

Oracle® Solaris Cluster Data Service for Oracle Guide

Copyright © 2000, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Preface	7
1 Installing and Configuring HA for Oracle	11
Overview of the Installation and Configuration Process for HA for Oracle	11
Planning the HA for Oracle Installation and Configuration	13
Configuration Requirements	13
Configuration Planning Questions	14
Preparing the Nodes and Disks	15
▼ How to Prepare the Nodes	15
▼ How to Configure the Oracle Database Access Using Solaris Volume Manager	17
▼ How to Configure the Oracle Database Access Using Oracle ASM	17
▼ How to Configure an Oracle Grid Infrastructure for Clusters SCAN Listener	18
Installing the Oracle ASM Software	19
Verifying the Oracle ASM Software Installation	20
Installing the Oracle Software	20
▼ How to Install the Oracle Software	20
▼ How to Set the Oracle Kernel Parameters	21
Verifying the Oracle Installation and Configuration	21
▼ How to Verify the Oracle Installation	21
Creating an Oracle Database	22
▼ How to Create a Primary Oracle Database	22
Setting Up Oracle Database Permissions	23
▼ How to Set Up Oracle Database Permissions	23
Installing the HA for Oracle Package	27
▼ How to Install the HA for Oracle Package	27
Registering and Configuring HA for Oracle	28
Tools for Registering and Configuring HA for Oracle	28
Setting HA for Oracle Extension Properties	28

▼ How to Register and Configure HA for Oracle (clsetup)	28
▼ How to Register and Configure HA for Oracle Without Oracle ASM (CLI)	33
▼ How to Create an Oracle Grid Infrastructure Resource With Clustered Oracle ASM Disk Groups and a Third-Party Volume Manager (CLI)	40
▼ How to Register and Configure HA for Oracle With Clustered Oracle ASM Instance (CLI)	43
Where to Go From Here	47
Verifying the HA for Oracle Installation	47
▼ How to Verify the HA for Oracle Installation	47
Oracle Clients	48
Location of HA for Oracle Log Files	48
Tuning the HA for Oracle Fault Monitors	49
Operation of the Oracle Server Fault Monitor	50
Operation of the Oracle Listener Fault Monitor	52
Obtaining Core Files for Troubleshooting DBMS Timeouts	53
Customizing the HA for Oracle Server Fault Monitor	53
Defining Custom Behavior for Errors	54
Propagating a Custom Action File to All Nodes in a Cluster	62
Specifying the Custom Action File That a Server Fault Monitor Should Use	62
Changing the Role of an Oracle Data Guard Instance	63
▼ How to Change the Role of an Oracle Data Guard Instance	63
A HA for Oracle Extension Properties	65
SUNW.oracle_server Extension Properties	65
SUNW.oracle_listener Extension Properties	69
B Preset Actions for DBMS Errors and Logged Alerts	71
C Sample Configurations for Oracle ASM with HA for Oracle	79
Clustered Oracle ASM with Clustered Disk Groups	79
Index	83

Tables

TABLE 1-1	Task Map: Installing and Configuring HA for Oracle	12
TABLE 1-2	Task Map: Installing and Configuring HA for Oracle with Clustered Oracle ASM	12
TABLE 1-3	Resource Types for HA for Oracle Fault Monitors	49
TABLE B-1	Preset Actions for DBMS Errors	71
TABLE B-2	Preset Actions for Logged Alerts	78

Preface

Oracle Solaris Cluster Data Service for Oracle Guide explains how to install and configure Oracle Solaris Cluster data services.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures. In this document, “x86” refers to the larger family of x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

This document is intended for system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or presales guide. Before reading this document, you should have already determined your system requirements and purchased the appropriate equipment and software.

The instructions in this book assume knowledge of the Oracle Solaris Operating System and expertise with the volume-manager software that is used with Oracle Solaris Cluster software.

Bash is the default shell for Oracle Solaris 11. Machine names shown with the Bash shell prompt are displayed for clarity.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Oracle Solaris Cluster data services. The document does *not* contain comprehensive information about basic UNIX commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

- Online documentation for the Oracle Solaris Operating System
- Oracle Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Description	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

Topic	Documentation
Hardware installation and administration	<i>Oracle Solaris Cluster 4.0 Hardware Administration Manual</i> Individual hardware administration guides
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
Data service installation and administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> and individual data service guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i> <i>Oracle Solaris Cluster Quick Reference</i>
Software upgrade	<i>Oracle Solaris Cluster Upgrade Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i> <i>Oracle Solaris Cluster Geographic Edition Reference Manual</i> <i>Oracle Solaris Cluster Quorum Server Reference Manual</i>

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Getting Help

If you have problems installing or using Oracle Solaris Cluster, contact your service provider and provide the following information.

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 11)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 4.0)

Use the following commands to gather information about your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>pkg list</code>	Reports which packages are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information for each node

Also have available the contents of the `/var/adm/messages` file.

Installing and Configuring HA for Oracle

This chapter explains how to install and configure Oracle Solaris Cluster HA for Oracle (HA for Oracle).

This chapter contains the following sections.

- “Overview of the Installation and Configuration Process for HA for Oracle” on page 11
- “Planning the HA for Oracle Installation and Configuration” on page 13
- “Preparing the Nodes and Disks” on page 15
- “Installing the Oracle Software” on page 20
- “Verifying the Oracle Installation and Configuration” on page 21
- “Creating an Oracle Database” on page 22
- “Setting Up Oracle Database Permissions” on page 23
- “Installing the HA for Oracle Package” on page 27
- “Registering and Configuring HA for Oracle” on page 28
- “Verifying the HA for Oracle Installation” on page 47
- “Tuning the HA for Oracle Fault Monitors” on page 49
- “Customizing the HA for Oracle Server Fault Monitor” on page 53
- “Changing the Role of an Oracle Data Guard Instance” on page 63

Overview of the Installation and Configuration Process for HA for Oracle

The following task maps summarize the tasks for installing and configuring HA for Oracle. The tables also provide cross-references to detailed instructions for performing the tasks.

- [Task Map: Installing and Configuring HA for Oracle](#)
- [Task Map: Installing and Configuring HA for Oracle with Clustered Oracle ASM](#)

Perform these tasks in the order that they are listed. If you are using HA for Oracle with Oracle Data Guard, perform these tasks on each cluster where your Oracle database instances are running.

Note – Single instance Oracle ASM is not supported with Oracle 11g release 2.

TABLE 1-1 Task Map: Installing and Configuring HA for Oracle

Task	Cross-Reference
Plan the HA for Oracle installation and configuration	“Planning the HA for Oracle Installation and Configuration” on page 13
Prepare the nodes and disks	“Preparing the Nodes and Disks” on page 15
Install the Oracle software	“How to Install the Oracle Software” on page 20
Verify the Oracle installation	“How to Verify the Oracle Installation” on page 21
Create an Oracle database	“How to Create a Primary Oracle Database” on page 22
Set up Oracle database permissions	“How to Set Up Oracle Database Permissions” on page 23
Install the HA for Oracle packages	“Installing the HA for Oracle Package” on page 27
Register and configure HA for Oracle	“How to Register and Configure HA for Oracle Without Oracle ASM (CLI)” on page 33
Verify the HA for Oracle installation	“Verifying the HA for Oracle Installation” on page 47
Tune the HA for Oracle fault monitor	“Tuning the HA for Oracle Fault Monitors” on page 49
(Optional) Customize the HA for Oracle server fault monitor	“Customizing the HA for Oracle Server Fault Monitor” on page 53
(Optional) Change the role of an Oracle Data Guard instance	“Changing the Role of an Oracle Data Guard Instance” on page 63

TABLE 1-2 Task Map: Installing and Configuring HA for Oracle with Clustered Oracle ASM

Task	Cross-Reference
Plan the HA for Oracle installation and configuration	“Planning the HA for Oracle Installation and Configuration” on page 13
Prepare the nodes and disks	“Preparing the Nodes and Disks” on page 15
Install the Oracle ASM software	“Installing the Oracle ASM Software” on page 19
Install the Oracle software	“How to Install the Oracle Software” on page 20
Verify the Oracle installation	“How to Verify the Oracle Installation” on page 21
Create an Oracle database	“How to Create a Primary Oracle Database” on page 22
Set up Oracle database permissions	“How to Set Up Oracle Database Permissions” on page 23

TABLE 1-2 Task Map: Installing and Configuring HA for Oracle with Clustered Oracle ASM
(Continued)

Task	Cross-Reference
Install the HA for Oracle packages	“Installing the HA for Oracle Package” on page 27
Register and configure HA for Oracle	“How to Register and Configure HA for Oracle With Clustered Oracle ASM Instance (CLI)” on page 43
Verify the HA for Oracle installation	“Verifying the HA for Oracle Installation” on page 47
Tune the HA for Oracle fault monitor	“Tuning the HA for Oracle Fault Monitors” on page 49
(Optional) Customize the HA for Oracle server fault monitor	“Customizing the HA for Oracle Server Fault Monitor” on page 53
(Optional) Change the role of an Oracle Data Guard instance	“Changing the Role of an Oracle Data Guard Instance” on page 63

Planning the HA for Oracle Installation and Configuration

This section contains the information that you need to plan your HA for Oracle installation and configuration.

Configuration Requirements



Caution – Your data service configuration might not be supported if you do not adhere to these requirements.

Use the requirements in this section to plan the installation and configuration of HA for Oracle. These requirements apply to HA for Oracle only. You can also configure HA for Oracle in a zone cluster.

For requirements that apply to all data services, see “Configuration Guidelines for Oracle Solaris Cluster Data Services” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

- Oracle Grid Infrastructure software requirements** - If you will use Oracle Grid Infrastructure (Oracle ASM and Oracle Clusterware), ensure that the cluster meets Oracle Grid Infrastructure software requirements. See “Identifying Software Requirements” in *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2) for Oracle Solaris* (http://docs.oracle.com/cd/E11882_01/install.112/e24616/presolar.htm#CHDEFJCB).

- **Oracle application files** – These files include Oracle binaries, configuration files, and parameter files. You can install these files either on the local file system, the highly available local file system, or on the cluster file system.
See “[Configuration Guidelines for Oracle Solaris Cluster Data Services](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for the advantages and disadvantages of placing the Oracle binaries on the local file system, highly available local file system, and the cluster file system.
- **Database related files**– These files include the control file, redo logs, and data files. You must install these files on either raw devices or as regular files on the highly available local or cluster file system. Raw devices from Oracle Solaris Cluster device groups are not supported in non-global zones.
- **Oracle ASM configuration** – Oracle Automatic Storage Management (Oracle ASM) is a storage option that provides the services of a file system, logical volume manager, and software redundant array of independent disks (RAID) in a platform independent manner. For more information on Oracle ASM, see the Oracle documentation corresponding to the Oracle database version you are using.

Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of HA for Oracle.

- What resource groups will you use for network addresses and application resources and the dependencies between them?
- What is the logical hostname for clients that will access the data service?
- Where will the system configuration files reside?
See “[Configuration Guidelines for Oracle Solaris Cluster Data Services](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for the advantages and disadvantages of placing the Oracle binaries on the local file system rather than the cluster file system.
- Does your database setup require standby instances?
If you use the `clsetup` utility to register and configure HA for Oracle, some of these questions are answered automatically by the utility.
For information about standby databases, see your Oracle documentation.
- Are you planning to use Oracle ASM storage?
For information about standby databases and Oracle ASM storage, see your Oracle documentation.

Preparing the Nodes and Disks

This section contains the procedures that you need to prepare the nodes and disks.

- “How to Prepare the Nodes” on page 15
- “How to Configure the Oracle Database Access Using Solaris Volume Manager” on page 17
- “How to Configure the Oracle Database Access Using Oracle ASM” on page 17
- “How to Configure an Oracle Grid Infrastructure for Clusters SCAN Listener” on page 18

▼ How to Prepare the Nodes

Use this procedure to prepare for the installation and configuration of Oracle software.



Caution – Perform all of the steps in this section on all the nodes. If you do not perform all of the steps on all of the nodes, the Oracle installation is incomplete. An incomplete Oracle installation causes HA for Oracle to fail during startup.

Note – Consult the Oracle documentation before you perform this procedure.

The following steps prepare your nodes and install the Oracle software.

- 1 **Become superuser on all of the cluster members.**
- 2 **Configure the cluster file system for HA for Oracle.**



Caution – Raw devices from Oracle Solaris Cluster device groups are not supported in non-global zones.

- If raw devices contain the databases, configure the global devices for raw device access. See the *Oracle Solaris Cluster Software Installation Guide* for information about how to configure global devices.
- If you use the Solaris Volume Manager software, configure the Oracle software to use UNIX file system (UFS) logging on mirrored metadevices or raw-mirrored metadevices. See the Solaris Volume Manager documentation for more information about how to configure raw-mirrored metadevices.
- If you use Oracle Solaris ZFS for Oracle files, configure a highly available local ZFS file system. For more information, see “How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS File System Highly Available” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

3 Prepare the \$ORACLE_HOME directory on a local or multihost disk.

Note – If you install the Oracle binaries on a local disk, use a separate disk if possible. Installing the Oracle binaries on a separate disk prevents the binaries from overwrites during operating environment reinstallation. When you use Oracle ASM, create two \$ORACLE_HOME directories, one \$ORACLE_HOME directory for the Oracle database and the additional \$ORACLE_HOME directory for Oracle ASM.

4 On each node, create an entry for the database administrator (DBA) group in the /etc/group file, and add potential users to the group.

You typically name the DBA group dba. Verify that the oracle users are members of the dba group, and add entries as necessary for other DBA users. Ensure that the group IDs are the same on all of the nodes that run HA for Oracle, as the following example illustrates.

```
dba:*:520:root,oracle
```

If you use Oracle ASM, you should add an additional entry for the DBA group in the /etc/group file for Oracle ASM.

```
dba:*:520:root,oracle,oraasm
```

5 On each node, create an entry for the Oracle user ID (oracle).

You typically name the Oracle user ID oracle. The following command updates the /etc/passwd and /etc/shadow files with an entry for the Oracle user ID.

```
# useradd -u 120 -g dba -d /Oracle_home oracle
```

Ensure that the oracle user entry is the same on all the nodes that run HA for Oracle.

If you use Oracle ASM, you should create an additional Oracle user ID for Oracle ASM.

```
# useradd -u 121 -g dba -d /asm-home oraasm
```

6 Configure kernel parameters in the global cluster and, if used, the zone cluster that runs HA for Oracle.

You might need to reboot the cluster to initiate certain parameter changes. For information about tuning Oracle Solaris kernel parameters, see [“Tuning the Oracle Solaris Kernel” in Oracle Solaris Tunable Parameters Reference Manual](#).

For information about shared memory requirements, see [“Configuring Kernel Parameters in Oracle Solaris 10” in Oracle Database Installation Guide 11g Release 2 \(11.2\) for Oracle Solaris](#). This information is valid for Oracle Solaris Cluster 4.0 software on the Oracle Solaris 11 OS.

7 If running HA for Oracle in a zone cluster, configure the limitpriv property.

The limitpriv property is required.

```
#clzonecluster configure zcname
```

```
clzonecluster:zcname> set limitpriv="default,proc_priocntl,proc_clock_highres,sys_time"
```

```
clzonecluster:zcname> commit
clzonecluster:zcname> exit
```

▼ How to Configure the Oracle Database Access Using Solaris Volume Manager

Use this procedure to configure the Oracle database using Solaris Volume Manager.

Note – You can run this procedure only in the global zone.

