle
  Sid = vrts
  Home = "/orahome/Oracle"
)
```

```
requires group cvm_grp online local firm
Ora_oraprod requires ASM_asmdg
ASM_asmdg requires oradata_vol
```

Sample configuration of Oracle pluggable database (PDB) resource in main.cf

Figure B-12 Sample Oracle CDB-PDB configuration



Review the sample configuration with two PDB resources that are configured in the main.cf file.

```

group CDB (
    SystemList = { sys1 = 1, sys2 = 1 }
    AutoStartList = { sys2 }
)

DiskGroup diskgrpres (
    DiskGroup = cdb_dg
)

Mount mntresora (
    MountPoint = "/db01"
    BlockDevice = "/dev/vx/dsk/cdb_dg/oravol"
)
  
```

```
)

IP IP_oraprod (
    Device = eth0
    Address = "192.168.1.22"
    NetMask = "255.255.240.0"
)

NIC NIC_oraprod (
    Device = eth0
    NetworkHosts = {"192.168.1.1"}
)

Netlsnr lisner_res(
    Owner = oracle
    Home = "/u01/orahome"
    TnsAdmin = "/u01/orahome"
    Listener = LISTENER_PROD
    MonScript = "./bin/Netlsnr/LsnrTest.pl"
    LsnrPwd = cqfOdoOolOo
)

Oracle orares_cdb (
    Critical = 0
    Sid = cdb
    Owner = oracle
    Home = "/u01/oraHome"
    StartUpOpt = STARTUP
    DBName = cdb
    User = "c##vcsuser"
    Pword = hvnTkvKvmVktMvj
    Table = vcs
    LevelTwoMonitorFreq = 5
)

Oracle pdb1 (
    Critical = 0
    Sid = cdb
    Owner = oracle
    Home = "/u01/oraHome"
    StartUpOpt = STARTUP
    PDBName = pdb1
```

Sample configuration of Oracle pluggable database (PDB) resource in main.cf

```
    User = "c##vcsuser"
    Pword = hvnTkvKvmVktMvj
    Table = vcs
    LevelTwoMonitorFreq = 5
)

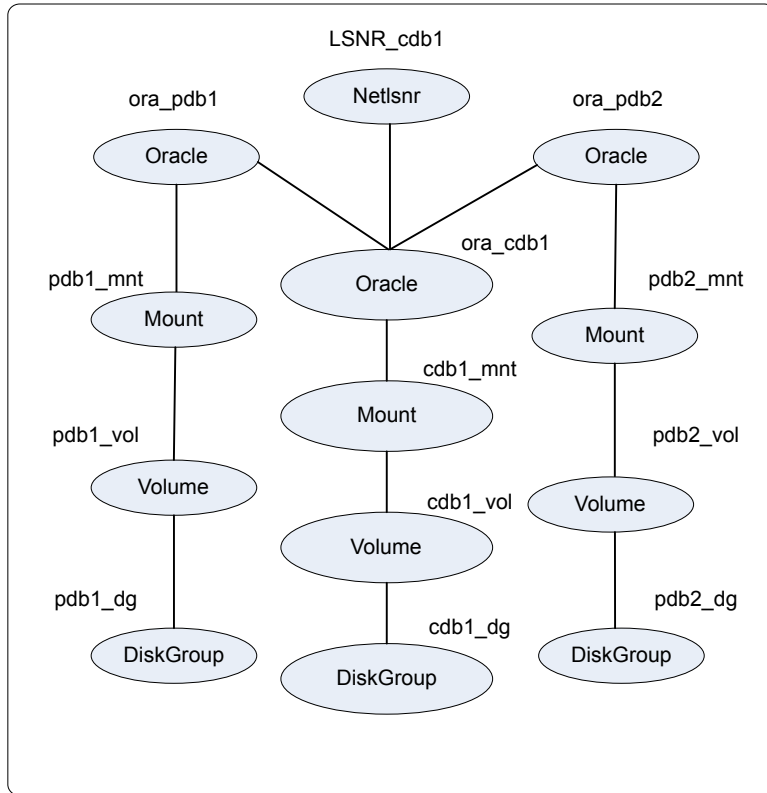
Oracle pdb2 (
    Critical = 1
    Sid = cdb
    Owner = oracle
    Home = "/u01/oraHome"
    PDBName = PDB2
    User = "c##vcsuser"
    Pword = hvnTkvKvmVktMvj
    Table = vcs
    LevelTwoMonitorFreq = 5
)

mntresora requires diskgrpres
orares_cdb requires mntresora
pdb1 requires orares_cdb
pdb2 requires orares_cdb
lisner_res requires orares_cdb
IP_oraprod requires NIC_oraprod
lisner_res requires IP_oraprod
```

Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf

Figure B-13 Sample migratable Oracle CDB-PDB configuration

ORA_CDB1_PROD_Group



