Oracle® Solaris Cluster Data Service for Oracle Guide

ORACLE®

110714@25097

Contents

Tables

Preface

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

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit 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.

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	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your .login file.
		Use ls -a to list all files.
		machine_name% you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	machine_name% **su**
		Password:
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is rm *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*.
		A *cache* is a copy that is stored locally.
		Do *not* 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://docs.sun.com`.

Topic	Documentation
Data service administration	*Oracle Solaris Cluster Data Services Planning and Administration Guide*
	Individual data service guides
Concepts	*Oracle Solaris Cluster Concepts Guide*
Overview	*Oracle Solaris Cluster Overview*
Software installation	*Oracle Solaris Cluster Software Installation Guide*
System administration	*Oracle Solaris Cluster System Administration Guide*
Hardware administration	*Oracle Solaris Cluster 3.3 Hardware Administration Manual*
	Individual hardware administration guides
Data service development	*Oracle Solaris Cluster Data Services Developer's Guide*
Error messages	*Oracle Solaris Cluster Error Messages Guide*
Command and function reference	*Oracle Solaris Cluster Reference Manual*

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster at `http://docs.sun.com`.

Related Third-Party Web Site References

Third-party URLs that are referenced in this document provide additional related information.

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

See the following web sites for additional resources:

- Documentation (http://docs.sun.com)
- Support (http://www.oracle.com/us/support/systems/index.html)
- Training (http://education.oracle.com) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to http://docs.sun.com and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

Oracle Technology Network (http://www.oracle.com/technetwork/index.html) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the Discussion Forums (http://forums.oracle.com).
- Get hands-on step-by-step tutorials with Oracle By Example (http://www.oracle.com/technology/obe/start/index.html).
- Download Sample Code (http://www.oracle.com/technology/sample_code/index.html).

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 Oracle Solaris Operating System (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

Command	Function
prtconf -v	Displays the size of the system memory and reports information about peripheral devices
psrinfo -v	Displays information about processors
showrev –p	Reports which patches are installed
prtdiag -v	Displays system diagnostic information
/usr/cluster/bin/clnode show-rev	Displays Oracle Solaris Cluster release and package version information

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 HA for Oracle.

This chapter contains the following sections.

Note – You can use Oracle Solaris Cluster Manager to configure this data service. See the Oracle Solaris Cluster Manager online help for details.

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 Single Instance Oracle ASM
- 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   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 18
Prepare the nodes and disks	"Preparing the Nodes and Disks" on page 21
Install the Oracle software	"How to Install the Oracle Software" on page 26
Verify the Oracle installation	"How to Verify the Oracle Installation" on page 28
Create an Oracle database	"How to Create a Primary Oracle Database" on page 29
Set up Oracle database permissions	"How to Set Up Oracle Database Permissions" on page 30
Install the HA for Oracle packages	"Installing the HA for Oracle Packages" on page 33
Register and configure HA for Oracle	"How to Register and Configure HA for Oracle Without Oracle ASM Using Oracle Solaris Cluster Maintenance Commands" on page 41
Verify the HA for Oracle installation	"Verifying the HA for Oracle Installation" on page 59
Tune the HA for Oracle fault monitor	"Tuning the HA for Oracle Fault Monitors" on page 61
(Optional) Customize the HA for Oracle server fault monitor	"Customizing the HA for Oracle Server Fault Monitor" on page 65
(Optional) Upgrade HA for Oracle resource types	"Upgrading HA for Oracle Resource Types" on page 75
(Optional) Change the role of an Oracle Data Guard instance	"Changing the Role of an Oracle Data Guard Instance" on page 78

TABLE 2   Task Map: Installing and Configuring HA for Oracle with Single Instance Oracle ASM

Task	Cross-Reference
Plan the HA for Oracle installation and configuration	"Planning the HA for Oracle Installation and Configuration" on page 18

TABLE 2 Task Map: Installing and Configuring HA for Oracle with Single Instance Oracle ASM *(Continued)*

Task	Cross-Reference
Prepare the nodes and disks	"Preparing the Nodes and Disks" on page 21
Install the Oracle ASM software	"How to Install the Oracle ASM Software" on page 25
Install the Oracle software	"How to Install the Oracle Software" on page 26
Verify the Oracle installation	"How to Verify the Oracle Installation" on page 28
Create an Oracle database	"How to Create a Primary Oracle Database" on page 29
Set up Oracle database permissions	"How to Set Up Oracle Database Permissions" on page 30
Install the HA for Oracle packages	"Installing the HA for Oracle Packages" on page 33
Register and configure HA for Oracle	"How to Register and Configure HA for Oracle with Single Instance Oracle ASM Using Oracle Solaris Cluster Maintenance Commands" on page 48
Verify the HA for Oracle installation	"Verifying the HA for Oracle Installation" on page 59
Tune the HA for Oracle fault monitor	"Tuning the HA for Oracle Fault Monitors" on page 61
(Optional) Customize the HA for Oracle server fault monitor	"Customizing the HA for Oracle Server Fault Monitor" on page 65
(Optional) Upgrade HA for Oracle resource types	"Upgrading HA for Oracle Resource Types" on page 75
(Optional) Change the role of an Oracle Data Guard instance	"Changing the Role of an Oracle Data Guard Instance" on page 78