1 Configure the disk devices for the Solaris Volume Manager software to use.

See *Oracle Solaris Cluster Software Installation Guide* for information about how to configure the Solaris Volume Manager software.

2 If you use raw devices to contain the databases, run the following commands to change each raw-mirrored metadevice's owner, group, and mode.

If you do not use raw devices, do not perform this step.

a. If you create raw devices, run the following commands for each device on each node that can master the Oracle resource group.

```
# chown oracle /dev/md/metaset/rdisk/dn
# chgrp dba /dev/md/metaset/rdisk/dn
# chmod 600 /dev/md/metaset/rdisk/dn
```

metaset Specifies the name of the disk set

/rdisk/dn Specifies the name of the raw disk device within the *metaset* disk set

b. Verify that the changes are effective.

```
# ls -lL /dev/md/metaset/rdisk/dn
```

▼ How to Configure the Oracle Database Access Using Oracle ASM

Use this procedure to configure the Oracle database access using Oracle ASM. You can use Oracle ASM on Solaris Volume Manager.

1 Configure the disk devices for the Oracle ASM software to use.

See “Using Oracle ASM” in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide* for information about how to configure Oracle ASM.

- 2 Set up the permissions for the Oracle ASM disks to be used by Oracle ASM disk groups.
 - a. Run the following commands for each device on each node that can master the Oracle resource group.

```
# chown oraasm:dba /dev/did/rdisk/dn
# chmod 660 /dev/did/rdisk/dn
```

- b. Verify that the changes are effective.

```
# ls -lhL /dev/did/rdisk/dn
```

▼ How to Configure an Oracle Grid Infrastructure for Clusters SCAN Listener

When the HA for Oracle data service is deployed using Grid Infrastructure for Clusters, the preferred listener for a single-instance database is an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener. To leverage a SCAN in the HA for Oracle configuration, you must modify the single-instance database `remote_listener` parameter to reflect the SCAN name and port number.

If you configure a SCAN listener, you can ignore subsequent instructions in this manual about creating an Oracle listener resource with the `SUNW.oracle_listener` resource type. An Oracle listener resource is unnecessary when the SCAN listener is configured.

For more information about the SCAN, see *Oracle Grid Infrastructure Installation Guide 11g Release 2 (11.2) for Solaris Operating System* (http://docs.oracle.com/cd/E11882_01/install.112/e24346/toc.htm).

The following example procedure shows the commands and sample output to configure the SCAN name and port. Perform this procedure on each node of the cluster.

- 1 Become superuser on the node.

- 2 Configure a SCAN listener.

```
# srvctl config scan
SCAN name: scanname, Network: 1/ipaddress/netmask/adaptername
SCAN VIP name: scanvip, IP: /ipalias/vipaddress
```

```
# srvctl config scan_listener
SCAN Listener LISTENER_SCAN1 exists. Port: TCP:portnumber
#
```

```
SQL> show parameters listener
```

NAME	TYPE	VALUE
listener_networks	string	

```

local_listener      string      (DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=nodename)
remote_listener     string
                    (PORT=portnumber))))

```

```
SQL>
```

```
SQL> alter system set remote_listener = 'scanname:portnumber' scope=both;
```

```
System altered.
```

```
SQL> show parameters listener
```

NAME	TYPE	VALUE
listener_networks	string	
local_listener	string	(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=nodename) (PORT=portnumber))))
remote_listener	string	scanname:portnumber

```
SQL>
```

Installing the Oracle ASM Software

Oracle ASM installation consists of installing and creating an Oracle ASM instance and configuring the required Oracle ASM disk groups. An Oracle ASM disk group is a collection of disk devices to store data files that Oracle ASM instances manage as a unit. Oracle ASM instances mount disk groups to make Oracle ASM files available to database instances.

When you run Oracle Universal Installer, you have the option to configure and install Oracle ASM and create disk groups. For detailed information on installing and configuring Oracle ASM, see the Oracle documentation corresponding to the Oracle database version you are using.

Note – Before starting the Oracle Universal Installer, ensure that the Oracle Solaris `library/motif` package is installed.

For zone clusters, also ensure that the `group/system/solaris-large-server` package is installed.

Note – Set up and configure the file system, where the Oracle ASM `ORACLE_HOME` resides, before performing the Oracle ASM installation. Use only a local file system for the Oracle ASM home.

Verifying the Oracle ASM Software Installation

After installing Oracle ASM software, you can verify that Oracle ASM software is installed and Oracle ASM disk group is mounted on a cluster node by performing the following steps.

```
# sqlplus "/ as sysdba"
sql> select * from v$sga;
sql> select name,state from v$asm_diskgroup;
sql> exit;
#
```

Installing the Oracle Software

This section contains the procedure that you need to install Oracle software.

▼ How to Install the Oracle Software

- 1 **Become superuser on a cluster member.**
- 2 **If you plan to install the Oracle software on a cluster file system, start the Oracle Solaris Cluster software and become the owner of the device group.**

If you plan to install the Oracle software at another location, omit this step.

For more information about installation locations, see [“Preparing the Nodes and Disks” on page 15](#).

- 3 **Install the Oracle software.**

Before you start the Oracle installation, ensure that the system resources required for Oracle have been configured. Log in as `oracle` to ensure ownership of the entire directory before you perform this step. See the appropriate Oracle installation and configuration guides for instructions about how to install Oracle software.

You could use Oracle Solaris Resource Management (SRM) to ensure that the kernel parameters are set to at least the minimum values that Oracle requires. For more information about setting the Oracle kernel parameters, see [“How to Set the Oracle Kernel Parameters” on page 21](#). After the system resources have been configured for Oracle you can start the installation itself.

▼ How to Set the Oracle Kernel Parameters

The default project is modified to contain the resources required for Oracle as the RGM uses the default project for running the data service. If you want to use a specific SRM project for running Oracle, you must create that project and configure the system resources in that project using the same procedure. Specify the project name instead of default. When you configure the resource group or resource for the Oracle server, specify that project name in the corresponding property of the resource group or resource.

1 Display the settings for the default project.

```
phys-X# prctl -i project default
```

2 If no kernel parameters are set, or if any kernel parameters are not set to the minimum required value for Oracle as shown in the following table, set the parameter.

```
phys-X# projmod -s -K "parameter=(priv,value,deny)" default
```

Oracle Kernel Parameter	Minimum Required Value
process.max-sem-nsems	256
project.max-sem-ids	100
project.max-shm-ids	100
project.max-shm-memory	4294967295

See *Oracle Database Installation Guide 11g Release 2 (11.2) for Oracle Solaris* (http://docs.oracle.com/cd/E11882_01/install.112/e24346/title.htm) for more information about these parameters.

3 Verify the new settings.

```
phys-X# prctl -i project default
```

Verifying the Oracle Installation and Configuration

This section contains the procedure that you need to verify the Oracle installation and configuration.

▼ How to Verify the Oracle Installation

This procedure does not verify that your application is highly available because you have not yet installed your data service.

1 Confirm that the owner, group, and mode of the \$ORACLE_HOME/bin/oracle file are as follows:

- Owner: oracle
- Group: dba
- Mode: -rwsr-s--x

```
# ls -l $ORACLE_HOME/bin/oracle
```

2 Verify that the listener binaries exist in the \$ORACLE_HOME/bin directory.

Next Steps When you have completed the work in this section, go to [“Creating an Oracle Database” on page 22](#).

Creating an Oracle Database

After verifying the Oracle installation and configuration, create the Oracle databases that you require.

- If you are using Oracle *without* standby databases, perform the procedure [“How to Create a Primary Oracle Database” on page 22](#).
- If you are using Oracle Data Guard, create the following database instances:
 - **Primary database instance.** For instructions for creating a primary database, see [“How to Create a Primary Oracle Database” on page 22](#).
 - **Standby database instances.** A standby database instance can be either a physical standby database instance or a logical standby database instance. For instructions for creating standby database instances, see your Oracle documentation.

▼ How to Create a Primary Oracle Database

1 Prepare database configuration files.

Place all of the database-related files (data files, redo log files, and control files) on either shared raw global devices or on the cluster file system. See [“Preparing the Nodes and Disks” on page 15](#) for information about installation locations.

Note – If the database exists in the non-global zone, do not place the database-related files on the shared raw devices.

Within the `init$ORACLE_SID.ora` or `config$ORACLE_SID.ora` file, you might need to modify the assignments for `control_files` and `background_dump_dest` to specify the locations of the control files and alert files.

Note – If you use Oracle Solaris authentication for database logins, set the `remote_os_authent` variable in the `init$ORACLE_SID.ora` file to `True`.

2 Start the creation of the database by using one utility from the following list:

- The Oracle installer
- The Oracle `sqlplus(1M)` command
- The Oracle Database Configuration Assistant

During creation, ensure that all the database-related files are placed in the appropriate location, either on shared global devices, on the cluster file system, or on a highly available local file system.

3 Verify that the file names of your control files match the file names in your configuration files.

4 Create the `v$sysstat` view.

Run the catalog scripts that create the `v$sysstat` view. The HA for Oracle fault monitor uses this view.

Next Steps When you have completed the work in this section, go to [“Setting Up Oracle Database Permissions” on page 23](#).

Setting Up Oracle Database Permissions



Caution – Do not perform the steps in this section for an Oracle physical standby database.

Perform the procedure in this section to set up database permissions for an Oracle primary database or an Oracle logical standby database.

▼ How to Set Up Oracle Database Permissions

1 Enable access for the user and password to be used for fault monitoring.

- To use the Oracle authentication method, grant to this user authority on the `v_$sysstat` view and the `v_$archive_dest` view.

```
# sqlplus "/ as sysdba"
```

```
sql> create user user identified by passwd;
sql> alter user user default tablespace system quota 1m on system;
```

```

sql> grant select on v_$sysstat to user;
sql> grant select on v_$archive_dest to user;
sql> grant select on v_$database to user;
sql> grant create session to user;
sql> grant create table to user;
sql> create profile profile limit PASSWORD_LIFE_TIME UNLIMITED;
sql> alter user user identified by passwd profile profile;

sql> exit;
#

```

You can use this method for all the supported Oracle releases.

- To use the Oracle Solaris authentication method, perform the following steps:

- a. Confirm that the `remote_os_authent` parameter is set to `TRUE`.

```

# sqlplus "/ as sysdba"
sql> show parameter remote_os_authent

```

NAME	TYPE	VALUE
remote_os_authent	boolean	TRUE

- b. Determine the setting of the `os_authent_prefix` parameter.

```

# sql> show parameter os_authent_prefix

```

NAME	TYPE	VALUE
os_authent_prefix	string	ops\$

- c. Grant permission for the database to use Oracle Solaris authentication.

```

sql> create user prefix user identified by externally default
tablespace system quota 1m on system;
sql> grant connect, resource to prefix user;
sql> grant select on v_$sysstat to prefix user;
sql> grant select on v_$archive_dest to prefix user;
sql> grant select on v_$database to prefix user;
sql> grant create session to prefix user;
sql> grant create table to prefix user;
sql> exit;
#

```

The replaceable items in these commands are as follows:

- `prefix` is the setting of the `os_authent_prefix` parameter. The default setting of this parameter is `ops$`.
- `user` is the user for whom you are enabling Oracle Solaris authentication. Ensure that this user owns the files under the `$ORACLE_HOME` directory.

Note – Do not type a space between `prefix` and `user`.

2 Configure Oracle Net for the Oracle Solaris Cluster software.

The `listener.ora` file must be accessible from all the nodes that are in the cluster. Place these files either under the cluster file system or in the local file system of each node that can potentially run the Oracle resources.

Note – If you place the `listener.ora` file in a location other than the `/var/opt/oracle` directory or the `$ORACLE_HOME/network/admin` directory, you must specify the `TNS_ADMIN` variable or an equivalent Oracle variable in a user-environment file. For information about Oracle variables, see the Oracle documentation.

You must also run the `clresource` command to set the resource extension parameter `User_env`, which sources the user-environment file. See “[SUNW.oracle_listener Extension Properties](#)” on page 69 or “[SUNW.oracle_server Extension Properties](#)” on page 65 for format details.

HA for Oracle imposes no restrictions on the listener name; it can be any valid Oracle listener name.

The following code sample identifies the lines in `listener.ora` that are updated.

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC))
        (ADDRESS = (PROTOCOL = TCP) (HOST = logical-hostname) (PORT = port-used))
      )
    )
  )
```

The following code sample identifies the lines in `tnsnames.ora` that are updated on client machines.

```
service_name =
  .
  .
  .
  (ADDRESS =
    (PROTOCOL = TCP)
    (HOST = logicalhostname) <- logical hostname
    (PORT = 1527) <- must match port in LISTENER.ORA
  )
  (CONNECT_DATA =
    (SID = <SID>)) <- database name, default is ORCL
```

The following example shows how to update the `listener.ora` and `tnsnames.ora` files for the following Oracle instances.

Instance	Logical Host	Listener
ora8	hadbms3	LISTENER-ora8

Instance	Logical Host	Listener
ora9	hadbms4	LISTENER-ora9

The corresponding `listener.ora` entries are the following entries.

```

LISTENER-ora9 =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = hadbms4)
      (PORT = 1530)
    )
  )
SID_LIST_LISTENER-ora9 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ora9)
    )
  )
LISTENER-ora8 =
  (ADDRESS_LIST =
    (ADDRESS= (PROTOCOL=TCP) (HOST=hadbms3) (PORT=1806))
  )
SID_LIST_LISTENER-ora8 =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ora8)
    )
  )

```

The corresponding `tnsnames.ora` entries are the following entries.

```

ora8 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)
      (HOST = hadbms3)
      (PORT = 1806))
    )
    (CONNECT_DATA = (SID = ora8))
  )
ora9 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS =
        (PROTOCOL = TCP)
        (HOST = hadbms4)
        (PORT = 1530))
    )
    (CONNECT_DATA = (SID = ora9))
  )

```

3 Verify that the Oracle Solaris Cluster software is installed and running on all the nodes.

```
# cluster status clustername
```

Next Steps Go to “Installing the HA for Oracle Package” on page 27 to install the HA for Oracle packages.

Installing the HA for Oracle Package

If you did not install the HA for Oracle package during your initial Oracle Solaris Cluster installation, perform this procedure to install the package.

▼ How to Install the HA for Oracle Package

Perform this procedure on each cluster node where you want the HA for Oracle software to run.

- 1 On the cluster node where you are installing the data service package, become superuser.
- 2 Ensure that the `solaris` and `ha-cluster` publishers are valid.

```
# pkg publisher
PUBLISHER          TYPE    STATUS  URI
solaris            origin online  solaris-repository
ha-cluster         origin online  ha-cluster-repository
```

For information about setting the `solaris` publisher, see “Set the Publisher Origin To the File Repository URI” in *Copying and Creating Oracle Solaris 11 Package Repositories*.

- 3 Install the HA for Oracle software package.

```
# pkg install ha-cluster/data-service/oracle-database
```

- 4 Verify that the package installed successfully.