```
include "types.cf"
include "OracleTypes.cf"

cluster vcs (
)

system SysA (
)

system SysB(
```

Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf

```
)

group ORA_CDB1_PROD_Group(
  SystemList = { SysA = 0, SysB = 1 }
)

    DiskGroup cdb1_dg (
        DiskGroup = cdb_dg1
    )

Volume cdb1_vol (
  Volume = cdbvol1
  DiskGroup = cdb_dg1
)

Mount cdb1_mnt (
  MountPoint = "/cdb1"
  BlockDevice = "/dev/vx/dsk/cdb_dg1/cdbvol1"
  FSType = vxfs
  FsckOpt = "-y"
)

DiskGroup pdb1_dg (
    DiskGroup = pdb_dg1
)

    Volume pdb1_vol (
        Volume = pdbvol1
        DiskGroup = pdb_dg1
    )

    Mount pdb1_mnt (
        MountPoint = "/pdb1"
        BlockDevice = "/dev/vx/dsk/pdb_dg1/pdbvol1"
        FSType = vxfs
        FsckOpt = "-y"
    )

DiskGroup pdb2_dg (
    DiskGroup = pdb_dg2
)

    Volume pdb2_vol (
```



```

        Volume = pdbvol2
        DiskGroup = pdb_dg2
    )

    Mount pdb2_mnt (
        MountPoint = "/pdb2"
        BlockDevice = "/dev/vx/dsk/pdb_dg2/pdbvol2"
        FSType = vxfs
        FsckOpt = "-y"
    )

IP IP_cdb1 (
    Device = hme0
    Address = "10.212.102.13"
    NetMask = "255.255.252.0"
)

NIC NIC_cdb1 (
    Device = hme0
    NetworkHosts = { "10.212.102.1" }
)

Netlsnr LSNR_cdb1 (
    Owner = oracle
    Home = "/u01/oraHome1"
)

Oracle ora_cdb1 (
    Sid = cdb1
    Owner = oracle
    Home = "/u01/oraHome1"
    StartUpOpt = STARTUP
    DBName = cdb1
)

Oracle ora_pdb1 (
    Sid = cdb1
    Owner = oracle
    Home = "/u01/oraHome1"
    PDBName = PDB1
)

Oracle ora_pdb2 (

```

Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf

```

Critical = 0
Sid = cdb1
Owner = oracle
Home = "/u01/oraHome1"
PDBName = PDB2
)

IP_cdb1 requires NIC_cdb1
LSNR_cdb1 requires IP_cdb1
LSNR_cdb1 requires ora_cdb1
cdb1_mnt requires cdb1_vol
cdb1_vol requires cdb1_dg
pdb1_mnt requires pdb1_vol
pdb1_vol requires pdb1_dg
pdb2_mnt requires pdb2_vol
pdb2_vol requires pdb2_dg
ora_cdb1 requires cdb1_mnt
ora_pdb1 requires ora_cdb1
ora_pdb1 requires pdb1_mnt
ora_pdb2 requires ora_cdb1
ora_pdb2 requires pdb2_mnt

group ORA_CDB2_PROD_Group(
  SystemList = { SysA = 0, SysB = 1 }
)

    DiskGroup cdb2_dg (
        DiskGroup = cdb_dg2
    )

Volume cdb2_vol (
  Volume = cdbvol2
  DiskGroup = cdb_dg2
)

Mount cdb2_mnt (
  MountPoint = "/cdb2"
  BlockDevice = "/dev/vx/dsk/cdb_dg2/cdbvol2"
  FSType = vxfs
  FsckOpt = "-y"
)

DiskGroup pdb3_dg (

```

Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf

```
        DiskGroup = pdb_dg3
    )

    Volume pdb3_vol (
        Volume = pdbvol3
        DiskGroup = pdb_dg3
    )

    Mount pdb3_mnt (
        MountPoint = "/pdb3"
        BlockDevice = "/dev/vx/dsk/pdb_dg3/pdbvol3"
        FSType = vxfs
        FsckOpt = "-y"
    )