TABLE 3 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 18
Prepare the nodes and disks	"Preparing the Nodes and Disks" on page 21
Install the Oracle ASM software	"How to Install the Oracle ASM Software" on page 25
Install the Oracle software	"How to Install the Oracle Software" on page 26
Verify the Oracle installation	"How to Verify the Oracle Installation" on page 28
Create an Oracle database	"How to Create a Primary Oracle Database" on page 29
Set up Oracle database permissions	"How to Set Up Oracle Database Permissions" on page 30
Install the HA for Oracle packages	"Installing the HA for Oracle Packages" on page 33

TABLE 3 Task Map: Installing and Configuring HA for Oracle with Clustered Oracle ASM *(Continued)*

Task	Cross-Reference
Register and configure HA for Oracle	"How to Register and Configure HA for Oracle with Clustered Oracle ASM Instance Using Oracle Solaris Cluster Maintenance Commands" on page 52
Verify the HA for Oracle installation	"Verifying the HA for Oracle Installation" on page 59
Tune the HA for Oracle fault monitor	"Tuning the HA for Oracle Fault Monitors" on page 61
(Optional) Customize the HA for Oracle server fault monitor	"Customizing the HA for Oracle Server Fault Monitor" on page 65
(Optional) Upgrade HA for Oracle resource types	"Upgrading HA for Oracle Resource Types" on page 75
(Optional) Change the role of an Oracle Data Guard instance	"Changing the Role of an Oracle Data Guard Instance" on page 78

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.

Note – Unless otherwise stated, instructions and information for Oracle 11*g* apply to both Oracle 11*g* release 1 and Oracle 11*g* release 2.

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 must meet these requirements before you proceed with your HA for Oracle installation and configuration. HA for Oracle can be configured to run in non-global zones if required. If you are configuring HA for Oracle to run in a non-global zone, you must use a highly available local file system. You can also configure the HA for Oracle in a zone cluster.

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

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 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 new storage option introduced with Oracle database 10*g* that provides the services of a file system, logical volume manager, and software redundant array of independent disks (RAID) in a platform independent manner. If you are planning to install Oracle ASM, you should choose the appropriate Oracle ASM instance and disk group for your Oracle database installation. There are two types of Oracle ASM instances, single Oracle ASM instance and clustered Oracle ASM instance. See "Choosing the Appropriate Oracle ASM Instance" on page 95 for information on choosing the appropriate Oracle ASM instance. For more information on Oracle ASM, see the Oracle documentation corresponding to the Oracle database version you are using.

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

  If the single instance Oracle database is installed on a node using an Oracle ASM disk group, then certain files are not included within the Oracle ASM disk group and they resided only locally. You should ensure that these files are copied to the other cluster nodes from the node where the single instance Oracle database is created.

  If you are using an Oracle 11*g* release 1 database, you copy the following directories to other cluster nodes:

  - `${ORACLE_BASE}/diag/rdbms/`*sid*
  - `${ORACLE_BASE}/admin`
  - `${ORACLE_HOME}/dbs/pfile.ora`

If you are using an Oracle 10*g* database, you copy the following directories to other cluster nodes:

- ${ORACLE_HOME}/admin
- ${ORACLE_HOME}/dbs/init*sid*.org

In the following example, you use the /usr/sfw/sbin/wget command from the target cluster node to recursively copy the ${ORACLE_HOME}/admin directory in Oracle 10*g*.

```
# su dbuser
cd ORACLE_HOME
/usr/sfw/bin/wget -r -nH --Cut-dirs=n ftp://user:password@host:ORACLE_HOME/admin
```

In the /usr/sfw/sbin/wget command, the -r option is for recursive copy and -nH option is for disabling prefixing of host directories. Specify the --Cut-dirs option to ignore directory components. For the detailed description of options for /usr/sfw/sbin/wget, see wget man page.

Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of HA for Oracle. Write the answers to these questions in the space that is provided on the data service worksheets in "Configuration Worksheets" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

- What resource groups will you use for network addresses and application resources and the dependencies between them?

- What is the logical hostname (for failover services) or shared address (for scalable services) 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?

  If you are planning to use Oracle ASM storage, you should decide if you require single instance or clustered instance Oracle ASM. See "Choosing the Appropriate Oracle ASM Instance" on page 95 for information on how to choose the appropriate Oracle ASM installation.

  For information about standby databases, 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

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 the Solaris Zettabyte File System (ZFS) for Oracle files, configure a highly available local ZFS. For more information, see "How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS Highly Available" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

If you use the Sun QFS file system, configure the Sun QFS file system for using with Oracle. For more information on configuring the Sun QFS file system, see "Using the Sun QFS Shared File System" in *Oracle Solaris Cluster Data Service for Oracle Real Application Clusters 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 you should 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 or zone, 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 or zones 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

You can create group entries in a network name service (for example, NIS or NIS+). If you create group entries in this way, add your entries to the local /etc/inet/hosts file to eliminate dependency on the network name service.

5   On each node or zone, 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 or zones 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   If you are using Oracle database 10.2.0.3 or higher, you should perform the following steps on each Solaris zone.