```
$ pkg info ha-cluster/data-service/oracle-database
```

Installation is successful if output shows that State is Installed.

- 5 Perform any necessary updates to the Oracle Solaris Cluster software.

For instructions on updating single or multiple packages, see Chapter 11, “Updating Your Software,” in *Oracle Solaris Cluster System Administration Guide*.

Registering and Configuring HA for Oracle

Tools for Registering and Configuring HA for Oracle

Oracle Solaris Cluster provides the following tools for registering and configuring HA for Oracle:

- **The `clsetup` utility.** For more information, see [“How to Register and Configure HA for Oracle \(`clsetup`\)”](#) on page 28.
- **Oracle Solaris Cluster maintenance commands.** For more information, see [“How to Register and Configure HA for Oracle Without Oracle ASM \(CLI\)”](#) on page 33.

The `clsetup` utility provides a wizard for configuring HA for Oracle. The wizard reduces the possibility for configuration errors that might result from command syntax errors or omissions. This wizard also ensures that all required resources are created and that all required dependencies between resources are set.

Setting HA for Oracle Extension Properties

Use the extension properties in [Appendix A, “HA for Oracle Extension Properties,”](#) to create your resources. To set an extension property of a resource, include the option `-p property=value` in the `clresource` command that creates or modifies the resource. Use the procedure in [Chapter 2, “Administering Data Service Resources,”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure the extension properties if you have already created your resources.

You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable entries indicate when you can update each property. See the `r_properties(5)` man page for details about all Oracle Solaris Cluster resource properties.

[“SUNW.oracle_server Extension Properties”](#) on page 65 describes the extension properties that you can set for the Oracle server. For the Oracle server, you are required to set only the following extension properties:

- `Oracle_home`
- `Oracle_sid`
- `Db_unique_name`

▼ How to Register and Configure HA for Oracle (`clsetup`)

This procedure provides the long forms of the Oracle Solaris Cluster maintenance commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

Note – The `clsetup` utility does not support configuration of HA for Oracle with ZFS.

Before You Begin Ensure that the following prerequisites are met:

- The volume manager of the cluster is configured to provide volumes on shared storage that are accessible from any node where Oracle could potentially run.
- Raw devices and file systems on the storage volumes that Oracle will use for its database are created.
- The Oracle software is installed to be accessible from all nodes where Oracle could potentially run.
- Kernel variables for the UNIX operating system are configured for Oracle.
- The Oracle software is configured for all nodes that could potentially run Oracle.
- The data service packages are installed.

Ensure that you have the following information:

- The names of the cluster nodes that master the data service.
- The path to the Oracle application binaries for the resources that you plan to configure.
- The database type.

1 Become superuser on any cluster node.

2 Start the `clsetup` utility.

```
# clsetup
```

The `clsetup` main menu is displayed.

3 Type the number that corresponds to the option for data services and press Return.

The Data Services menu is displayed.

4 Type the number that corresponds to the option for configuring HA for Oracle and press Return.

The `clsetup` utility displays the list of prerequisites for performing this task.

5 Verify that the prerequisites are met, and press Return.

The `clsetup` utility displays a list of the cluster nodes.

6 Select the nodes where you require Oracle to run.

- To accept the default selection of all listed nodes in an arbitrary order, press Return.

- **To select a subset of the listed nodes, type a comma-separated or space-separated list of the numbers that correspond to the nodes. Then press Return.**

Ensure that the nodes are listed in the order in which the nodes are to appear in the node list of the resource group in which the Oracle resource is to be placed.

- **To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the numbers that correspond to the nodes and press Return.**

Ensure that the nodes are listed in the order in which the nodes are to appear in the node list of the resource group in which the Oracle resource is to be placed.

7 To confirm your selection of nodes, type `d` and press Return.

The `clsetup` utility displays the types of Oracle components that are to be configured.

8 Type the numbers of the Oracle components you want to configure and press Return.

The `clsetup` utility lists the Oracle home directory.

9 Specify the Oracle home directory for your installation of the Oracle software.

- **If the directory is listed, select the directory as follows:**

a. Type the number that corresponds the directory that you are selecting.

The `clsetup` utility displays a list of Oracle system identifiers that are configured on the cluster. The utility also prompts you to specify the system identifier for your installation of Oracle.

- **If the directory is not listed, specify the directory explicitly.**

a. Type `e` and press Return.

The `clsetup` utility prompts you for the Oracle home directory.

b. Type the full path to the Oracle home directory and press Return.

The `clsetup` utility displays a list of Oracle system identifiers that are configured on the cluster. The utility also prompts you to specify the system identifier for your installation of Oracle.

10 Specify the Oracle SID of the Oracle database that you are configuring.

- **If the SID is listed, select the SID as follows:**

a. Type the number that corresponds the SID that you are selecting.

The `clsetup` utility displays the properties of the Oracle Solaris Cluster resources that the utility will create.

- **If the SID is not listed, specify the SID explicitly.**

- a. **Type e and press Return.**

The `clsetup` utility prompts you for the SID.

- b. **Type the SID and press Return.**

The `clsetup` utility displays the properties of the Oracle Solaris Cluster resources that the utility will create.

The `clsetup` utility displays the properties of the Oracle Solaris Cluster resources that the utility will create.

- 11 If you require a different name for any Oracle Solaris Cluster resources properties, change each value as follows.**

- a. **Type the number that corresponds to the name that you are changing and press Return.**

The `clsetup` utility displays a screen where you can specify the new name.

- b. **At the New Value prompt, type the new name and press Return.**

The `clsetup` utility returns you to the list of the properties of the Oracle Solaris Cluster resource that the utility will create.

- 12 To confirm your selection of Oracle Solaris Cluster resource properties, type d and press Return.**

The `clsetup` utility displays a list of existing storage resources. If no storage resources are available, the `clsetup` utility displays a list of shared storage types where data is to be stored.

- 13 Type the numbers that correspond to type of shared storage that you are using for storing the data and press Return.**

The `clsetup` utility displays the file-system mount points that are configured in the cluster.

- 14 Select the file system mount points as follows.**

- **To accept the default selection of all listed file-system mount points in an arbitrary order, type a. Then press Return.**
- **To select a subset of the listed file system mount points, type a comma-separated or space-separated list of the numbers that correspond to the file-system mount points. Then press Return.**

The `clsetup` utility displays the global disk sets and device groups that are configured in the cluster.

15 Select the device groups as follows.

- **To accept the default selection of all listed device groups in an arbitrary order, type a and press Return.**
- **To select a subset of the listed device groups, type a comma-separated or space-separated list of the numbers that correspond to the device groups and press Return.**

The `clsetup` utility returns to you the list of highly available storage resources.

16 Type a comma-separated or space-separated list of the numbers that correspond to the storage resources that your data service requires, and press Return.

17 To confirm your selection of Oracle Solaris Cluster storage resources, type d and press Return.

The `clsetup` utility displays all the existing logical hostname resources in the cluster. If there are no logical hostname resources available, the `clsetup` utility prompts for the logical hostname that the resource will make highly available.

18 Specify the logical hostname and press Return.

The `clsetup` utility returns to you the list of available logical hostname resources.

19 Type a comma-separated or space-separated list of the numbers that correspond to the logical hostname resources that your data service requires, and press Return.

20 To confirm your selection of Oracle Solaris Cluster logical hostname resources, type d and press Return.

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects that the utility will create.

21 If you require a different name for any Oracle Solaris Cluster objects, change each name as follows.

- **Type the number that corresponds to the name that you are changing and press Return.**

The `clsetup` utility displays a screen where you can specify the new name.

- **At the New Value prompt, type the new name and press Return.**

The `clsetup` utility returns you to the list of the names of the Oracle Solaris Cluster objects that the utility will create.

22 To confirm your selection of Oracle Solaris Cluster object names, type d and press Return.

23 To create the configuration, type c and Press Return.

The `clsetup` utility displays a progress message to indicate that the utility is running commands to create the configuration. When configuration is complete, the `clsetup` utility displays the commands that the utility ran to create the configuration.

24 Press Return to continue.**25 (Optional) Type q and press Return repeatedly until you quit the clsetup utility.**

If you prefer, you can leave the `clsetup` utility running while you perform other required tasks before using the utility again.

▼ How to Register and Configure HA for Oracle Without Oracle ASM (CLI)

This procedure provides the long forms of the Oracle Solaris Cluster maintenance commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

Before You Begin Ensure that the following prerequisites are met:

- The volume manager of the cluster is configured to provide volumes on shared storage that are accessible from any node where Oracle could potentially run.
- Raw devices and file systems on the storage volumes that Oracle will use for its database are created.
- The Oracle software is installed to be accessible from all nodes where Oracle could potentially run.
- Kernel variables for the UNIX operating system are configured for Oracle.
- The Oracle software is configured for all nodes that could potentially run Oracle.
- The data service packages are installed.

Ensure that you have the following information:

- The names of the cluster nodes that master the data service.
- The logical hostname that clients use to access the data service. Normally, you set up this IP address when you install the cluster. See the *Oracle Solaris Cluster Concepts Guide* for details about network resources.
- The path to the Oracle application binaries for the resources that you plan to configure.
- The database type.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.**

2 Register the resource types for the data service.

For HA for Oracle, you register two resource types, `SUNW.oracle_server` and `SUNW.oracle_listener`, as follows.

Note – If you are using an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener, omit registration of the `SUNW.oracle_listener` resource type.

```
# clresourcetype register SUNW.oracle_server
# clresourcetype register SUNW.oracle_listener
```

3 Create a failover resource group to hold the network and application resources.

This step is not required if you use the Oracle Solaris ZFS file system, because the resource group was created when the highly available local ZFS file system was configured in [“How to Prepare the Nodes” on page 15](#). The resources that are created in other steps in this procedure are to be added to this resource group.

You can optionally select the set of nodes on which the data service can run with the `-n` option, as follows.

```
# clresourcegroup create [-n node-zone-list] resource-group
resource-group    Specifies the name of the resource group. This name can be your choice but
                  must be unique for resource groups within the cluster.
```

4 Verify that all of the network resources that you use have been added to your name service database.

You should have performed this verification during the Oracle Solaris Cluster installation.

Note – Ensure that all of the network resources are present in the server's and client's `/etc/inet/hosts` file to avoid any failures because of name service lookup.

5 Add a logical hostname resource to the failover resource group.

```
# clreslogicalhostname create -g resource-group [-h logicalhostname] logicalhostname-rs
logicalhostname    Specifies a logical hostname. This logical hostname must present in your
                  name service database. If logicalhostname and logicalhostname-rs are
                  identical, logicalhostname is optional.

logicalhostname-rs Specifies the name that you are assigning to the logical hostname
                  resource that you are creating.
```

6 Register the HASStoragePlus resource type with the cluster.

```
# clresourcetype register SUNW.HASStoragePlus
```

7 Add a resource of type HASStoragePlus to the failover resource group.

Note – If you use the Oracle Solaris ZFS file system for Oracle files, omit this step. The HASStoragePlus resource was created when the highly available local ZFS file system was configured. For more information, see [“How to Prepare the Nodes” on page 15](#).



Caution – Raw devices from Oracle Solaris Cluster device groups are not supported in non-global zones.

```
# clresource create -g resource-group -t SUNW.HASStoragePlus \
-p GlobalDevicePaths=device-path \
-p FilesystemMountPoints=mount-point-list \
-p AffinityOn=TRUE hastp-rs
```

You must set either the GlobalDevicePaths extension property or the FilesystemMountPoints extension property:

- If your database is on a raw device, set the GlobalDevicePaths extension property to the global device path.
- If your database is on the cluster file system, specify mount points of the cluster file system and the local file system.

Note – AffinityOn must be set to TRUE and the local file system must reside on global disk groups to be failover.

The resource is created in the enabled state.

8 Bring online the failover resource group in a managed state on a cluster node.

```
# clresourcegroup online -M resource-group
```

-M Places the resource group that is brought online in a managed state.

9 Create Oracle application resources in the failover resource group.

Note – If you are using an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener, omit configuration of the SUNW.oracle_listener resource.

- Oracle server resource:

```
# clresource create -g resourcegroup \
-t SUNW.oracle_server \
-p Oracle_sid=instance \
-p Oracle_home=Oracle_home \
-p Db_unique_name=db-unique-name \
-p Restart_type=entity-to-restart \
[-p Dataguard_role=role] \
```

```
[-p Standby_mode=mode] \  
-p Resource_dependencies_offline_restart=storageplus-resource \  
resource
```

- Oracle listener resource:

```
# clresource create -g resource-group \  
-t SUNW.oracle_listener \  
-p Listener_name=listener \  
-p Oracle_home=Oracle_home \  
-p Resource_dependencies_offline_restart=storageplus-resource \  
resource
```

-g *resource-group*

Specifies the name of the resource group into which the resources are to be placed.

-t

Specifies the type of the resource to add.

-p *Oracle_sid=instance*

Sets the Oracle system identifier.

-p *Listener_name=listener*

Sets the name of the Oracle listener instance. This name must match the corresponding entry in `listener.ora`.

-p *Oracle_home=Oracle_home*

Sets the path to the Oracle home directory.

-p *Db_unique_name=db-unique-name*

Sets the unique name of the database.

-p *Restart_type=entity-to-restart*

Specifies the entity that the server fault monitor restarts when the response to a fault is restart. Set *entity-to-restart* as follows:

- To specify that only this resource is restarted, set *entity-to-restart* to `RESOURCE_RESTART`. By default, only this resource is restarted.
- To specify that all resources in the resource group that contains this resource are restarted, set *entity-to-restart* to `RESOURCE_GROUP_RESTART`.

If you set *entity-to-restart* to `RESOURCE_GROUP_RESTART`, all other resources (such as Apache or DNS) in the resource group are restarted, even if they are not faulty. Therefore, include in the resource group only the resources that you require to be restarted when the Oracle server resource is restarted.

-p *Dataguard_role=role*

Specifies the role of the database instance. Change *role* as follows:

- To create a resource for a primary database instance that does not have standby instances configured, change *role* to `NONE`. This value is the default value.
- To create a resource for a primary database instance that has standby database instances configured, change *role* to `PRIMARY`.

- To create a resource for a standby database instance, change *role* to STANDBY.
- p Standby_mode=*mode*
Specifies the mode for the standby database instance. If you change Dataguard_role to NONE or PRIMARY, the value of the Standby_mode is ignored.
- To specify a logical standby database, change *mode* to LOGICAL. This value is the default value.
 - To specify a physical standby database, change *mode* to PHYSICAL.
 - To specify a snapshot standby database, change *mode* to SNAPSHOT.

resource

Specifies the name of the resource that you are creating.

Note – Optionally, you can set additional extension properties that belong to the Oracle data service to override their default values. See “[Setting HA for Oracle Extension Properties](#)” on [page 28](#) for a list of extension properties.

The resources are created in the enabled state.

Example 1-1 Registering HA for Oracle to Run in the Global Zone

This example shows how to register HA for Oracle on a two-node cluster. The following are the sample names used in the commands:

Node names

phys-schost-1, phys-schost-2

Logical hostname

schost-1

Resource group

resource-group-1 (failover resource group)

HAStoragePlus resource

hastp-rs

Oracle resources

oracle-server-1, oracle-listener-1

Oracle instances

ora-lsnr (listener), ora-srvr (server)

Oracle database unique name

dbcloud

Create the failover resource group to contain all of the resources.