    DiskGroup pdb4_dg (
        DiskGroup = pdb_dg4
    )

    Volume pdb4_vol (
        Volume = pdbvol4
        DiskGroup = pdb_dg4
    )

    Mount pdb4_mnt (
        MountPoint = "/pdb4"
        BlockDevice = "/dev/vx/dsk/pdb_dg4/pdbvol4"
        FSType = vxfs
        FsckOpt = "-y"
    )

    IP IP_cdb2 (
        Device = hme0
        Address = "10.212.102.14"
        NetMask = "255.255.252.0"
    )

    NIC NIC_cdb2 (
        Device = hme0
        NetworkHosts = { "10.212.102.1" }
    )

    Netlsnr LSNR_cdb2 (
```

Sample configuration of migratable Oracle pluggable database (PDB) resource in main.cf

```
        Owner = oracle
        Home = "/u01/oraHome2"
    )

Oracle ora_cdb2 (
    Sid = cdb2
    Owner = oracle
    Home = "/u01/oraHome2"
    StartUpOpt = STARTUP
    DBName = cdb2
)

Oracle ora_pdb3 (
    Sid = cdb2
    Owner = oracle
    Home = "/u01/oraHome2"
    PDBName = PDB3
)

Oracle ora_pdb4 (
    Sid = cdb2
    Owner = oracle
    Home = "/u01/oraHome2"
    PDBName = PDB4
)

IP_cdb2 requires NIC_cdb2
LSNR_cdb2 requires IP_cdb2
LSNR_cdb2 requires ora_cdb2
cdb2_mnt requires cdb2_vol
cdb2_vol requires cdb2_dg
pdb3_mnt requires pdb3_vol
pdb3_vol requires pdb3_dg
pdb4_mnt requires pdb4_vol
pdb4_vol requires pdb4_dg
ora_cdb2 requires cdb2_mnt
ora_pdb3 requires ora_cdb2
ora_pdb3 requires pdb3_mnt
ora_pdb4 requires ora_cdb2
ora_pdb4 requires pdb4_mnt
```

Best practices

This appendix includes the following topics:

- [Best practices for multiple Oracle instance configurations in a VCS environment](#)

Best practices for multiple Oracle instance configurations in a VCS environment

Review some of the best practices for using multiple Oracle instances in a VCS environment:

- For each SID to be configured, create Linux accounts with DBA privileges.
- Make sure that each Oracle instance has a separate disk group and is configured as a separate service group.
- Define the system parameters such that the allocation of semaphore and shared memory is appropriate on all systems.
- Use a dedicated set of binaries for each Oracle instance, even if each instance uses the same Oracle version.
- If your configuration uses the same Oracle version for all instances, install a version on the root disk or preferably on a secondary disk. Locate the pfiles in the default location and define several listener processes to ensure clean failover.
- If your configuration has different versions of Oracle, create a separate \$ORACLE_HOME for each Oracle version.
- Follow the Optimal Flexible Architecture (OFA) standard (/uwx/<SID>). In cluster configurations, you could adapt the standard to make it more application-specific. For example, /app/uwx/<SID>.
- Listeners accompanying different versions of Oracle may not be backward-compatible. So, if you want to create a single listener.ora file, you

must verify that the listener supports the other versions of Oracle in the cluster. You must also create a separate Envfile for each version of Oracle.

- Make sure that each listener listens to a different virtual address. Also, assign different names to listeners and make sure that they do not listen to the same port.
- The pfiles must be coordinated between systems. For the same instance of a database, ensure that the pfiles referenced are identical across the nodes.

Using the SPFILE in a VCS cluster for Oracle

This appendix includes the following topics:

- [About the Oracle initialization parameter files](#)
- [Starting an Oracle instance](#)
- [Using the SPFILE in a VCS cluster](#)

About the Oracle initialization parameter files

Oracle9i introduced the SPFILE, which is a binary file stored on the database server. With this feature, changes to the instance parameters can be set to be persistent across all startup and shutdown procedures.

Starting an Oracle instance

You can start an Oracle instance in the following ways:

- Using the default SPFILE `spfileSID.ora`
- Using the default init.ora file `initSID.ora`
- By specifying an initialization file `init.ora`
- By specifying an SPFILE in the initialization file `init.ora`

When you run the `startup` command without a PFILE clause, Oracle reads the initialization parameters from the SPFILE. On Linux platforms, the default location for the SPFILE or PFILE is `$ORACLE_HOME/dbs`.

Oracle locates the initialization parameter file by examining file names in the following order:

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

Using the SPFILE in a VCS cluster

Oracle versions earlier than Oracle9i used an initialization file `initSID.ora`, a text file, to start database instances. Changes that were applied to instance parameters during a database session were not saved to the file. You had to manually apply them to the initialization file.

When using the Cluster Server agent for Oracle, you can start a database instance by specifying a PFILE. If you do not specify the PFILE, the database instance starts up using the default SPFILE.

The agent attribute `Pfile` must specify the location of the PFILE. If your configuration uses the SPFILE, the contents of the PFILE must specify the location of the SPFILE, which must be created from the PFILE.

Note: If you want the SPFILE's session parameter changes be persistent across an instance failover, then recommends you to save the SPFILE on shared storage.

To create the SPFILE from a PFILE

- ◆ The SPFILE must be created from the PFILE. You must have the `sysdba` or the `sysoper` system privileges to create an SPFILE.

You can run the following command to create the SPFILE:

```
CREATE SPFILE [= spfile_name] FROM PFILE [= pfile_name ];
```

If you do not specify the complete path for the SPFILE, this command creates an SPFILE at the default location (`$ORACLE_HOME/dbs` on Linux).

To specify the SPFILE location in the PFILE

- ◆ To specify the location of the SPFILE in a PFILE, create a PFILE and specify the following entry in the PFILE:

```
SPFILE = spfile_location
```

The variable *spfile_location* represents the complete path of the SPFILE. For example:

```
SPFILE = /database/startup/spfileora1.ora
```

In this case, to start the database use the following command:

```
startup pfile=location_of_pfile
```

OHASD in a single instance database environment

This appendix includes the following topics:

- [About OHASD in a single instance database environment](#)
- [Configuring the Application agent to make OHASD highly available](#)

About OHASD in a single instance database environment

The Oracle Grid Infrastructure software provides system support for an Oracle database including Oracle Automatic Restart and Oracle ASM features.

On nodes that have Oracle Grid Infrastructure installed in a single instance database environment, ensure that the Oracle High Availability Services daemon (OHASD) process comes online before the Oracle resources.

Administration of Oracle resources on these nodes is dependent on the OHASD process. So, the process needs to be online first to ensure that the resources come online.

Configuring the Application agent to make OHASD highly available

You can configure the Application agent to ensure the OHASD process is made highly available in the VCS cluster.

To configure the Application agent for the Ohasd process to ensure the service is online

1. Configure the Application agent