> **Note** – You must perform these steps while creating a zone. If you fail to perform these steps, it results in an error. To resolve the error, you should set the necessary privileges and then run *ORACLE_HOME*/bin/localconfig reset *ORACLE_HOME*.

```
# zoneadm -Z zcname halt
# zonecfg -Z zcname info limitpriv
# zoneadm -Z zcname set limitpriv=default,proc_priocntl
# zoneadm -Z zcname info limitpriv
# zoneadm -Z zcname boot
```

▼ 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/rdsk/dn
# chgrp dba /dev/md/metaset/rdsk/dn
# chmod 600 /dev/md/metaset/rdsk/dn
```

metaset   Specifies the name of the disk set

/rdsk/d*n*   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/rdsk/dn
```

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

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

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

---

1 Configure the disk devices for the VxVM software to use.

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

2 If you use raw devices to contain the databases, run the following commands on the current disk-group primary to change each device'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 command for each raw device.

```
# vxedit -g diskgroup set user=oracle group=dba mode=600 volume
```

diskgroup   Specifies the name of the disk group

volume   Specifies the name of the raw device within the disk group

b. Verify that the changes are effective.

```
# ls -lL /dev/vx/rdsk/diskgroup/volume
```

c. Synchronize the device group with the cluster to keep the VxVM namespace consistent throughout the cluster.

```
# cldevicegroup sync diskgroup
```

▼ 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 VxVM or Solaris Volume Manager.

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

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

2 Setup 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 Install 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.

You should decide the type of Oracle ASM instance you want to use. You can use Oracle ASM instance as single instance Oracle ASM or clustered Oracle ASM using Oracle Clusterware instances. For information on how to choose the appropriate Oracle ASM instance, see "Choosing the Appropriate Oracle ASM Instance" on page 95.

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

If you are using clustered Oracle ASM, you should first install the Oracle Clusterware software and then install Oracle ASM. If you are using clustered Oracle ASM on SPARC, you should also install the ORCLudlm package before installing Oracle ASM.

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.

When you create the same disk group with Oracle ASM on different cluster nodes, you must create the disk group one node at a time and dismount the disk group before creating the same disk group on another cluster node.

Note – You should 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.

How to Verify 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;
#
```

Note – If you are using Oracle database 11*g*, you should use `sqlplus "/ as sysasm"` instead of `sqlplus "/ as sysdba"`.

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 21.

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 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 27. After the system resources have been configured for Oracle you can start the installation itself.

4 **(Optional) If you are using HA for Oracle with Oracle 10***g* **release 1, prevent the Oracle `cssd` daemon from being started.**

If you are using HA for Oracle with Oracle 10*g* release 2, omit this step.

Remove the entry for the Oracle `cssd` daemon from the `/etc/inittab` file on the node where the Oracle software is installed. To remove this entry, remove the following line from the `/etc/inittab` file:

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

HA for Oracle does not require the Oracle cssd daemon. Therefore, removal of this entry does *not* affect the operation of Oracle 10*g* release 1 with HA for Oracle. If your Oracle installation changes so that the Oracle cssd daemon is required, restore the entry for this daemon to the /etc/inittab file.

> ⚠️ **Caution** – If you are using Oracle RAC 10*g* release 1, do *not* remove the entry for the cssd daemon from the /etc/inittab file.

If you remove the entry for the Oracle cssd daemon from the /etc/inittab file, you prevent unnecessary error messages from being displayed. Otherwise, an attempt by the init(1M) command to start the Oracle cssd daemon might cause such error messages to be displayed. These error messages are displayed if the Oracle binary files are installed on a highly available local file system or on the cluster file system. The messages are displayed repeatedly until the file system where the Oracle binary files are installed is mounted. These error messages are as follows:

```
INIT: Command is respawning too rapidly. Check for possible errors.
id:  h1 "/etc/init.d/init.cssd run >/dev/null 2>&1 >/dev/null"

Waiting for filesystem containing $CRSCTL.
```

If you are using HA for Oracle on the x86 platform, unnecessary error messages about the unavailability of the UNIX Distributed Lock Manager (Oracle UDLM) might also be displayed.

These messages are displayed if the following events occur:

- A node is running in noncluster mode. In this situation, file systems that Oracle Solaris Cluster controls are never mounted.

- A node is booting. In this situation, the messages are displayed repeatedly until Oracle Solaris Cluster mounts the file system where the Oracle binary files are installed.

- Oracle is started on or fails over to a node or zone where the Oracle installation was *not* originally run. In such a configuration, the Oracle binary files are installed on a highly available local file system. In this situation, the messages are displayed on the console of the node or zone where the Oracle installation was run.

▼ 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 the *Oracle10g Installation Guide* 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 29.

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 29. This procedure is not required for any additional databases that you might create and configure.
- 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 29.
  - **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 21 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 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

⚠️ **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>    exit;
  #
  ```

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

- **To use the 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 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 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 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 NET8 for the Oracle Solaris Cluster software.