```
# clresourcegroup create resource-group-1
```

```

    Add the logical hostname resource to the resource group.
# clreslogicalhostname create -g resource-group-1 schost-1

    Register the SUNW.HAStoragePlus resource type.
# clresourcetype register SUNW.HAStoragePlus

    Add a resource of type SUNW.HAStoragePlus to the resource group.
# clresource create -g resource-group-1 \
-t SUNW.HAStoragePlus \
-p FileSystemMountPoints=/global/oracle,/global/ora-data/logs,/local/ora-data \
-p AffinityOn=TRUE \
hastp-rs
    Bring the resource group online in a managed state
# clresourcegroup online -M resource-group-1

    Register the Oracle resource types.
# clresourcetype register SUNW.oracle_server
# clresourcetype register SUNW.oracle_listener

    Add the Oracle application resources to the resource group.
# clresource create -g resource-group-1 \
-t SUNW.oracle_server \
-p Oracle_home=/global/oracle \
-p Oracle_sid=ora-srvr \
-p Db_unique_name=dbcloud \
-p Dataguard_role=STANDBY \
-p Standby_mode=PHYSICAL \
-p Resource_dependencies_offline_restart=hastp-rs \
oracle-server-1

# clresource create -g resource-group-1 \
-t SUNW.oracle_listener \
-p Oracle_home=/global/oracle \
-p Listener_name=ora-lsnr \
oracle-listener-1

```

Example 1-2 Registering HA for Oracle to Run in a Zone Cluster

This example shows how to register HA for Oracle in a zone cluster. The following are the sample names used in the commands, which are issued from the global cluster:

Node names

phys-schost-1, phys-schost-2

Zone cluster names

zonecluster1, zonecluster2

Logical hostname

zchost-1

Resource group
 resource-group-1 (failover resource group)

HAStoragePlus resource
 hastp-rs

Oracle resources
 oracle-server-1, oracle-listener-1

Oracle instances
 ora-lsnr (listener), ora-srvr (server)

Oracle database unique name
 dbcloud

```

    Create the failover resource group to contain all of the resources.
# clresourcegroup create -Z zonecluster1 resource-group-1

    Add the logical hostname resource to the resource group.
# clreslogicalhostname create -Z zonecluster1 -g resource-group-1 zchost-1

    Register the SUNW.HAStoragePlus resource type.
# clresourcetype register -Z zonecluster1 SUNW.HAStoragePlus

    Add a resource of type SUNW.HAStoragePlus to the resource group.
# clresource create -Z zonecluster1 \
-g resource-group-1 \
-t SUNW.HAStoragePlus \
-p FileSystemMountPoints=/global/oracle,/global/ora-data/logs,/local/ora-data \
-p AffinityOn=TRUE \
hastp-rs

    Bring the resource group online in a managed state
# clresourcegroup online -Z zonecluster1 -M resource-group-1

    Register the Oracle resource types.
# clresourcetype register -Z zonecluster1 SUNW.oracle_server
# clresourcetype register -Z zonecluster1 SUNW.oracle_listener

    Add the Oracle application resources to the resource group.
# clresource create -Z zonecluster1 \
-g resource-group-1 \
-t SUNW.oracle_server \
-p Oracle_home=/global/oracle \
-p Oracle_sid=ora-srvr \
-p Db_unique_name=dbcloud \
-p Dataguard_role=STANDBY \
-p Standby_mode=PHYSICAL \
oracle-server-1

```

```
# clresource create -Z zonecluster1 -g resource-group-1 \
-t SUNW.oracle_listener \
-p Oracle_home=/global/oracle \
-p Listener_name=ora-lsnr \
oracle-listener-1
```

▼ How to Create an Oracle Grid Infrastructure Resource With Clustered Oracle ASM Disk Groups and a Third-Party Volume Manager (CLI)

Perform this procedure to include a third-party volume manager with clustered Oracle ASM disk groups in your configuration that uses clustered Oracle ASM. This procedure uses Oracle Grid Infrastructure (Grid Infrastructure) commands and Oracle Solaris Cluster maintenance commands.

Use this procedure to manually create an Oracle Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` resource. In this procedure, the Grid Infrastructure resource is named `sun.resource`. You configure `sun.resource` to ensure that the corresponding Oracle ASM disk group is not mounted until `sun.resource` is online. The `sun.resource` resource comes online only if the corresponding `SUNW.ScalDeviceGroup` resource is online. And the `SUNW.ScalDeviceGroup` resource only comes online if the actual volume-manager disk set or disk group is online.

To ensure that the Oracle ASM disk group benefits from this dependency chain, after you define `sun.resource`, you modify the appropriate Grid Infrastructure Oracle ASM disk-group resource so that the hard-start dependency includes `sun.resource`. Modifying the hard-start dependency of the Grid Infrastructure Oracle ASM disk-group resource can only be performed by the `SUNW.scalable_asm_diskgroup_proxy` resource by using the `VALIDATE` method. Therefore, you must set an offline-restart dependency between the `SUNW.scalable_asm_diskgroup_proxy` and `SUNW.ScalDeviceGroup` resources.

Before You Begin Configure the disk set or disk group for clustered Oracle ASM to use.

- If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. Follow procedures in [“How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database”](#) in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide*.

- 1 **Become superuser.**
- 2 **Create the Grid Infrastructure `sun.storage_proxy.type` resource type.**

```
# /Grid_home/bin/crsctl add type sun.storage_proxy.type -basetype cluster_resource
```

3 Create a Grid Infrastructure `sun.resource` resource of type `sun.storage_proxy.type`.

Note – Ensure that all attribute values are enclosed in single quotes (`'`). Otherwise, the `VALIDATE` method of the `SUNW.scalable_asm_diskgroup_proxy` resource will fail the validation.

```
# /Grid_home/bin/crsctl add res sun.scal-asmdg1-rs -type sun.storage_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
HOSTING_MEMBERS='node1 node2' \
CARDINALITY='2' \
PLACEMENT='restricted' \
ACL='owner:root:rxw,pgrp:oinstall:rxw,other::r--' \
SCRIPT_TIMEOUT='20' \
RESTART_ATTEMPTS='60'"
```

<code>sun.scal-asmdg1-rs</code>	The <code>SUNW.ScalDeviceGroup</code> resource name.
<code>-type sun.storage_proxy.type</code>	Specifies the <code>sun.storage_proxy.type</code> resource type.
<code>ACTION_SCRIPT</code>	Specifies the <code>/opt/SUNWscor/dsconfig/bin/scproxy_crs_action</code> action script.
<code>HOSTING_MEMBERS</code>	Specifies the Oracle Solaris Cluster resource group's node list entries that contain the <code>SUNW.ScalDeviceGroup</code> resource.
<code>CARDINALITY</code>	Sets the number of hosts that are defined in <code>HOSTING_MEMBERS</code> .
<code>PLACEMENT</code>	Set to <code>restricted</code> .
<code>ACL</code>	Sets the owner equal to <code>root</code> and the group equal to the <code>ACL</code> group entry for the Oracle ASM disk group. The following example command displays the <code>ACL</code> group entry, <code>oinstall</code> :

```
# /Grid_home/bin/crsctl stat res ora.DATA1.dg -p | grep ACL=
ACL=owner:oragrid:rxw,pgrp:oinstall:rxw,other::r--
```

<code>SCRIPT_TIMEOUT</code>	Set to 20.
<code>RESTART_ATTEMPTS</code>	Set to 60.

4 Verify that `sun.resource` is correctly defined.

Output is similar to the following:

```
# /Grid_home/bin/crsctl stat res sun.scal-asmdg1-rs -p
NAME=sun.scal-asmdg1-rs
TYPE=sun.storage_proxy.type
ACL=owner:root:rxw,pgrp:oinstall:rxw,other::r--
ACTION_FAILURE_TEMPLATE=
```

```

ACTION_SCRIPT=/opt/SUNWscor/dsconfig/bin/scproxy_crs_action
ACTIVE_PLACEMENT=0
AGENT_FILENAME=%CRS_HOME%/bin/scriptagent
AUTO_START=restore
CARDINALITY=2
CHECK_INTERVAL=60
DEFAULT_TEMPLATE=
DEGREE=1
DESCRIPTION=
ENABLED=1
FAILOVER_DELAY=0
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=node1 node2
LOAD=1
LOGGING_LEVEL=1
NOT_RESTARTING_TEMPLATE=
OFFLINE_CHECK_INTERVAL=0
PLACEMENT=restricted
PROFILE_CHANGE_TEMPLATE=
RESTART_ATTEMPTS=60
SCRIPT_TIMEOUT=20
SERVER_POOLS=
START_DEPENDENCIES=
START_TIMEOUT=0
STATE_CHANGE_TEMPLATE=
STOP_DEPENDENCIES=
STOP_TIMEOUT=0
UPTIME_THRESHOLD=1h
#

```

5 Set the offline-restart dependency to add a dependency.

a. Display the current dependency.

```

# /Grid_home/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(ora.asm) pullup(ora.asm)

# clresource show -p Resource_dependencies_offline_restart asm-data1-rs

=== Resources ===

Resource: asm-data1-rs
Resource_dependencies_offline_restart: asm-inst-rs

-- Standard and extension properties --

```

b. Set the new dependency.

- If *asm-data1-rs* already exists, use the following command to set the dependency.

Note that the command includes the plus (+) symbol:

```
# clresource set -p Resource_dependencies_offline_restart+=scal-asm1-rs asm-data1-rs
```

- If `asm-data1-rs` does not yet exist, use the following command to create the resource with the offline-restart dependency:

```
# clresource create -g asm-dg-rg \  
-t SUNW.scalable_asm_diskgroup_proxy \  
-p asm_diskgroups=data1 \  
-p Resource_dependencies_offline_restart=asm-inst-rs,scal-asm1-rs \  
-d asm-data1-rs
```

c. Verify the configured dependency.

```
# /Grid_home/bin/crsctl stat res ora.DATA1.dg -p | grep START_DEPENDENCIES  
START_DEPENDENCIES=hard(ora.asm,sun.scal-asm1-rs) pullup(ora.asm)  
# clresource show -p Resource_dependencies_offline_restart asm-data1-rs  
=== Resources ===  
  
Resource: asm-data1-rs  
Resource_dependencies_offline_restart: asm-inst-rs scal-asm1-rs  
  
-- Standard and extension properties --
```

Next Steps Go to “How to Register and Configure HA for Oracle With Clustered Oracle ASM Instance (CLI)” on page 43.

▼ How to Register and Configure HA for Oracle With Clustered Oracle ASM Instance (CLI)

This procedure explains the steps to register and configure HA for Oracle with clustered Oracle ASM instance using Oracle Solaris Cluster maintenance commands.

Note – You should install the Oracle Clusterware software before performing the steps in this procedure.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.

Note – You can also use the `clsetup` utility to perform the steps 1 to 6 in the procedure.

- 2 Register the resource types for Oracle RAC framework and Oracle Clusterware framework.

```
# clresourcetype register SUNW.rac_framework  
# clresourcetype register SUNW.crs_framework
```

- 3 Create a scalable resource group `rac-fmwk-rg`.

```
# clresourcegroup create -S rac-fmwk-rg
```

- 4 Add a resource of type SUNW.rac_framework to the rac-fmwk-rg resource group.**

```
# clresource create -g rac-fmwk-rg -t SUNW.rac_framework rac-fmwk-rs
```

- 5 Add a resource of type SUNW.crs_framework to the rac-fmwk-rg resource group.**

```
# clresource create -g rac-fmwk-rg -t SUNW.crs_framework \  
-p Resource_dependencies_offline_restart=rac-fmwk-rs \  
-d crs-fmwk-rs
```

- 6 Create a failover resource group ora-db-rg for the Oracle database.**

```
# clresourcegroup create ora-db-rg
```

- 7 Register the SUNW.HASStoragePlus resource type and create a resource group.**
-

Note – For information on how to set up a file system, see “Preparing the Nodes and Disks” on page 15.

- 8 Register the Oracle ASM resource types for the data service.**

- a. Register the scalable ASM instance proxy resource type.**

```
# clresourcetype register SUNW.scalable_asm_instance_proxy
```

- b. Register the appropriate ASM disk-group resource type.**

```
# clresourcetype register SUNW.scalable_asm_diskgroup_proxy
```

- 9 Create resource groups asm-inst-rg and asm-dg-rg.**

```
# clresourcegroup create -S asm-inst-rg asm-dg-rg
```

- 10 Set a strong positive affinity on rac-fmwk-rg by asm-inst-rg.**

```
# clresourcegroup set -p Rg_affinities=++rac-fmwk-rg asm-inst-rg
```

- 11 Set a strong positive affinity on asm-inst-rg by asm-dg-rg.**

```
# clresourcegroup set -p Rg_affinities=++asm-inst-rg asm-dg-rg
```

- 12 Add a resource of type SUNW.crs_framework to the rac-fmwk-rg resource group.**

```
# clresource create -g rac-fmwk-rg \  
-t SUNW.crs_framework \  
-p Resource_dependencies_offline_restart=rac-fmwk-rs \  
-d crs-fmwk-rs
```

- 13 Add a resource of type SUNW.scalable_asm_instance_proxy to the asm-inst-rg resource group.**

```
# clresource create -g asm-inst-rg \  
-t SUNW.asm_instance_proxy \  
-p Oracle_home=Oracle_home \  
-p "oracle_sid{node1}"=instance \  
-d asm-inst-rg
```

```

-p "oracle_sid{node2}"=instance \
-p Resource_dependencies_offline_restart=crs-fmwk-rs \
-d asm-inst-rs

-g asm-inst-rg
    Specifies the name of the resource group into which the resources are to be placed.

-t SUNW.asm_inst_proxy
    Specifies the type of the resource to add.

-p oracle_sid{node}=instance
    Sets the Oracle system identifier.

-p Oracle_home=Oracle_home
    Sets the path to the Oracle home directory.

-p Resource_dependencies_offline_restart=crs-fmwk-rs
    Specifies the offline restart dependency.

-d asm-inst-rs
    Specifies the name of the resource that you are creating.

```

14 Add an ASM disk-group resource to the *asm-dg-rg* resource group.

Use the `SUNW.scalable_asm_diskgroup_proxy` resource type.

```

# clresource create -g asm-dg-rg -t SUNW.scalable_asm_diskgroup_proxy \
-p Asm_diskgroups=dg[,dg...] \
-p Resource_dependencies_offline_restart=asm-inst-rs,asm-stor-rs \
-d asm-dg-rs

```

15 Bring online the *asm-inst-rg* resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM asm-inst-rg
```

16 Bring online the *asm-dg-rg* resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM asm-dg-rg
```

17 Verify the Oracle ASM installation by issuing the status command.

```
# clresource status +
```

18 Register the resource types for the HA for Oracle data service.

For HA for Oracle, you register two resource types, `SUNW.oracle_server` and `SUNW.oracle_listener`, as follows.

Note – If you are using an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener, omit registration of the `SUNW.oracle_listener` resource type.

```

# clresourcetype register SUNW.oracle_server
# clresourcetype register SUNW.oracle_listener

```

19 Add a logical hostname resource to the failover resource group for Oracle database.

```
# clreslogicalhostname create -g ora-db-rg [-h logicalhostname] \
logicalhostname-rs
```

logicalhostname Specifies a logical hostname. This logical hostname must present in your name service database. If *logicalhostname* and *logicalhostname-rs* are identical, *logicalhostname* is optional.

logicalhostname-rs Specifies the name that you are assigning to the logical hostname resource that you are creating.

20 Bring online the failover resource group in a managed state on a cluster node.