```
Application ohasd_res (
  StartProgram = "<GRID_HOME>/bin/crsctl start has"
  StopProgram = "<GRID_HOME>/bin/crsctl stop has"
  PidFiles @node1 = {
    "<GRID_HOME>/ohasd/init/node1.pid" }
  PidFiles @node2 = {
    "<GRID_HOME>/ohasd/init/node2.pid" }
)
```

2. Set up the resource dependency on the Ohasd daemon if the resource is configured:

- For databases that are running in oracle restart mode
- For the ASMInst resource

3. In the following dependency tree, ohasd resource is configured in a failover service group or parallel service group.

Run the # `$GRID_HOME/bin/srvctl disable has` command to disable the automatic startup of ohasd on all the nodes in the VCS cluster.

Figure E-1 shows OHASD resource configured in Parallel service group with the ASMInst resource in the Failover service group




Figure E-2 shows OHASD resource configured in Failover service group with the ASMInst resource.

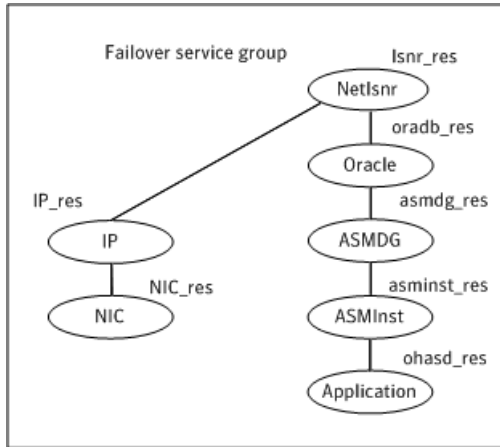



Figure E-3 shows the dependency tree of databases that are running in oracle restart mode.



Index

Symbols

`$ORACLE_HOME`
limitation for 37
location 37, 41
on local disk 41
on shared disk 41

A

agent operations
 Netlsnr agent 28
 Oracle agent 13

ASM
 configuring on a remote node 54
 starting on a remote node 54

ASMDG agent
 attribute definitions 161
 resource type 161

ASMDG agent attributes
 AgentDirectory 161
 DBAPword 161
 DBAUser 161
 DiskGroups 161
 Encoding 161
 EnvFile 161
 Home 161
 Owner 161
 Sid 161

ASMIInst agent
 attribute definitions 157
 resource type 157

ASMIInst agent attributes
 AgentDirectory 157
 DBAPword 157
 DBAUser 157
 Encoding 157
 EnvFile 157
 Home 157
 Owner 157
 Pfile 157
 ShutDownOpt 157
 Sid 157

ASMIInst agent attributes *(continued)*
 StartUpOpt 157

ASMIInst monitoring
 health check 31
 process 31

attribute definitions
 ASMDG agent 161
 ASMIInst 157
 Netlsnr agent 153
 Oracle agent 143

attributes
 editing 80
 modifying 80

B

basic monitoring 20
 health check 20
 process 20

C

cluster manager 78
configuration
 typical setup 33
configuration files 164
cssd daemon 57
 disabling 57

D

database tablespaces 37
detail monitoring 20
 about 84
 disabling for Netlsnr 88
 disabling for Oracle 87
 enabling for Netlsnr 89
 enabling for Oracle 87
 setting up 84

E

editing attributes 80
encrypting passwords 82

- enterprise agent
 - about 9
 - configuring using cluster manager 78
 - configuring using command-line 80
 - disabling 65
 - error messages 128
 - installing 62
 - removing 66
 - troubleshooting 127
 - typical setup 33
 - upgrading 64
- error messages 128

F

- failover 10

H

- health check APIs 20, 31
- health check monitoring 20, 31
- hot backup 37, 42

I

- intelligent resource monitoring
 - disabling manually 89
 - enabling manually 89

M

- message log 128
- migrating
 - PDB 122
- monitoring
 - basic 20
 - detail 20
- monitoring options 10

N

- Netlsnr agent
 - attribute definitions 153
 - operations 28
 - resource type 153
- Netlsnr agent attributes
 - AgentDebug 153
 - AgentDirectory 153
 - Encoding 153
 - EnvFile 153
 - Home 153
 - IntentionalOffline 153

- Netlsnr agent attributes *(continued)*

- Listener 153
- LsnrPwd 153
- MonScript 153
- Owner 153
- TnsAdmin 153

- NLS information
 - defining 37

O

- operations
 - Netlsnr agent 28
 - Oracle agent 13
- Oracle
 - database tablespaces 37
 - error handling 22
 - installing 36
 - multiple instances 37
 - shutdown options 15, 31
 - startup options 15, 30
 - supported configurations 68
- Oracle agent
 - attribute definitions 143
 - operations 13
 - resource type 142
- Oracle agent attributes
 - AgentDebug 143
 - AgentDirectory 143
 - AutoEndBkup 143
 - DBAPword 143
 - DBAUser 143
 - Encoding 143
 - EnvFile 143
 - Home 143
 - IMF 143
 - IntentionalOffline 143
 - MonitorOption 143
 - MonScript 143
 - Owner 143
 - Pfile 143
 - Pword 143
 - ShutDownOpt 143
 - Sid 143
 - StartUpOpt 143
 - Table 143
 - User 143
- Oracle configuration
 - multiple instances (multiple listeners) 71
 - multiple instances (single listener) 168

Oracle configuration *(continued)*
 single instance 165
Oracle database
 hot backup 37, 42

P

passwords
 encrypting 82
process monitoring 20, 31

R

removing agent 120
resource attributes
 editing 80
resource type
 ASMDG 161
 ASMInst 157
 Netsnr 153
 Oracle 142

S

sample configuration files 164
sample Oracle configuration 33
service group
 bringing online 120
 switching 120
 taking offline 120
shutdown options 15
SPFILE
 creating 208
 location of 208
 using in VCS 207
startup options 15
switching service group 120

T

troubleshooting 127

V

vcsencrypt utility 82
virtual environment
 before configuring monitoring 97
 configuring Oracle for monitoring 100
 infrastructure service groups 107
 launching the wizard 99
 resource dependency 106
 sample configurations 110

virtual environment *(continued)*
 troubleshooting 109
 wizard limitations 109
virtual fire drill 25