The listener.ora file must be accessible from all the nodes or zones that are in the cluster. Place these files either under the cluster file system or in the local file system of each node or zone 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(1CL) command to set the resource extension parameter User_env, which sources the user-environment file. See "SUNW.oracle_listener Extension Properties" on page 85 or "SUNW.oracle_server Extension Properties" on page 81 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 =
    (ADDRESS_LIST =
            (ADDRESS =
                (PROTOCOL = TCP)
                    (HOST = logical-hostname) <- use logical hostname
                (PORT = 1527)
            )
    )
.
.
SID_LIST_LISTENER =
    .
            .
                        (SID_NAME = SID) <- Database name,
default is ORCL
```

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
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 to install the HA for Oracle packages.

Installing the HA for Oracle Packages

If you did not install the HA for Oracle packages during your initial Oracle Solaris Cluster installation, perform this procedure to install the packages. To install the packages, use the installer program.

> **Note –** You need to install the HA for Oracle packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for Oracle Packages

Perform this procedure on each cluster node where you are installing the HA for Oracle packages.

You can run the `installer` program with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar.

> **Note –** Even if you plan to configure this data service to run in non-global zones, install the packages for this data service in the global zone. The packages are propagated to any existing non-global zones and to any non-global zones that are created after you install the packages.

Before You Begin Ensure that you have the Oracle Solaris Cluster installation media.

If you intend to run the `installer` program with a GUI, ensure that your `DISPLAY` environment variable is set.

1 On the cluster node where you are installing the data service packages, become superuser.

2 Load the Oracle Solaris Cluster installation media into the DVD-ROM drive.

If the Volume Management daemon vold(1M) is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `/cdrom` directory.

3 Change to the installation wizard directory of the DVD-ROM.

- **If you are installing the data service packages on the SPARC platform, type the following command:**

  ```
  # cd /cdrom/cdrom0/Solaris_sparc
  ```

- **If you are installing the data service packages on the x86 platform, type the following command:**

  ```
  # cd /cdrom/cdrom0/Solaris_x86
  ```

4 Start the installation wizard.

```
# ./installer
```

5 When you are prompted, accept the license agreement.

6   From the list of Oracle Solaris Cluster agents under Availability Services, select the data service for Oracle.

7   If you require support for languages other than English, select the option to install multilingual packages.

English language support is always installed.

8   When prompted whether to configure the data service now or later, choose Configure Later.

Choose Configure Later to perform the configuration after the installation.

9   Follow the instructions on the screen to install the data service packages on the node.

The installation wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.

10   (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.

The Product Registration option is not available with the CLI. If you are running the installation wizard with the CLI, omit this step.

11   Exit the installation wizard.

12   Unload the installation media from the DVD-ROM drive.

    a.   To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

    b.   Eject the DVD-ROM.

```
# eject cdrom
```

Next Steps   See "Registering and Configuring HA for Oracle" on page 36 to register HA for Oracle and to configure the cluster for the data service.

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(1CL) **utility.** For more information, see "How to Register and Configure HA for Oracle by Using the clsetup Utility" on page 37.
- **Oracle Solaris Cluster Manager.** For more information, see the Oracle Solaris Cluster Manager online help.
- **Oracle Solaris Cluster maintenance commands.** For more information, see "How to Register and Configure HA for Oracle Without Oracle ASM Using Oracle Solaris Cluster Maintenance Commands" on page 41.

The clsetup utility and Oracle Solaris Cluster Manager each provide a wizard for configuring HA for Oracle. The wizards reduce the possibility for configuration errors that might result from command syntax errors or omissions. These wizards also ensure 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 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(1CL) 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 Appendix A, "Standard Properties," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about all Solaris Cluster properties.

"SUNW.oracle_server Extension Properties" on page 81 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
- Alert_log_file
- Connect_string

▼ How to Register and Configure HA for Oracle by Using the `clsetup` Utility

This procedure provides the long forms of the 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 or zones where Oracle could potentially run.

- Kernel variables for the UNIX operating system are configured for Oracle.

- The Oracle software is configured for all nodes or zones that could potentially run Oracle.

- The data service packages are installed.

Ensure that you have the following information:

- The names of the cluster nodes or zones 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   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 configuringHA 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 or zones where you require Oracle to run.

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

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

    Ensure that the nodes or zones are listed in the order in which the nodes or zones are to appear in the node list of the resource group in which the Oracle resource is to be placed. The first node or zone in the list is the primary node or zone of this resource group.

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

    Ensure that the nodes or zones are listed in the order in which the nodes or zones are to appear in the node list of the resource group in which the Oracle resource is to be placed. The first node or zone in the list is the primary node or zone of this resource group.

7   To confirm your selection of nodes or zones, 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 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 Solaris Cluster resources that the utility will create.

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

11   If you require a different name for any Oracle Solaris Cluster resources properties, change each name 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 filea 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 is to 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 Solaris Cluster objects that the utility will create.

21   If you require a different name for any 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 Using Oracle Solaris Cluster Maintenance Commands

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 or zones where Oracle could potentially run.

- Kernel variables for the UNIX operating system are configured for Oracle.

- The Oracle software is configured for all nodes or zones that could potentially run Oracle.

- The data service packages are installed.

Ensure that you have the following information:

- The names of the cluster nodes or zones 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.