```
# clresourcegroup online -eM ora-db-rg
```

21 Create Oracle application resources in the failover resource group.

Note – If you are using an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener, omit configuration of the SUNW.oracle_listener resource.

- Oracle server resource:

```
# clresource create -g ora-db-rg \
-t SUNW.oracle_server \
-p Oracle_sid=instance \
-p Oracle_home=Oracle_home \
-p Db_unique_home=db-unique-home \
-p Resource_dependencies_offline_restart=asm-dg-rs \
-d ora-db-rs
```

- Oracle listener resource:

```
# clresource create -g ora-db-rg \
-t SUNW.oracle_listener \
-p Oracle_home=Oracle_home \
-p Listener_name=listener \
-p Resource_dependencies_offline_restart=asm-dg-rs \
-d ora-ls-rs
```

-g *ora-db-rg*

Specifies the name of the resource group into which the resources are to be placed.

-t SUNW.oracle_server/listener

Specifies the type of the resource to add.

-p Oracle_sid=*instance*

Sets the Oracle system identifier.

-p Oracle_home=*Oracle_home*

Sets the path to the Oracle home directory.

-p Db_unique_name=*db-unique-name*

Sets the Oracle database unique name.

-p `listener_name=listener`

Sets the name of the Oracle listener instance. This name must match the corresponding entry in `listener.ora`.

-d `ora-ls-rs`

Specifies the name of the resource that you are creating.

Note – Optionally, you can set additional extension properties that belong to the Oracle data service to override their default values. See [“Setting HA for Oracle Extension Properties” on page 28](#) for a list of extension properties.

22 Bring online the Oracle server resource.

```
# clresource enable ora-db-rs
```

23 If configured, bring online the Oracle listener resource.

```
# clresource enable ora-ls-rs
```

Where to Go From Here

Go to [“Verifying the HA for Oracle Installation” on page 47](#) after you register and configure HA for Oracle.

Verifying the HA for Oracle Installation

Perform the following verification tests to make sure that you have correctly installed HA for Oracle.

These sanity checks ensure that all the nodes that run HA for Oracle can start the Oracle instance and that the other nodes in the configuration can access the Oracle instance. Perform these sanity checks to isolate any problems in starting the Oracle software from HA for Oracle.

▼ How to Verify the HA for Oracle Installation

- 1 Log in as `oracle` to the node that currently masters the Oracle resource group.
- 2 Set the environment variables `ORACLE_SID` and `ORACLE_HOME`.
- 3 Confirm that you can start the Oracle instance from this node.

4 Confirm that you can connect to the Oracle instance.

Use the `sqlplus` command with the `user/password` variable that is defined in the `connect_string` property.

```
# sqlplus sysdba/passwd@tns_service
```

5 Shut down the Oracle instance.

The Oracle Solaris Cluster software restarts the Oracle instance because the Oracle instance is under Oracle Solaris Cluster control.

6 Switch the resource group that contains the Oracle database resource to another cluster member.

```
# clresourcegroup switch -n node-zone-list resource-group
```

`resource-group` Specifies the name of the resource group that you are switching.

7 Log in as `oracle` to the node that now contains the resource group.**8 Repeat [Step 3](#) and [Step 4](#) to confirm interactions with the Oracle instance.**

Oracle Clients

Clients must always refer to the database by using the network resource, not the physical hostname. The network resource is an IP address that can move between physical nodes during failover. The physical hostname is a machine name.

For example, in the `tnsnames.ora` file, you must specify the network resource as the host on which the database instance is running. See [“How to Set Up Oracle Database Permissions” on page 23](#).

Note – Oracle client-server connections cannot survive a HA for Oracle switchover. The client application must be prepared to handle disconnection and reconnection or recovery as appropriate. A transaction monitor might simplify the application. Further, HA for Oracle node recovery time is application dependent.

Location of HA for Oracle Log Files

Each instance of the HA for Oracle data service maintains log files in subdirectories of the `/var/opt/SUNWscor` directory.

- The `/var/opt/SUNWscor/oracle_server` directory contains log files for the Oracle server.

- The `/var/opt/SUNWscor/oracle_listener` directory contains log files for the Oracle listener.
- The `/var/opt/SUNWscor/oracle_asm` directory contains log file for Oracle ASM.

These files contain information about actions that the HA for Oracle data service performs. Refer to these files to obtain diagnostic information for troubleshooting your configuration or to monitor the behavior of the HA for Oracle data service.

Tuning the HA for Oracle Fault Monitors

Fault monitoring for the HA for Oracle data service is provided by the following fault monitors:

- The Oracle server fault monitor
- The Oracle listener fault monitor

Note – If you are using an Oracle Grid Infrastructure for Clusters Single Client Access Name (SCAN) listener, no fault monitoring is provided for the SCAN listener by Oracle Solaris Cluster software.

Each fault monitor is contained in a resource whose resource type is shown in the following table.

TABLE 1-3 Resource Types for HA for Oracle Fault Monitors

Fault Monitor	Resource Type
Oracle server	SUNW.oracle_server
Oracle listener	SUNW.oracle_listener

System properties and extension properties of these resources control the behavior of the fault monitors. The default values of these properties determine the preset behavior of the fault monitors. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the HA for Oracle fault monitors *only* if you need to modify this preset behavior.

Tuning the HA for Oracle fault monitors involves the following tasks:

- Setting the interval between fault monitor probes
- Setting the timeout for fault monitor probes
- Defining the criteria for persistent faults
- Specifying the failover behavior of a resource

For more information, see [“Tuning Fault Monitors for Oracle Solaris Cluster Data Services”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*. Information about the HA for Oracle fault monitors that you need to perform these tasks is provided in the subsections that follow.

Tune the HA for Oracle fault monitors when you register and configure HA for Oracle. For more information, see [“Registering and Configuring HA for Oracle”](#) on page 28.

Operation of the Oracle Server Fault Monitor

The fault monitor for the Oracle server uses a request to the server to query the health of the server.

The server fault monitor is started through `pmfadm` to make the monitor highly available. If the monitor is killed for any reason, the Process Monitor Facility (PMF) automatically restarts the monitor.

The server fault monitor consists of the following processes.

- A main fault monitor process
- A database client fault probe

This section contains the following information about the server fault monitor:

- [“Operation of the Main Fault Monitor”](#) on page 50
- [“Operation of the Database Client Fault Probe”](#) on page 50
- [“Actions by the Server Fault Monitor in Response to a Database Transaction Failure”](#) on page 51
- [“Scanning of Logged Alerts by the Server Fault Monitor”](#) on page 52

Operation of the Main Fault Monitor

The main fault monitor determines that an operation is successful if the database is online and no errors are returned during the transaction.

Operation of the Database Client Fault Probe

The database client fault probe performs the following operations:

1. Monitoring the partition for archived redo logs. See [“Operations to Monitor the Partition for Archived Redo Logs”](#) on page 51.
2. If the partition is healthy, determining whether the database is operational. See [“Operations to Determine Whether the Database is Operational”](#) on page 51.

The probe uses the timeout value that is set in the resource property `Probe_timeout` to determine how much time to allocate to successfully probe Oracle.

Operations to Monitor the Partition for Archived Redo Logs

The database client fault probe queries the dynamic performance view `v$aarchive_dest` to determine all possible destinations for archived redo logs. For every active destination, the probe determines whether the destination is healthy and has sufficient free space for storing archived redo logs.

- If the destination is healthy, the probe determines the amount of free space in the destination's file system. If the amount of free space is less than 10% of the file system's capacity and is less than 20 Mbytes, the probe prints a message to `syslog`.
- If the destination is in `ERROR` status, the probe prints a message to `syslog` and disables operations to determine whether the database is operational. The operations remain disabled until the error condition is cleared.

Operations to Determine Whether the Database is Operational

If the partition for archived redo logs is healthy, the database client fault probe queries the dynamic performance view `v$sysstat` to obtain database performance statistics. Changes to these statistics indicate that the database is operational. If these statistics remain unchanged between consecutive queries, the fault probe performs database transactions to determine if the database is operational. These transactions involve the creation, updating, and dropping of a table in the user table space.

The database client fault probe performs all its transactions as the Oracle user. The ID of this user is specified during the preparation of the nodes as explained in [“How to Prepare the Nodes” on page 15](#).

Actions by the Server Fault Monitor in Response to a Database Transaction Failure

If a database transaction fails, the server fault monitor performs an action that is determined by the error that caused the failure. To change the action that the server fault monitor performs, customize the server fault monitor as explained in [“Customizing the HA for Oracle Server Fault Monitor” on page 53](#).

If the action requires an external program to be run, the program is run as a separate process in the background.

Possible actions are as follows:

- **Ignore.** The server fault monitor ignores the error.
- **Stop monitoring.** The server fault monitor is stopped without shutting down the database.
- **Restart.** The server fault monitor stops and restarts the entity that is specified by the value of the `Restart_type` extension property:

- If the `Restart_type` extension property is set to `RESOURCE_RESTART`, the server fault monitor restarts the database server resource. By default, the server fault monitor restarts the database server resource.
- If the `Restart_type` extension property is set to `RESOURCE_GROUP_RESTART`, the server fault monitor restarts the database server resource group.

Note – The number of attempts to restart might exceed the value of the `Retry_count` resource property within the time that the `Retry_interval` resource property specifies. If this situation occurs, the server fault monitor attempts to switch over the resource group to another node.

- **Switch over.** The server fault monitor switches over the database server resource group to another node. If no nodes are available, the attempt to switch over the resource group fails. If the attempt to switch over the resource group fails, the database server is restarted.

Scanning of Logged Alerts by the Server Fault Monitor

The Oracle software logs alerts in an alert log file. The absolute path of this file is specified by the `alert_log_file` extension property of the `SUNW.oracle_server` resource. The server fault monitor scans the alert log file for new alerts at the following times:

- When the server fault monitor is started
- Each time that the server fault monitor queries the health of the server

If an action is defined for a logged alert that the server fault monitor detects, the server fault monitor performs the action in response to the alert.

Preset actions for logged alerts are listed in [Table B-2](#). To change the action that the server fault monitor performs, customize the server fault monitor as explained in [“Customizing the HA for Oracle Server Fault Monitor” on page 53](#).

Operation of the Oracle Listener Fault Monitor

The Oracle listener fault monitor checks the status of an Oracle listener.

If the listener is running, the Oracle listener fault monitor considers a probe successful. If the fault monitor detects an error, the listener is restarted.

Note – The listener resource does not provide a mechanism for setting the listener password. If Oracle listener security is enabled, a probe by the listener fault monitor might return Oracle error TNS-01169. Because the listener is able to respond, the listener fault monitor treats the probe as a success. This action does not cause a failure of the listener to remain undetected. A failure of the listener returns a different error, or causes the probe to time out.

The listener probe is started through `pmfadm` to make the probe highly available. If the probe is killed, PMF automatically restarts the probe.

If a problem occurs with the listener during a probe, the probe tries to restart the listener. The value that is set for the resource property `retry_count` determines the maximum number of times that the probe attempts the restart. If, after trying for the maximum number of times, the probe is still unsuccessful, the probe stops the fault monitor and does not switch over the resource group.

Obtaining Core Files for Troubleshooting DBMS Timeouts

To facilitate troubleshooting of unexplained DBMS timeouts, you can enable the fault monitor to create a core file when a probe timeout occurs. The contents of the core file relate to the fault monitor process. The fault monitor creates the core file in the root (`/`) directory. To enable the fault monitor to create a core file, use the `coreadm` command to enable set-id core dumps. For more information, see the [coreadm\(1M\)](#) man page.

Customizing the HA for Oracle Server Fault Monitor

Customizing the HA for Oracle server fault monitor enables you to modify the behavior of the server fault monitor as follows:

- Overriding the preset action for an error
- Specifying an action for an error for which no action is preset



Caution – Before you customize the HA for Oracle server fault monitor, consider the effects of your customizations, especially if you change an action from restart or switch over to ignore or stop monitoring. If errors remain uncorrected for long periods, the errors might cause problems with the database. If you encounter problems with the database after customizing the HA for Oracle server fault monitor, revert to using the preset actions. Reverting to the preset actions enables you to determine if the problem is caused by your customizations.

Customizing the HA for Oracle server fault monitor involves the following activities:

1. [Defining custom behavior for errors](#)
2. [Propagating a custom action file to all nodes in a cluster](#)
3. [Specifying the custom action file that a server fault monitor should use](#)

Defining Custom Behavior for Errors

The HA for Oracle server fault monitor detects the following types of errors:

- DBMS errors that occur during a probe of the database by the server fault monitor
- Alerts that Oracle logs in the alert log file
- Timeouts that result from a failure to receive a response within the time that is set by the `Probe_timeout` extension property

To define custom behavior for these types of errors, create a custom action file. This section contains the following information about custom action files:

- [“Custom Action File Format” on page 54](#)
- [“Changing the Response to a DBMS Error” on page 57](#)
- [“Changing the Response to Logged Alerts” on page 58](#)
- [“Changing the Maximum Number of Consecutive Timed-Out Probes” on page 60](#)

Custom Action File Format

A custom action file is a plain text file. The file contains one or more entries that define the custom behavior of the HA for Oracle server fault monitor. Each entry defines the custom behavior for a single DBMS error, a single timeout error, or several logged alerts. A maximum of 1024 entries is allowed in a custom action file.

Note – Each entry in a custom action file overrides the preset action for an error, or specifies an action for an error for which no action is preset. Create entries in a custom action file *only* for the preset actions that you are overriding or for errors for which no action is preset. Do *not* create entries for actions that you are not changing.

An entry in a custom action file consists of a sequence of keyword-value pairs that are separated by semicolons. Each entry is enclosed in braces.

The format of an entry in a custom action file is as follows:

```
{
[ERROR_TYPE=DBMS_ERROR|SCAN_LOG|TIMEOUT_ERROR;]
ERROR=error-spec;
[ACTION=SWITCH|RESTART|STOP|NONE;]
[CONNECTION_STATE=co|di|on|*;]
[NEW_STATE=co|di|on|*;]
[MESSAGE="message-string"]
}
```

White space may be used between separated keyword-value pairs and between entries to format the file.

The meaning and permitted values of the keywords in a custom action file are as follows:

ERROR_TYPE

Indicates the type of the error that the server fault monitor has detected. The following values are permitted for this keyword:

- DBMS_ERROR Specifies that the error is a DBMS error.
- SCAN_LOG Specifies that the error is an alert that is logged in the alert log file.
- TIMEOUT_ERROR Specifies that the error is a timeout.

The ERROR_TYPE keyword is optional. If you omit this keyword, the error is assumed to be a DBMS error.

ERROR

Identifies the error. The data type and the meaning of *error-spec* are determined by the value of the ERROR_TYPE keyword as shown in the following table.

ERROR_TYPE	Data Type	Meaning
DBMS_ERROR	Integer	The error number of a DBMS error that is generated by Oracle
SCAN_LOG	Quoted regular expression	A string in an error message that Oracle has logged to the Oracle alert log file
TIMEOUT_ERROR	Integer	The number of consecutive timed-out probes since the server fault monitor was last started or restarted

You must specify the ERROR keyword. If you omit this keyword, the entry in the custom action file is ignored.

ACTION

Specifies the action that the server fault monitor is to perform in response to the error. The following values are permitted for this keyword:

- NONE Specifies that the server fault monitor ignores the error.
- STOP Specifies that the server fault monitor is stopped.
- RESTART Specifies that the server fault monitor stops and restarts the entity that is specified by the value of the Restart_type extension property of the SUNW.oracle_server resource.
- SWITCH Specifies that the server fault monitor switches over the database server resource group to another node.

The **ACTION** keyword is optional. If you omit this keyword, the server fault monitor ignores the error.

CONNECTION_STATE

Specifies the required state of the connection between the database and the server fault monitor when the error is detected. The entry applies only if the connection is in the required state when the error is detected. The following values are permitted for this keyword:

- * Specifies that the entry always applies, regardless of the state of the connection.
- co Specifies that the entry applies only if the server fault monitor is attempting to connect to the database.
- on Specifies that the entry applies only if the server fault monitor is online. The server fault monitor is online if it is connected to the database.
- di Specifies that the entry applies only if the server fault monitor is disconnecting from the database.

The **CONNECTION_STATE** keyword is optional. If you omit this keyword, the entry always applies, regardless of the state of the connection.

NEW_STATE

Specifies the state of the connection between the database and the server fault monitor that the server fault monitor must attain after the error is detected. The following values are permitted for this keyword:

- * Specifies that the state of the connection must remain unchanged.
- co Specifies that the server fault monitor must disconnect from the database and reconnect immediately to the database.
- di Specifies that the server fault monitor must disconnect from the database. The server fault monitor reconnects when it next probes the database.

The **NEW_STATE** keyword is optional. If you omit this keyword, the state of the database connection remains unchanged after the error is detected.

MESSAGE

Specifies an additional message that is printed to the resource's log file when this error is detected. The message must be enclosed in double quotes. This message is additional to the standard message that is defined for the error.

The **MESSAGE** keyword is optional. If you omit this keyword, no additional message is printed to the resource's log file when this error is detected.

Changing the Response to a DBMS Error

The action that the server fault monitor performs in response to each DBMS error is preset as listed in [Table B-1](#). To determine whether you need to change the response to a DBMS error, consider the effect of DBMS errors on your database to determine if the preset actions are appropriate. For examples, see the subsections that follow:

- “Responding to an Error Whose Effects Are Major” on page 57
- “Ignoring an Error Whose Effects Are Minor” on page 58

To change the response to a DBMS error, create an entry in a custom action file in which the keywords are set as follows:

- `ERROR_TYPE` is set to `DBMS_ERROR`.
- `ERROR` is set to the error number of the DBMS error.
- `ACTION` is set to the action that you require.

Responding to an Error Whose Effects Are Major

If an error that the server fault monitor ignores affects more than one session, action by the server fault monitor might be required to prevent a loss of service.

For example, no action is preset for Oracle error 4031: unable to allocate *num-bytes* bytes of shared memory. However, this Oracle error indicates that the shared global area (SGA) has insufficient memory, is badly fragmented, or both states apply. If this error affects only a single session, ignoring the error might be appropriate. However, if this error affects more than one session, consider specifying that the server fault monitor restart the database.

The following example shows an entry in a custom action file for changing the response to a DBMS error to restart.

EXAMPLE 1-3 Changing the Response to a DBMS Error to Restart

```
{
ERROR_TYPE=DBMS_ERROR;
ERROR=4031;
ACTION=restart;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="Insufficient memory in shared pool.";
}
```

This example shows an entry in a custom action file that overrides the preset action for DBMS error 4031. This entry specifies the following behavior:

- In response to DBMS error 4031, the action that the server fault monitor performs is restart.
- This entry applies regardless of the state of the connection between the database and the server fault monitor when the error is detected.
- The state of the connection between the database and the server fault monitor must remain unchanged after the error is detected.

EXAMPLE 1-3 Changing the Response to a DBMS Error to Restart (Continued)

- The following message is printed to the resource's log file when this error is detected:
Insufficient memory in shared pool.

Ignoring an Error Whose Effects Are Minor

If the effects of an error to which the server fault monitor responds are minor, ignoring the error might be less disruptive than responding to the error.

For example, the preset action for Oracle error 4030: out of process memory when trying to allocate *num-bytes* bytes is restart. This Oracle error indicates that the server fault monitor could not allocate private heap memory. One possible cause of this error is that insufficient memory is available to the operating system. If this error affects more than one session, restarting the database might be appropriate. However, this error might not affect other sessions because these sessions do not require further private memory. In this situation, consider specifying that the server fault monitor ignore the error.

The following example shows an entry in a custom action file for ignoring a DBMS error.

EXAMPLE 1-4 Ignoring a DBMS Error

```
{
ERROR_TYPE=DBMS_ERROR;
ERROR=4030;
ACTION=none;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="";
}
```

This example shows an entry in a custom action file that overrides the preset action for DBMS error 4030. This entry specifies the following behavior:

- The server fault monitor ignores DBMS error 4030.
- This entry applies regardless of the state of the connection between the database and the server fault monitor when the error is detected.
- The state of the connection between the database and the server fault monitor must remain unchanged after the error is detected.
- No additional message is printed to the resource's log file when this error is detected.

Changing the Response to Logged Alerts

The Oracle software logs alerts in a file that is identified by the `alert_log_file` extension property. The server fault monitor scans this file and performs actions in response to alerts for which an action is defined.

Logged alerts for which an action is preset are listed in [Table B-2](#). Change the response to logged alerts to change the preset action, or to define new alerts to which the server fault monitor responds.

To change the response to logged alerts, create an entry in a custom action file in which the keywords are set as follows:

- `ERROR_TYPE` is set to `SCAN_LOG`.
- `ERROR` is set to a quoted regular expression that identifies a string in an error message that Oracle has logged to the Oracle alert log file.
- `ACTION` is set to the action that you require.

The server fault monitor processes the entries in a custom action file in the order in which the entries occur. Only the first entry that matches a logged alert is processed. Later entries that match are ignored. If you are using regular expressions to specify actions for several logged alerts, ensure that more specific entries occur before more general entries. Specific entries that occur after general entries might be ignored.

For example, a custom action file might define different actions for errors that are identified by the regular expressions `ORA-65` and `ORA-6`. To ensure that the entry that contains the regular expression `ORA-65` is not ignored, ensure that this entry occurs before the entry that contains the regular expression `ORA-6`.

The following example shows an entry in a custom action file for changing the response to a logged alert.

EXAMPLE 1-5 Changing the Response to a Logged Alert

```
{
ERROR_TYPE=SCAN_LOG;
ERROR="ORA-00600: internal error";
ACTION=RESTART;
}
```

This example shows an entry in a custom action file that overrides the preset action for logged alerts about internal errors. This entry specifies the following behavior:

- In response to logged alerts that contain the text `ORA-00600: internal error`, the action that the server fault monitor performs is restart.
- This entry applies regardless of the state of the connection between the database and the server fault monitor when the error is detected.
- The state of the connection between the database and the server fault monitor must remain unchanged after the error is detected.
- No additional message is printed to the resource's log file when this error is detected.

Changing the Maximum Number of Consecutive Timed-Out Probes

By default, the server fault monitor restarts the database after the second consecutive timed-out probe. If the database is lightly loaded, two consecutive timed-out probes should be sufficient to indicate that the database is hanging. However, during periods of heavy load, a server fault monitor probe might time out even if the database is functioning correctly. To prevent the server fault monitor from restarting the database unnecessarily, increase the maximum number of consecutive timed-out probes.



Caution – Increasing the maximum number of consecutive timed-out probes increases the time that is required to detect that the database is hanging.

To change the maximum number of consecutive timed-out probes allowed, create one entry in a custom action file for each consecutive timed-out probe that is allowed *except* the first timed-out probe.

Note – You are not required to create an entry for the first timed-out probe. The action that the server fault monitor performs in response to the first timed-out probe is preset.

For the last allowed timed-out probe, create an entry in which the keywords are set as follows:

- `ERROR_TYPE` is set to `TIMEOUT_ERROR`.
- `ERROR` is set to the maximum number of consecutive timed-out probes that are allowed.
- `ACTION` is set to `RESTART`.

For each remaining consecutive timed-out probe except the first timed-out probe, create an entry in which the keywords are set as follows:

- `ERROR_TYPE` is set to `TIMEOUT_ERROR`.
- `ERROR` is set to the sequence number of the timed-out probe. For example, for the second consecutive timed-out probe, set this keyword to 2. For the third consecutive timed-out probe, set this keyword to 3.
- `ACTION` is set to `NONE`.

Tip – To facilitate debugging, specify a message that indicates the sequence number of the timed-out probe.

The following example shows the entries in a custom action file for increasing the maximum number of consecutive timed-out probes to five.

EXAMPLE 1-6 Changing the Maximum Number of Consecutive Timed-Out Probes

```

{
ERROR_TYPE=TIMEOUT;
ERROR=2;
ACTION=NONE;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="Timeout #2 has occurred.";
}

{
ERROR_TYPE=TIMEOUT;
ERROR=3;
ACTION=NONE;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="Timeout #3 has occurred.";
}

{
ERROR_TYPE=TIMEOUT;
ERROR=4;
ACTION=NONE;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="Timeout #4 has occurred.";
}

{
ERROR_TYPE=TIMEOUT;
ERROR=5;
ACTION=RESTART;
CONNECTION_STATE=*;
NEW_STATE=*;
MESSAGE="Timeout #5 has occurred. Restarting.";
}

```

This example shows the entries in a custom action file for increasing the maximum number of consecutive timed-out probes to five. These entries specify the following behavior:

- The server fault monitor ignores the second consecutive timed-out probe through the fourth consecutive timed-out probe.
- In response to the fifth consecutive timed-out probe, the action that the server fault monitor performs is restart.
- The entries apply regardless of the state of the connection between the database and the server fault monitor when the timeout occurs.
- The state of the connection between the database and the server fault monitor must remain unchanged after the timeout occurs.
- When the second consecutive timed-out probe through the fourth consecutive timed-out probe occurs, a message of the following form is printed to the resource's log file:

Timeout *#number* has occurred.

EXAMPLE 1-6 Changing the Maximum Number of Consecutive Timed-Out Probes *(Continued)*

- When the fifth consecutive timed-out probe occurs, the following message is printed to the resource's log file:

```
Timeout #5 has occurred. Restarting.
```

Propagating a Custom Action File to All Nodes in a Cluster

A server fault monitor must behave consistently on all cluster nodes. Therefore, the custom action file that the server fault monitor uses must be identical on all cluster nodes. After creating or modifying a custom action file, ensure that this file is identical on all cluster nodes by propagating the file to all cluster nodes. To propagate the file to all cluster nodes, use the method that is most appropriate for your cluster configuration:

- Locating the file on a file system that all nodes share
- Locating the file on a highly available local file system
- Copying the file to the local file system of each cluster node by using operating system commands such as the `rcp` command or the `rdist` command.

Specifying the Custom Action File That a Server Fault Monitor Should Use

To apply customized actions to a server fault monitor, you must specify the custom action file that the fault monitor should use. Customized actions are applied to a server fault monitor when the server fault monitor reads a custom action file. A server fault monitor reads a custom action file when the you specify the file.

Specifying a custom action file also validates the file. If the file contains syntax errors, an error message is displayed. Therefore, after modifying a custom action file, specify the file again to validate the file.



Caution – If syntax errors in a modified custom action file are detected, correct the errors before the fault monitor is restarted. If the syntax errors remain uncorrected when the fault monitor is restarted, the fault monitor reads the erroneous file, ignoring entries that occur after the first syntax error.

▼ How to Specify the Custom Action File That a Server Fault Monitor Should Use

- 1 On a cluster node, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Set the `custom_action_file` extension property of the `SUNW.oracle_server` resource.

Set this property to the absolute path of the custom action file.

```
# clresource set -p custom_action_file=filepath server-resource
```

```
-p custom_action_file=filepath
```

Specifies the absolute path of the custom action file.

server-resource

Specifies the `SUNW.oracle_server` resource.

Changing the Role of an Oracle Data Guard Instance

Database role failover or switchover is possible between an Oracle primary database and an Oracle standby database. When you use Oracle commands to change the role of Oracle Data Guard instances, the changes are not propagated to the Oracle Solaris Cluster resources that represent these instances. Therefore, you must also use Oracle Solaris Cluster commands to change extension properties of these resources to ensure that database instances are started in the correct role.

▼ How to Change the Role of an Oracle Data Guard Instance

- 1 Prevent Oracle Solaris Cluster from starting the instance in an incorrect role.

If a node fails while you are changing the role of an Oracle Data Guard instance, Oracle Solaris Cluster might restart the instance in an incorrect role. To prevent this possibility, change the `Dataguard_role` extension property of the Oracle server resource that represents the instance to `IN_TRANSITION`.

```
# clresource set -p Dataguard_role=IN_TRANSITION server-rs
```

- 2 Perform the required operations on the Oracle database to convert the database to a new role.

- 3 Change the following extension properties of the Oracle server resource that represents the instance to reflect the new role of the instance:

- `Dataguard_role`

- Standby_mode

The required combination of `Dataguard_role` and `Standby_mode` depends on the change of role, as follows:

- **To change from a primary database to a physical standby database, run the following command:**

```
# clresource set -p Dataguard_role=STANDBY -p Standby_mode=PHYSICAL server-rs
```

- **To change from a primary database to a logical standby database, run the following command:**

```
# clresource set -p Dataguard_role=STANDBY \  
-p Standby_mode=LOGICAL server-rs
```

- **To change from a standby database to a primary database, run the following command:**

```
# clresource set -p Dataguard_role=PRIMARY server-rs
```

- **To change from a physical standby database to a snapshot standby database, run the following command:**