```
# 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 Solaris ZFS, because the resource group was created when the highly available local ZFS was configured. 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 or zones on which the data service can run with the -n option, as follows.

```
# clresourcegroup create [-n node-zone-list] resource-group
```

-n *node-zone-list*   Specifies a comma-separated, ordered list of zones that can master this resource group. The format of each entry in the list is *node*:*zone*. In this format, *node* specifies the name or ID of a node and *zone* specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

   The order in this list determines the order in which the nodes or zones are considered primary during failover. This list is optional. If you omit this list, the global zone of each cluster node can master the 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] \
logicalhotname-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 `SUNW.HAStoragePlus` resource type with the cluster.

```
# clresourcetype register SUNW.HAStoragePlus
```

7 Add a resource of type `SUNW.HAStoragePlus` to the failover resource group.

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

> **Note –** The `SUNW.HAStoragePlus` resource type must be version 4 if it is to be supported in non-global zones.

> **Note –** If you use the Solaris ZFS for Oracle files, omit this step. The `SUNW.HAStoragePlus` resource was created when the highly available local ZFS was configured. For more information, see "How to Prepare the Nodes" on page 21.

```
# clresource create -g resource-group -t SUNW.HAStoragePlus \
-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 or zone.

```
# 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.

■ Oracle server resource:

```
# clresource create  -g resourcegroup \
-t SUNW.oracle_server \
-p Connect_string=user/passwd \
-p ORACLE_SID=instance \
-p ORACLE_HOME=Oracle-home \
-p Alert_log_file=path-to-log \
-p Restart_type=entity-to-restart \
[-p Dataguard_role=role] \
[-p Standby_mode=mode] \-p resource_dependencies=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=storageplus-resource  resource
```

-g *resource-group*
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 Alert_log_file=*path-to-log*
Sets the path under $ORACLE_HOME for the server message log.

-p Connect_string=*user/passwd*
Specifies the user and password that the fault monitor uses to connect to the database. These settings must agree with the permissions that you set up in "How to Set Up Oracle Database Permissions" on page 30. If you use Solaris authorization, type a slash (/) instead of the user name and password.

-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 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.

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 36 for a list of extension properties.

The resources are created in the enabled state.

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

The following example shows how to register HA for Oracle on a two-node cluster.

```
Cluster Information
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)

(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 Alert_log_file=/global/oracle/message-log \
-p ORACLE_SID=ora-srvr -p Connect_string=scott/tiger \
-p Dataguard_role=STANDBY -p Standby_mode=PHYSICAL 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 2 Registering HA for Oracle to Run in the Non-Global Zone

The following example shows how to register HA for Oracle on a two-node cluster.

```
Cluster Information
Node names: phys-schost-1, phys-schost-2
Non-global zone names: sc1zone1, sc2zone1
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)

(Create the failover resource group to contain all of the resources.)
# clresourcegroup create phys-schost-1:sc1zone1,phys-schost-2:sc2zone1 \
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 Alert_log_file=/global/oracle/message-log \
-p ORACLE_SID=ora-srvr -p Connect_string=scott/tiger \
-p Dataguard_role=STANDBY -p Standby_mode=PHYSICAL 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 3   Registering HA for Oracle to Run in a Zone Cluster

The following example shows how to register HA for Oracle in a zone cluster.

```
Cluster Information
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)

(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 Alert_log_file=/global/oracle/message-log \
-p ORACLE_SID=ora-srvr -p Connect_string=scott/tiger \
-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 Register and Configure HA for Oracle with Single Instance Oracle ASM Using Oracle Solaris Cluster Maintenance Commands

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

This procedure explains the steps to register and configure single instance Oracle ASM using Oracle Solaris Cluster maintenance commands. The Oracle_Home for Oracle ASM is on a local file system only.

Note – Oracle ASM disk group and Oracle database can be configured in different resource groups. If the Oracle ASM disk group and Oracle database are configured in different resource groups, you should set a strong positive affinity (+++) between the Oracle ASM disk group resource group and Oracle database resource group.

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 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 ASM disk-group resource type.

- **For Oracle 10***g* **and 11***g* **release 1 only, use the `SUNW.asm_diskgroup` resource type.**

  ```
  # clresourcetype register SUNW.asm_diskgroup
  ```

- **For Oracle 11***g* **release 2 only, use the `SUNW.scalable_asm_diskgroup_proxy` resource type.**

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

3 Create a scalable resource group *asm-inst-rg* **for the Oracle ASM instance.**

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

-S *asm-inst-rg*   Specifies the name of the scalable resource group for the Oracle ASM instance.

4 Add a logical hostname resource to the `asm-dg-rg` resource group.

```
# clreslogicalhostname create -g asm-dg-rg [-h logicalhostname] \
logicalhostname-rs
```

-h *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.

5 Add a resource of type `SUNW.scalable_asm_instance` to the *asm-inst-rg* **resource group.**

```
# clresource create -g asm-inst-rg \
-t SUNW.scalable_asm_instance \
-p ORACLE_HOME=oracle-home \
-d asm-inst-rs
```

-g *asm-inst-rg*
  Specifies the name of the resource group for the Oracle ASM instance.

-p ORACLE_HOME=*oracle-home*
  Sets the path to the Oracle ASM home directory.

Note – You must use only a local file system for the Oracle ASM home directory.

-d *asm-inst-rs*
>   Specifies the name of the Oracle ASM instance resource that you are creating.

6   Bring online the scalable resource group in a managed state on a cluster node.

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

7   Create an Oracle ASM resource group *asm-dg–rg*.

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

8   Set a strong positive affinity on Oracle ASM disk group by the resource group *asm-dg-rg*.

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

9   Add a resource of type SUNW.asm_diskgroup to the asm-dg-rg resource group.

- **For Oracle 10*g* and 11*g* release 1 only, use the SUNW.asm_diskgroup resource type.**

    ```
    # clresource create -g asm-dg-rg -t SUNW.asm_diskgroup \
    -p asm_diskgroups=dg[,dg...] \
    -p resource_dependencies_offline_restart=asm-inst-rs,asm-stor-rs \
    -d asm-dg-rs
    ```

    -t SUNW.asm_diskgroup
    >   Specifies the type of the resource to add.

    -p asm_diskgroups=*dg*
    >   Specifies the ASM disk group.

    asm-stor-rs
    >   Specifies the name of the Oracle ASM storage resource.

    -d *asm-dg-rs*
    >   Specifies the name of the resource that you are creating.

- **For Oracle 11*g* release 2 only, 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
    ```

10   Bring online the Oracle ASM disk resource group in a managed state.

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

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

```
# clresource status +
```

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

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

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

13 Create Oracle application resources in the failover resource group.

- Oracle server resource:

  ```
  # clresource create -g asm-dg-rg \
  -t SUNW.oracle_server \
  -p Connect_string=user/passwd \
  -p ORACLE_SID=instance \
  -p ORACLE_HOME=Oracle-home \
  -p Alert_log_file=path-to-log \
  -p Restart_type=entity-to-restart \
  -p Resource_dependencies_offline_restart=asm-dg-rs \
  -d   ora-db-rs
  ```

- Oracle listener resource:

  ```
  # clresource create -g asm-dg-rg \
  -t SUNW.oracle_listener \
  -p LISTENER_NAME=listener \
  -p ORACLE_HOME=Oracle-home \
  -p resource_dependencies_offline_restart=asm-dg-rs logicalhostname-rs \
  -d ora-lsr-rs
  ```

-g *asm-dg-rg*
Specifies the name of the Oracle ASM disk resource group into which the resources are to be
placed.

-t SUNW.oracle_server/listener
Specifies the type of the resource to add.

-p Alert_log_file=*path-to-log*
Sets the path under $ORACLE_HOME for the server message log.

-p Connect_string=*user/passwd*
Specifies the user and password that the fault monitor uses to connect to the database. These
settings must agree with the permissions that you set up in "How to Set Up Oracle Database
Permissions" on page 30. If you use Solaris authorization, type a slash (/) instead of the user
name and password.

-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.

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 36 for a list of extension properties.

14   **Bring online the Oracle server resource.**

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

15   **Bring online the Oracle listener resource.**

```
# clresource enable ora-lsr-rs
```

▼ How to Register and Configure HA for Oracle with Clustered Oracle ASM Instance Using Oracle Solaris Cluster Maintenance Commands

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=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   SPARC: If you are using SPARC, perform the following steps. Other wise go to Step 8.

Note – You can perform this step only if you have installed the ORCLudlm package before performing the Oracle ASM software installation.

---

a.   Register the resource type for Oracle UDLM.

```
# clresourcetype register SUNW.rac_udlm
```

b.   Add a resource of type SUNW.rac_udlm to the *rac-fmwk-rg* **resource group.**

```
# clresource create -g rac-fmwk-rg -t SUNW.rac_udlm \
-p resource_dependencies=rac-fmwk-rs rac-udlm-rs
```

8   Register the appropriate resource type for the file system you are using and create a resource group.

Note – You should decide the file system you are going to use for Oracle_Home and set up the file system before performing this step. For information on how to setup a file system, see "Preparing the Nodes and Disks" on page 21.

If you are using a QFS shared file system, perform the following steps to register the QFS shared file system resource type and create the resource groups. Otherwise, you should use the SUNW.HAStoragePlus resource type.

- **If you are using QFS shared file system, perform the following steps to register the QFS shared file system resource type and create the resource groups. Otherwise, you should use the SUNW.HAStoragePlus resource type.**

  a.   Register the resource type for QFS shared file system.

  ```
  # clresourcetype register SUNW.qfs
  ```

  b.   Create a resource group *qfs-rg*.

  ```
  # clresourcegroup create qfs-rg
  ```

  c.   Add a resource of type SUNW.qfs to the *qfs-rg* **resource group.**

  ```
  # clresource create -g qfs-rg -t SUNW.qfs -p QFSFileSystem=qfs-mp qfs-rs
  ```

  -g *qfs-rg*

  Specifies the name of the resource group.

  -t SUNW.qfs

  Specifies the type of the resource to add.

-p QFSFileSystem=*qfs-mp*
Specifies the QFS shared file system mount point.

qfs-rs
Specifies the name of the QFS shared file system resource that you are creating.

d. **Bring online the** *qfs-rg* **resource group in a managed state on a cluster node.**

```
# clresourcegroup online -eM qfs-rg
```

e. **Register the resource type for SUNW.ScalMountPoint.**

```
# clresourcetype register SUNW.ScalMountPoint
```

f. **Create a scalable resource group** *scal-mp-rg*.