```
# clresource set -p Standby_mode=SNAPSHOT server-rs
```

HA for Oracle Extension Properties

The extension properties that you can set for each HA for Oracle resource type are listed in the following sections:

- “[SUNW.oracle_server Extension Properties](#)” on page 65
- “[SUNW.oracle_listener Extension Properties](#)” on page 69

See the [r_properties\(5\)](#) man page and the [rg_properties\(5\)](#) man page for details about all of the system-defined properties.

SUNW.oracle_server Extension Properties

Auto_End_Bkp (Boolean)

Specifies whether the following recovery actions are performed if an Oracle relational database management system (RDBMS) hot backup is interrupted.

- Recognizing when a database fails to open because of files that remain in hot backup mode. This verification process occurs when HA for Oracle starts.
- Identifying and releasing all files that remain in hot backup mode.
- Opening the database for use.

The permitted values for this property are as follows:

- `False` – Specifies that the recovery actions are *not* performed. This value is the default.
- `True` – Specifies that the recovery actions are performed.

Default: False

Range: None

Tunable: Any time

Connect_cycle (integer)

The number of probe cycles that the server fault monitor performs before disconnecting from the database.

Default: 5

Range: 0 – 99,999

Tunable: Any time

Custom_action_file (string)

The absolute path of the file that defines the custom behavior of the HA for Oracle server fault monitor.

Default: ""

Range: None

Tunable: Any time

Dataguard_role (string)

The role of the database. The permitted values for this property are as follows:

NONE	Specifies that no standby database instances configured for the database instance
PRIMARY	Specifies that the database is a primary database instance for which standby database instances are configured
STANDBY	Specifies that the database role is standby
IN_TRANSITION	Specifies that the database is undergoing a role reversal process

Default: NONE

Range: None

Tunable: Any time

Db_unique_name (string)

The unique name of the single-instance Oracle database that is being deployed.

Default: NONE

Range: None

Tunable: Any time

Debug_level (integer)

The level to which debug messages from the Oracle server component are logged. When the debug level is increased, more debug messages are written to the log files. These messages are

logged to the file `/var/opt/SUNWscor/oracle_server/message_log.rs`, where *rs* is the name of the resource that represents the Oracle server component.

Default: 1, which logs `syslog` messages

Range: 0– 100

Tunable: Any time

`Oracle_home` (string)

The path to the Oracle home directory.

Default: None

Range: Minimum = 1

Tunable: When disabled

`Oracle_sid` (string)

The Oracle system identifier.

Default: None

Range: Minimum = 1

Tunable: When disabled

`Parameter_file` (string)

The Oracle parameter file. If the Oracle parameter file is not specified, this property defaults to Oracle's default.

Default: ""

Range: Minimum = 0

Tunable: Any time

`Probe_timeout` (integer)

The timeout value (in seconds) that the server fault monitor uses to probe an Oracle server instance.

Default: 300

Range: 0– 99,999

Tunable: Any time

`Restart_type` (string)

Specifies the entity that the server fault monitor restarts when the response to a fault is restart. The permitted values for this property are as follows:

`RESOURCE_RESTART` Specifies that only this resource is restarted

RESOURCE_GROUP_RESTART Specifies that all resources in the resource group that contains this resource are restarted

Default: RESOURCE_RESTART

Range: None

Tunable: Any time

Standby_mode (string)

The mode of the standby database. The permitted values for this property are as follows:

LOGICAL Specifies a logical standby database

PHYSICAL Specifies a physical standby database

SNAPSHOT Specifies a snapshot standby database

Default: LOGICAL

Range: None

Tunable: Any time

User_env (string)

A file that contains environment variables to be set before server startup and shutdown. Those environment variables that have values that differ from Oracle defaults must be defined in this file.

For example, a user's `listener.ora` file might not reside under the `/var/opt/oracle` directory or the `$ORACLE_HOME/network/admin.` directory. In this situation, the `TNS_ADMIN` environment variable should be defined.

The definition of each environment variable that is defined must follow the format `VARIABLE_NAME=VARIABLE_VALUE`. Each of these environment variables must be specified, one per line in the environment file.

Default: NULL

Range: None

Tunable: Any time

Wait_for_online (Boolean)

Wait in the `START` method until the database is online.

Default: True

Range: None

Tunable: Any time

SUNW.oracle_listener Extension Properties

Listener_name (string)

The name of the Oracle listener. This name must match the corresponding entry in the listener.ora configuration file.

Default: LISTENER

Range: Not applicable

Tunable: When disabled

Oracle_home (string)

The path to the Oracle home directory.

Default: No default defined

Range: Not applicable

Tunable: When disabled

Probe_timeout (integer)

The timeout value in seconds that the fault monitor uses to probe an Oracle listener.

Default: 180

Range: 1– 99,999

Tunable: Any time

User_env (string)

A file that contains environment variables to be set before listener startup and shutdown. Those environment variables that have values that differ from Oracle defaults must be defined in this file.

For example, a user's listener.ora file might not reside under the /var/opt/oracle directory or the \$ORACLE_HOME/network/admin. directory. In this situation, the TNS_ADMIN environment variable should be defined.

The definition of each environment variable that is defined must follow the format VARIABLE_NAME=VARIABLE_VALUE. Each of these environment variables must be specified, one per line in the environment file.

Default: ""

Range: Not applicable

Tunable: Any time

Preset Actions for DBMS Errors and Logged Alerts

Preset actions for DBMS errors and logged alerts are listed as follows:

- DBMS errors for which an action is preset are listed in [Table B-1](#).
- Logged alerts for which an action is preset are listed in [Table B-2](#).

TABLE B-1 Preset Actions for DBMS Errors

Error Number	Action	Connection State	New State	Message
18	NONE	co	di	Max. number of DBMS sessions exceeded
20	NONE	co	di	Max. number of DBMS processes exceeded
28	NONE	on	di	Session killed by DBA, will reconnect
50	SWITCH	*	di	O/S error occurred while obtaining an enqueue. See o/s error.
51	NONE	*	di	timeout occurred while waiting for resource
55	NONE	*	*	maximum number of DML locks in DBMS exceeded
62	STOP	*	di	Need to set DML_LOCKS in init.ora file to value other than 0
107	RESTART	*	di	failed to connect to ORACLE listener process
257	NONE	*	di	archiver error. Connect internal only, until freed.
290	SWITCH	*	di	Operating system archival error occurred. Check alert log.
447	SWITCH	*	di	fatal error in background process
448	RESTART	*	di	normal completion of background process
449	RESTART	*	di	background process '%s' unexpectedly terminated with error %s
470	SWITCH	*	di	Oracle background process died
471	SWITCH	*	di	Oracle background process died

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
472	SWITCH	*	di	Oracle background process died
473	SWITCH	*	di	Oracle background process died
474	RESTART	*	di	SMON died, warm start required
475	SWITCH	*	di	Oracle background process died
476	SWITCH	*	di	Oracle background process died
477	SWITCH	*	di	Oracle background process died
480	RESTART	*	di	LCK* process terminated with error
481	RESTART	*	di	LMON process terminated with error
482	RESTART	*	di	LMD* process terminated with error
602	SWITCH	*	di	internal programming exception
604	NONE	on	di	Recursive error
705	RESTART	*	di	inconsistent state during start up
942	NONE	on	*	Warning - V\$SYSSTAT not accessible - check grant on V_\$SYSSTAT
1001	NONE	on	di	Lost connection to database
1002	NONE	on	*	Internal error in HA-DBMS Oracle
1003	NONE	on	di	Resetting database connection
1012	NONE	on	di	Not logged on
1012	RESTART	di	co	Not logged on
1014	NONE	*	*	ORACLE shutdown in progress
1017	STOP	*	*	Please correct login information in HA-DBMS Oracle database configuration
1031	NONE	on	*	Insufficient privileges to perform DBMS operations - check Oracle user privileges
1033	NONE	co	co	Oracle is in the shutdown or initialization process
1033	NONE	*	di	Oracle is in the shutdown or initialization process
1034	RESTART	co	co	Oracle is not available
1034	RESTART	di	co	Oracle is not available
1034	NONE	on	di	Oracle is not available

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
1035	RESTART	co	co	Access restricted - restarting database to reset
1041	NONE	on	di	
1041	NONE	di	co	
1045	NONE	co	*	Fault monitor user lacks CREATE SESSION privilege logon denied.
1046	RESTART	*	di	cannot acquire space to extend context area
1050	RESTART	*	di	cannot acquire space to open context area
1053	SWITCH	*	*	user storage address cannot be read or written
1054	SWITCH	*	*	user storage address cannot be read or written
1075	NONE	co	on	Already logged on
1089	NONE	on	di	immediate shutdown in progress
1089	NONE	*	*	Investigate! Could be hanging!
1090	NONE	*	di	shutdown in progress - connection is not permitted
1092	NONE	*	di	ORACLE instance terminated. Disconnection forced
1513	SWITCH	*	*	invalid current time returned by operating system
1542	NONE	on	*	table space is off-line - please correct!
1552	NONE	on	*	rollback segment is off-line - please correct!
1950	NONE	on	*	Insufficient privileges to perform DBMS operations - check Oracle user privileges
2701	STOP	*	*	HA-DBMS Oracle error - ORACLE_HOME did not get set!
2703	RESTART	*	di	
2704	RESTART	*	di	
2709	RESTART	*	di	
2710	RESTART	*	di	
2719	RESTART	*	di	
2721	RESTART	*	*	
2726	STOP	*	*	Could not locate ORACLE executables - check ORACLE_HOME setting
2735	RESTART	*	*	osnfp: cannot create shared memory segment
2811	SWITCH	*	*	Unable to attach shared memory segment

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
2839	SWITCH	*	*	Sync of blocks to disk failed.
2840	SWITCH	*	*	
2846	SWITCH	*	*	
2847	SWITCH	*	*	
2849	SWITCH	*	*	
2842	RESTART	*	*	Client unable to fork a server - Out of memory
3113	RESTART	co	di	lost connection
3113	NONE	on	di	lost connection
3113	NONE	di	di	lost connection
3114	NONE	*	co	Not connected?
4030	RESTART	*	*	
4032	RESTART	*	*	
4100	RESTART	*	*	communication area cannot be allocated insufficient memory
6108	STOP	co	*	Can't connect to remote database - make sure SQL*Net server is up
6114	STOP	co	*	Can't connect to remote database - check SQL*Net configuration
7205	SWITCH	*	di	
7206	SWITCH	*	di	
7208	SWITCH	*	di	
7210	SWITCH	*	di	
7211	SWITCH	*	di	
7212	SWITCH	*	di	
7213	SWITCH	*	di	
7214	SWITCH	*	di	
7215	SWITCH	*	di	
7216	SWITCH	*	di	
7218	SWITCH	*	di	
7219	RESTART	*	*	slspool: unable to allocate spooler argument buffer.

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
7223	RESTART	*	*	slspool: fork error, unable to spawn spool process. - Resource limit reached
7224	SWITCH	*	*	
7229	SWITCH	*	*	
7232	SWITCH	*	*	
7234	SWITCH	*	*	
7238	SWITCH	*	*	slemcl: close error.
7250	RESTART	*	*	
7251	RESTART	*	*	
7252	RESTART	*	*	
7253	RESTART	*	*	
7258	RESTART	*	*	
7259	RESTART	*	*	
7263	SWITCH	*	*	
7269	SWITCH	*	*	
7279	SWITCH	*	*	
7280	RESTART	*	*	
7296	SWITCH	*	*	
7297	SWITCH	*	*	
7306	RESTART	*	*	
7310	SWITCH	*	*	
7315	SWITCH	*	*	
7321	SWITCH	*	*	
7322	SWITCH	*	*	
7324	RESTART	*	*	
7325	RESTART	*	*	
7351	SWITCH	*	*	
7361	RESTART	*	*	

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
7404	SWITCH	*	*	
7414	RESTART	*	*	
7415	RESTART	*	*	
7417	SWITCH	*	*	
7418	SWITCH	*	*	