```
# clresourcegroup create -S scal-mp-rg
```

g. **Add a resource of type SUNW.ScalMountPoint to the** *scal-mp-rg* **resource group.**

```
# clresource create -g scal-mp-rg -t SUNW.ScalMountPoint \
-p mountpointdir=mount-point \
-p filesystemtype=s-qfs \
-p targetfilesystem=qfs-name \
-p resource_dependencies=qfs-rs \
-d scal-mp-rs
```

-g *scal-mp-rg*
Specifies the name of the scalable resource group.

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

-p mountpointdir=*mount-point*
Specifies the directory to which the file system is mounted.

-p filesystemtype=*s-qfs*
Specifies the file system type you are using.

-p targetfilesystem=*qfs-name*
Specifies the file system name that is mounted on this mount point.

-d *scal-mp-rs*
Specifies the name of the scalable mount point resource that you are creating.

h. **Bring online the** *scal-mp-rg* **resource group in a managed state on a cluster node.**

```
# clresourcegroup online -eM scal-mp-rg
```

■ **If you are using file systems other than QFS shared file system, perform the following steps to register the SUNW.HAStoragePlus resource type and create the resource groups.**

a. **Register the resource type for SUNW.HAStoragePlus resource type.**

```
# clresourcetype register SUNW.HAStoragePlus
```

    b. Add a resource of type `SUNW.HAStorgePlus` to the *ora-db-rg* **resource group.**

```
# clresource create -g ora-db-rg -t SUNW.HAStoragePlus \
-p filesystemmountpoints=mount-point-list -d hastp-rs
```

    c. Bring online the *ora-db-rg* **resource group in a managed state on a cluster node.**

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

9   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 ASM disk-group resource type.

      ■ **For Oracle 10***g* **and 11***g* **release 1 only, use the `SUNW.asm_diskgroup` resource type.**

```
# clresourcetype register SUNW.asm_diskgroup
```

      ■ **For Oracle 11***g* **release 2 only, use the `SUNW.scalable_asm_diskgroup_proxy` resource type.**

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

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

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

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

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

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

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

13 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=rac-fmwk-rs \
-d crs-fmwk-rs
```

14 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 crs_home=crs-home \
-p "oracle_sid{node1}"=instance \
-p "oracle_sid{node2}"=instance \
-p resource_dependencies_offline_restart=crs-fmwk-rs,qfs-mp-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=*instance*
   Sets the Oracle system identifier.

-p ORACLE_HOME=*Oracle-home*
   Sets the path to the Oracle home directory.

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

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

- **For Oracle 10*g* and 11*g* release 1 only, use the SUNW.asm_diskgroup resource type.**

   ```
   # clresource create -g asm-dg-rg -t SUNW.asm_diskgroup \
   -p asm_diskgroups=dg[,dg...] \
   -p resource_dependencies_offline_restart=asm-inst-rs,asm-stor-rs \
   -d asm-dg-rs
   ```

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

   -p asm_diskgroups=*dg*
      Specifies the ASM disk group.

   asm-stor-rs
      Specifies the name of the Oracle ASM storage resource.

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

- **For Oracle 11*g* release 2 only, 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
   ```

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

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

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

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

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

   ```
   # clresource status +
   ```

19   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.

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

20   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.

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

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

22   Create Oracle application resources in the failover resource group.

- Oracle server resource:

```
# clresource create -g ora-db-rg \
-t SUNW.oracle_server \
-p Connect_string=user/passwd \
-p ORACLE_SID=instance \
-p ORACLE_HOME=Oracle-home \
-p Alert_log_file=path-to-log \
-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,logicalhostname-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 Alert_log_file=*path-to-log*
  Sets the path under $ORACLE_HOME for the server message log.

-p Connect_string=*user*/*passwd*
Specifies the user and password that the fault monitor uses to connect to the database. These settings must agree with the permissions that you set up in "How to Set Up Oracle Database Permissions" on page 30. If you use Solaris authorization, type a slash (/) instead of the user name and password.

-p ORACLE_SID=*instance*
Sets the Oracle system identifier.

-p ORACLE_HOME=*Oracle-home*
Sets the path to the Oracle home directory.

-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 36 for a list of extension properties.

---

23   Bring online the Oracle server resource.

    # **clresource enable** *ora-db-rs*

24   Bring online the Oracle listener resource.

    # **clresource enable** *ora-ls-rs*

More Information   Cluster File System

Note – You can also install `Oracle_Home` for Oracle ASM on a cluster file system. If you are using a cluster file system, you should create the resources `SUNW.HAStoragePlus` and `SUNW.scalable_asm_instance` and set the appropriate dependencies between these two resources.