7419	SWITCH	*	*	
7430	SWITCH	*	*	
7455	SWITCH	*	*	
7456	SWITCH	*	*	
7466	SWITCH	*	*	
7470	SWITCH	*	*	
7475	SWITCH	*	*	
7476	SWITCH	*	*	
7477	SWITCH	*	*	
7478	SWITCH	*	*	
7479	SWITCH	*	*	
7481	SWITCH	*	*	
9706	SWITCH	*	*	
9716	SWITCH	*	*	
9718	RESTART	*	*	
9740	SWITCH	*	*	
9748	SWITCH	*	*	
9747	RESTART	*	*	
9749	RESTART	*	*	
9751	RESTART	*	*	
9755	RESTART	*	*	
9757	RESTART	*	*	

TABLE B-1 Preset Actions for DBMS Errors (Continued)

Error Number	Action	Connection State	New State	Message
9756	SWITCH	*	*	
9758	SWITCH	*	*	
9761	RESTART	*	*	
9765	RESTART	*	*	
9779	RESTART	*	*	
9829	RESTART	*	*	
9831	SWITCH	*	*	
9834	SWITCH	*	*	
9836	SWITCH	*	*	
9838	SWITCH	*	*	
9837	RESTART	*	*	
9844	RESTART	*	*	
9845	RESTART	*	*	
9846	RESTART	*	*	
9847	RESTART	*	*	
9853	SWITCH	*	*	
9854	SWITCH	*	*	
9856	RESTART	*	*	
9874	SWITCH	*	*	
9876	SWITCH	*	*	
9877	RESTART	*	*	
9878	RESTART	*	*	
9879	RESTART	*	*	
9885	RESTART	*	*	
9888	RESTART	*	*	
9894	RESTART	*	*	
9909	RESTART	*	*	

TABLE B-1 Preset Actions for DBMS Errors *(Continued)*

Error Number	Action	Connection State	New State	Message
9912	RESTART	*	*	
9913	RESTART	*	*	
9919	SWITCH	*	*	
9943	RESTART	*	*	
9947	RESTART	*	*	
9948	SWITCH	*	*	
9949	SWITCH	*	*	
9950	SWITCH	*	*	
12505	STOP	*	*	TNS:listener could not resolve SID given in connect descriptor.Check listener configuration file.
12541	STOP	*	*	TNS:no listener. Please verify connect_string property, listener and TNSconfiguration.
12545	SWITCH	*	*	Please check HA-Oracle parameters. Connect failed because target host or object does not exist
27100	STOP	*	*	Shared memory realm already exists
98765	STOP	*	*	Database role queried from database does not match the Oracle Solaris Cluster resource's dataguard role configuration.
99999	RESTART	*	di	Monitor detected death of Oracle background processes.

TABLE B-2 Preset Actions for Logged Alerts

Alert String	Action	Connection State	New State	Message
ORA-07265	SWITCH	*	di	Semaphore access problem
found dead multi-threaded server	NONE	*	*	Warning: Multi-threaded Oracle server process died (restarted automatically)
found dead dispatcher	NONE	*	*	Warning: Oracle dispatcher process died (restarted automatically)

Sample Configurations for Oracle ASM with HA for Oracle

This appendix contains diagrams that explain various sample configurations of Oracle Automatic Storage Management (Oracle ASM) with HA for Oracle. The diagrams in this section provide information about the dependencies between HA for Oracle resources and Oracle ASM services.

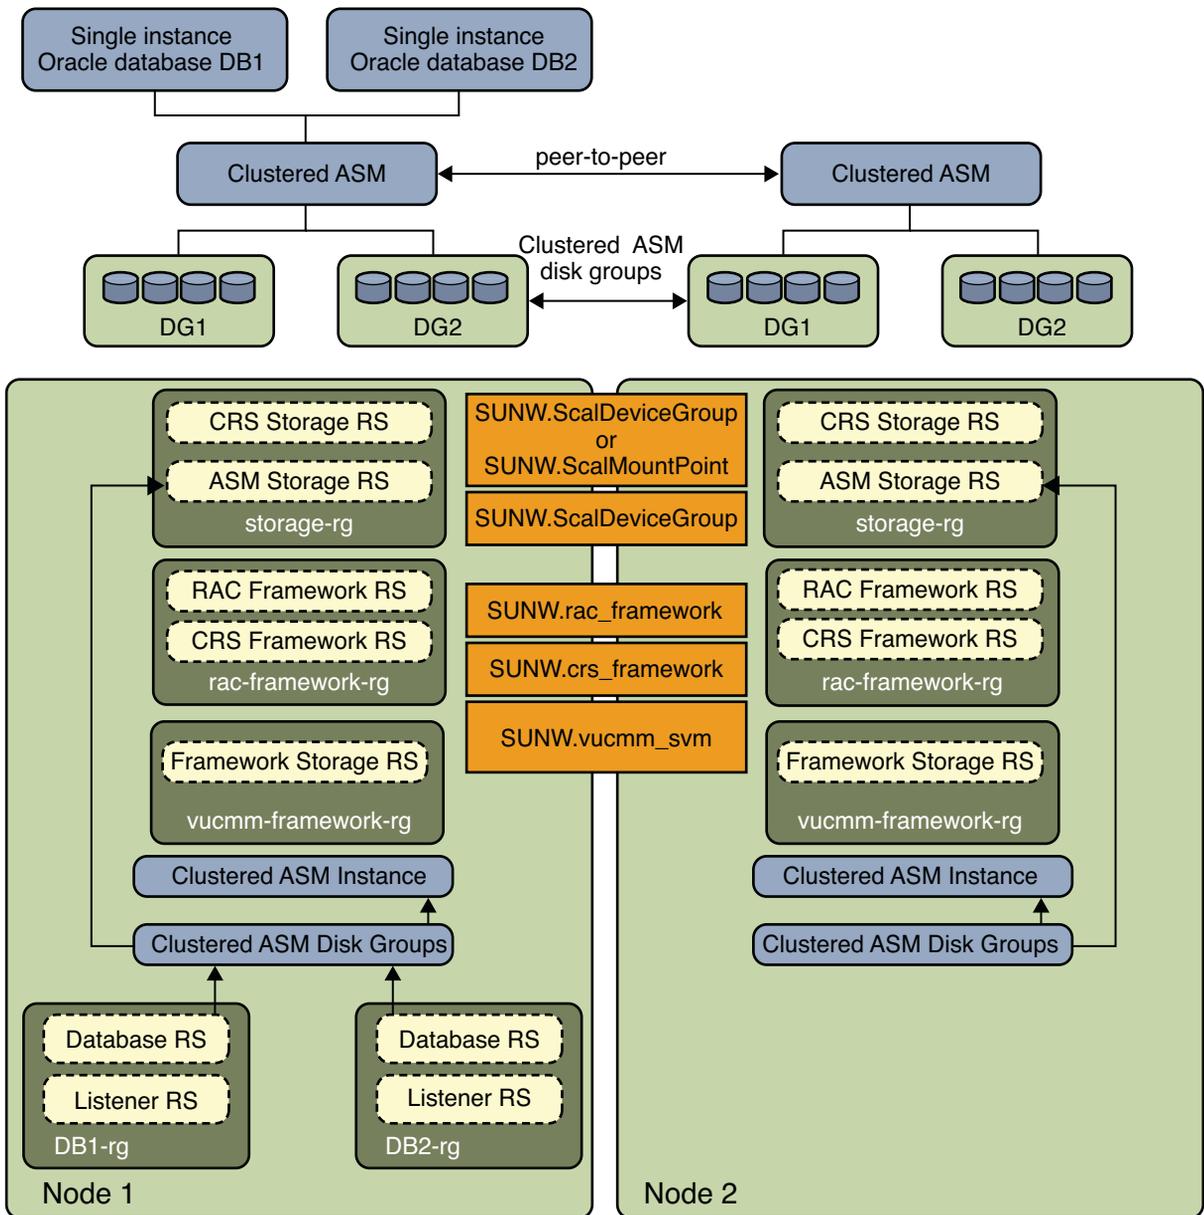
Clustered Oracle ASM with Clustered Disk Groups

This section contains sample configurations of single instance Oracle databases with clustered Oracle ASM instances and clustered Oracle ASM disk groups. There are two diagrams and the second diagram is a continuation of the first diagram.

The following diagram represents clustered Oracle ASM instances that serves two single instance Oracle databases, DB1 and DB2 on Node1. The databases DB1 or DB2 can use either one of the Oracle ASM disk groups DG1 and DG2 or both the disk groups as these are clustered Oracle ASM disk groups. The upper part of the diagram shows the relationships of the Oracle database instances with the clustered Oracle ASM instances on Node1 and Node2. The clustered Oracle ASM instances manage two Oracle ASM disk groups on both nodes at the same time. The lower part of the diagram represents the existing Oracle Solaris Cluster resource group and resources for single instance Oracle databases and their requirement for clustered Oracle ASM services.

The dotted box represents existing HA for Oracle resources with the new Oracle ASM resources. The arrows represent new dependencies between HA for Oracle and clustered Oracle ASM services.

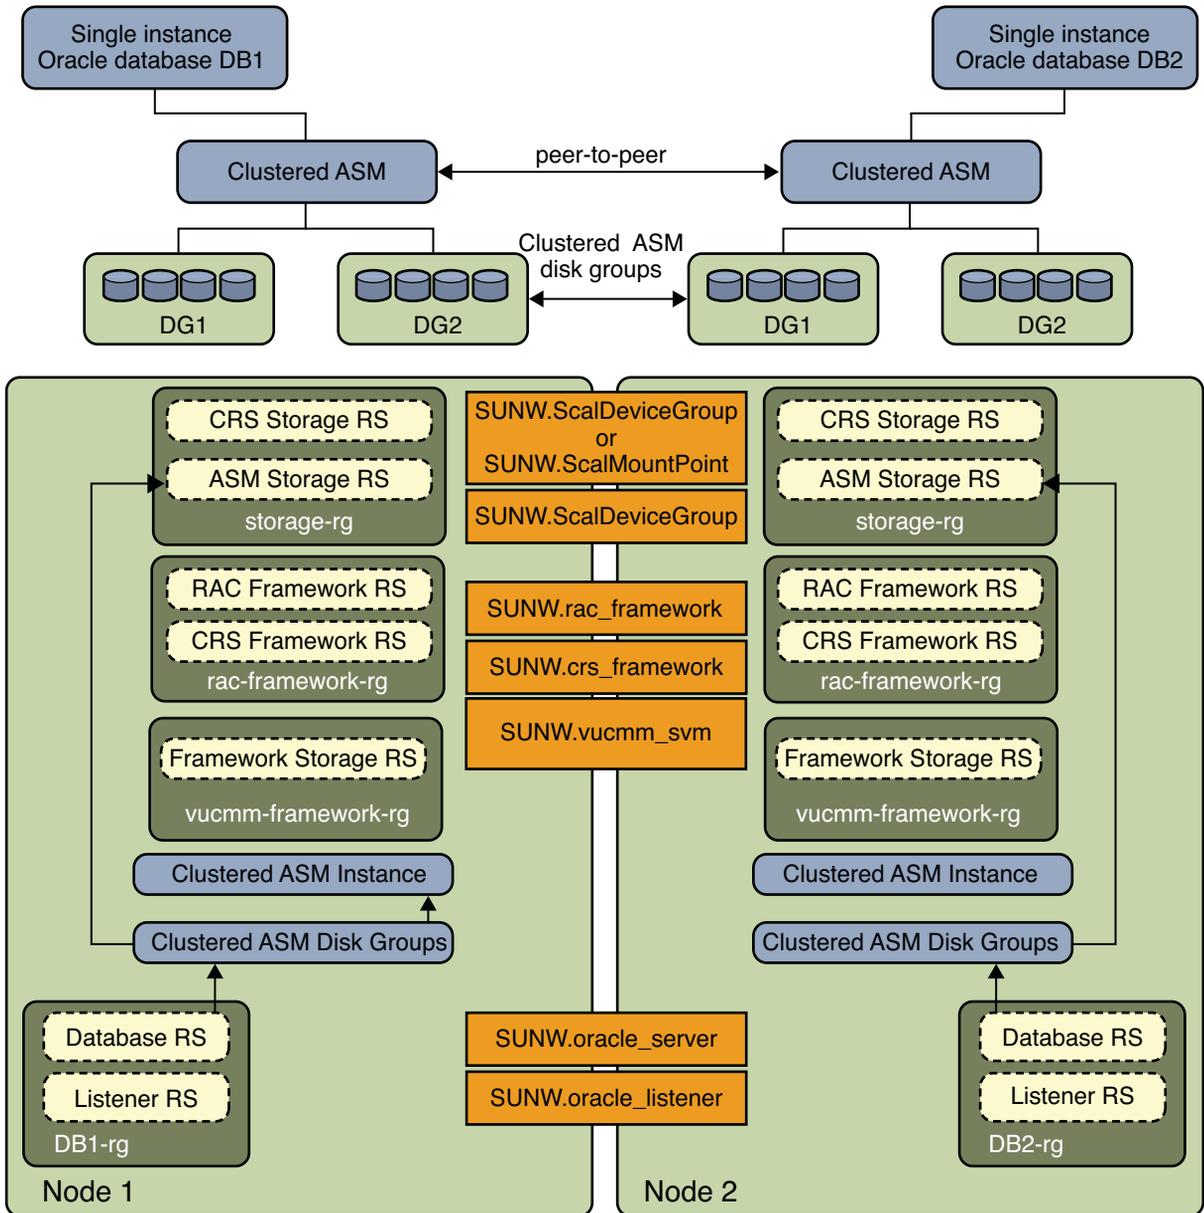
FIGURE C-1 Clustered Oracle ASM with Clustered Disk Groups [1]



In the following diagram, a single instance Oracle database DB2 that shares Oracle ASM disk groups DG1 and DG2 is now running on Node2 after a failover of resource groups DB2- rg to Node2. The upper part of the diagram shows the relationships of the Oracle database instances with the clustered Oracle ASM instances on Node1 and Node2. The clustered Oracle ASM

instances manage two Oracle ASM disk groups on both nodes at the same time. The lower part of the diagram represents the existing Oracle Solaris Cluster resource group and resources for single instance Oracle databases and their requirement for clustered Oracle ASM services. If the storage type is hardware RAID, the resource types `SUNW.ScaleDeviceGroup` and `SUNW.rac_svm` are not required.

FIGURE C-2 Clustered Oracle ASM with Clustered Disk Groups [2]



Index

A

action file, *See* custom action file

ACTION keyword, 55

actions

- listener fault monitor, 52–53

- preset for fault monitor, 71–78

- server fault monitor

 - definition, 51–52

 - modifying, 55

adding, messages to log files, 56

alert logs

- modifying response to errors, 58–59

- use of by fault monitor, 52

Auto_End_Bkp extension property, 65

C

caution notice, server fault monitor customizations, 53

changing, *See* modifying

configuring

- HA for Oracle

 - performing, 28–47

 - planning, 13–14

- Oracle database

 - using Oracle ASM, 17–18

 - using Solaris Volume Manager, 17

- Oracle Grid Infrastructure with clustered Oracle

 - ASM, 40–43

- SCAN listener, 18–19

Connect_cycle extension property, 65

CONNECTION_STATE keyword, 56

core files, fault monitors, 53

custom action file

- format, 54–56

- maximum number of entries in, 54

- order of entries in, 59

- propagating to cluster nodes, 62

- specifying, 62–63

- validating, 62

Custom_action_file extension property, 66

customizing, server fault monitor, 53–63

D

database instance, 36

database management system (DBMS)

- errors

 - modifying response to, 57–58

 - preset actions, 71–78

- timeouts, 53

database-related files, configuration requirements, 14

databases, Oracle, 22–23

Dataguard_role extension property, 36

DBMS (database management system)

- errors

 - modifying response to, 57–58

 - preset actions, 71–78

- timeouts, 53

Debug_level extension property, server, 66

E

- ERROR keyword, 55
- ERROR_TYPE keyword, 55
- errors
 - in custom action file, 62
 - DBMS
 - modifying response to, 57–58
 - preset actions for, 71–78
 - ignoring, 58
 - responding to, 57–58
 - SGA, 57
 - timeouts
 - core-file creation, 53
 - modifying maximum allowed, 60–62
 - types detected by fault monitor, 54
- extension properties
 - SUNW.oracle_listener resource type, 69
 - SUNW.oracle_server resource type, 65–68

F

- fault monitors
 - actions
 - listener fault monitor, 52–53
 - modifying, 55
 - server fault monitor, 51–52
 - alert logs, 52
 - caution notice, 53
 - core-file creation, 53
 - customizing, 53–63
 - error types detected by, 54
 - Oracle listener
 - resource type for, 49
 - Oracle server
 - resource type for, 49
 - preset actions, 71–78
 - tuning, 49–53
- files
 - alert log
 - modifying response to errors in, 58–59
 - use of by fault monitor, 52
 - core
 - fault monitors, 53

files (*Continued*)

- custom action
 - format, 54–56
 - order of entries in, 59
 - propagating to cluster nodes, 62
 - specifying, 62–63
 - validating, 62
- database, 14
- HA for Oracle logs
 - additional messages in, 56
 - location, 48–49
- Oracle application, 14
- fragmentation, of memory, 57

H

- HA for Oracle
 - See also* Oracle configuration
 - clustered Oracle ASM, 43–47
 - performing, 28–47
 - planning, 13–14
 - fault monitors, 49–53
 - installing, 27
 - overview of process, 11–13
 - planning, 13–14
 - log files
 - additional messages in, 56
 - location, 48–49
 - registering, 28–47
 - software package, installing, 27
 - SUNW.HASStoragePlus resource type, 34
 - verifying installation, 47–48
- heap memory, 58
 - help, 10
 - hot backup mode, 65

I

- ignoring, minor errors, 58
- installing
 - HA for Oracle, 27
 - planning, 13–14

installing (*Continued*)

Oracle ASM software, 19

insufficient memory errors, 57, 58

K

keywords, custom action file, 55

L

listener, extension properties, 69

listener fault monitor, 52–53

Listener_name extension property, 69

log files

HA for Oracle

additional messages in, 56

location, 48–49

logged alerts

modifying response to, 58–59

use of by fault monitor, 52

M

maximum values

entries in custom action file, 54

number of allowed timeouts, 60–62

memory

shortage of, 57, 58

MESSAGE keyword, 56

modifying

number of allowed timeouts, 60–62

Oracle Data Guard instance role, 63–64

response to DBMS errors, 57–58

response to logged alerts, 58–59

server fault monitor actions, 55

N

NEW_STATE keyword, 56

O

Oracle

application files, 14

clients, 48

database

configuring using Oracle ASM, 17–18

configuring using Solaris Volume Manager, 17

creating, 22–23

setting permissions, 23–27

error numbers, 71–78

preparing nodes for installation of, 15–17

verifying installation, 21–22

Oracle ASM

clustered configuration with Oracle Grid

Infrastructure, 40–43

configuration requirements, 14

installing, 19

software requirements, 13

Oracle Data Guard, 22–23

Dataguard_role extension property, 36

modifying instance role, 63–64

standby database modes, 37

Oracle Grid Infrastructure

configuring with clustered Oracle ASM, 40–43

software requirements, 13

Oracle_home extension property

listener, 69

server, 67

oracle_listener resource type, extension

properties, 69

oracle_server resource type, extension

properties, 65–68

Oracle_sid extension property, 67

Oracle Solaris, publisher, 27

Oracle Solaris Cluster, publisher, 27

order, entries in custom action file, 59

overriding, server fault monitor presets, 53–63

P

package, 27

Parameter_file extension property, 67

planning, HA for Oracle configuration, 13–14

preset actions, fault monitor, 71–78

preventing

- unnecessary restarts
 - for DBMS errors, 58
 - for timeouts, 60–62

primary database instances, 22–23

Probe_timeout extension property

- listener, 69
- server, 67

publisher

- Oracle Solaris, 27
- Oracle Solaris Cluster, 27

R

registering

- HA for Oracle
 - generally, 28–47
 - clustered Oracle ASM, 43–47

resource types

- fault monitors, 49
- SUNW.oracle_listener
 - extension properties, 69
- SUNW.oracle_server
 - extension properties, 65–68

responding, to major errors, 57–58

Restart_type extension property, 67

restarts

- preventing
 - for DBMS errors, 58
 - for timeouts, 60–62

restrictions, single instance Oracle ASM, 12

S

SCAN listener

- configuring, 18–19
- omitting SUNW.oracle_listener resource type, 34

server, extension properties, 65–68

server fault monitor

- actions
 - definition, 51–52
 - modifying, 55
- alert logs, 52

server fault monitor (*Continued*)

- caution notice, 53
- customizing, 53–63
- error types detected by, 54
- overview, 50–52
- preset actions, 71–78

sessions

- effect of errors on, 57, 58

SGA (shared global area), errors, 57

shared global area (SGA), errors, 57

software package, 27

software requirements, Oracle Grid Infrastructure, 13

sqlplus command, 23

standby database instances, 22–23, 37

Standby_mode extension property, 37, 68

SUNW.HASStoragePlus resource type, 34

SUNW.oracle_listener resource type

- extension properties, 69
- omitting for SCAN listener, 34

SUNW.oracle_server resource type, extension properties, 65–68

syntax errors, custom action file, 62

system properties, effect on fault monitors, 49

T

technical support, 10

timeouts

- core-file creation, 53
- modifying maximum allowed, 60–62

tuning, fault monitors, 49–53

U

User_env extension property

- listener, 69
- server, 68

V

validating, custom action file, 62

verifying

- HA for Oracle installation, 47–48

- Oracle installation, 21–22

W

- Wait_for_online extension property, 68