Sample commands to create the `SUNW.HAStoragePlus` resource and `SUNW.scalable_asm_instance` resource is given as follows:

```
# clresource create -g asm-inst-rg -t SUNW.HAStoragePlus \
-p FilesystemMountPoints=cluster-file-system \
hastp-rs

# clresource create -g asm-inst-rg -t SUNW.scalable_asm_instance \
-p ORACLE_HOME=oracle-home \
-p "oracle_sid{node1}"=instance \
-p "oracle_sid{node2}"=instance \
-p resource_dependencies_restart=hastp-rs \
-p resource_dependencies_offline_restart=crs-fmwk-rs \
asm-inst-rs
```

Where to Go From Here

Go to after you register and configure Solaris Cluster 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 or zones that run HA for Oracle can start the Oracle instance and that the other nodes or zones 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 or zone 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 or zone.**

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`**` `*`user`*`/`*`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`*

`-n` *node-zone-list*   Specifies a comma-separated, ordered list of zones that can master this resource group. The format of each entry in the list is *node*`:`*zone*. In this format, *node* specifies the name or ID of a node and *zone* specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

                             The order in this list determines the order in which the nodes or zones are considered primary during failover. This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

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

7 Log in as `oracle` to the node or zone 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. The network resource is a logical hostname or a shared address. See "How to Set Up Oracle Database Permissions" on page 30.

> **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.

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

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

TABLE 4 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 36.

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 62
- "Operation of the Database Client Fault Probe" on page 62
- "Actions by the Server Fault Monitor in Response to a Database Transaction Failure" on page 63
- "Scanning of Logged Alerts by the Server Fault Monitor" on page 64

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 63.

2. If the partition is healthy, determining whether the database is operational. See "Operations to Determine Whether the Database is Operational" on page 63.

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$archive_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 or zones as explained in "How to Prepare the Nodes" on page 21.

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 65.

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 or zone.

- **Switch over.** The server fault monitor switches over the database server resource group to another node or zone. If no nodes or zones 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 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 65.

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 / 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:

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

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 or zone.

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 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:

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 4   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 4   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 5   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 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 6   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 7   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 7   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 or zones. Therefore, the custom action file that the server fault monitor uses must be identical on all cluster nodes or zones. After creating or modifying a custom action file, ensure that this file is identical on all cluster nodes or zones by propagating the file to all cluster nodes or zones. To propagate the file to all cluster nodes or zones, use the method that is most appropriate for your cluster configuration:

■ Locating the file on a file system that all nodes or zones share

■ Locating the file on a highly available local file system

■ Copying the file to the local file system of each cluster node or zone by using operating system commands such as the rcp(1) command or the rdist(1) 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.

Upgrading HA for Oracle Resource Types

The existing resource types for the HA for Oracle data service are as follows:

- SUNW.oracle_listener, which represents an Oracle listener
- SUNW.oracle_server, which represents an Oracle server

In addition to the existing resource types, there are three additional resource types included in HA for Oracle data service for Oracle ASM. These resource types are as follows:

- SUNW.asm_diskgroup (Oracle 10*g* and 11*g* release 1 only) or
  SUNW.scalable_asm_diskgroup_proxy (Oracle 11*g* release 2 only), which represents the single or clustered instance Oracle ASM disk group in a Oracle Solaris Cluster configuration. For more information about the resource type, see asm_diskgroup(5)

- SUNW.scalable_asm_instance, which represents the single instance Oracle ASM in a Oracle Solaris Cluster configuration. For more information about the resource type, see SUNW.scalable_asm_instance(5).

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

- SUNW.scalable_asm_instance_proxy, which represents a proxy for the clustered Oracle ASM instance in a Oracle Solaris Cluster configuration. For more information about the resource type, see SUNW.scalable_asm_instance_proxy(5).

Upgrade the existing resource types if you are upgrading from an earlier version of HA for Oracle.

For general instructions that explain how to upgrade a resource type, see "Upgrading a Resource Type" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Upgrading the `SUNW.oracle_listener` Resource Type

The information that you require to complete the upgrade of the SUNW.oracle_listener resource type is provided in the subsections that follow.

Information for Registering the New Resource Type Version

To determine the version of the resource type that is currently registered, use one command from the following list:

- `clresourcetype list`
- `clresourcetype show`

The resource type registration (RTR) file for this resource type is /SUNW.oracle_listener.

To determine the version of the latest installed version of the /opt/cluster/lib/rgm/rtreg/SUNW.oracle_listener resource type, whether it is registered or not, use the following command:

```
# grep -i RT_VERSION /opt/cluster/lib/rgm/rtreg/SUNW.oracle_listener
```

If the version of the latest installed resource type is later than the registered version, migrate to the newer version to ensure full functionality.

Note – In the global zone, Oracle Solaris Cluster Manager or the Resource Group option of `clsetup` locates for you the available resource type versions you can upgrade to.

Information for Migrating Existing Instances of the Resource Type

The information that you require to edit each instance of the SUNW.oracle_listener resource type is as follows:

- You can perform the migration at any time.
- If you need to specify the timeout value in seconds that the fault monitor uses to probe an Oracle listener, set the Probe_timeout extension property. For more information, see "SUNW.oracle_listener Extension Properties" on page 85.

The following example shows a command for editing an instance of the SUNW.oracle_listener resource type.

EXAMPLE 8 Editing an Instance of the `SUNW.oracle_listener` Resource Type

```
# clresource set -p Type_version=N \
  -p probe_timeout=60 oracle-lrs
```

This command edits a `SUNW.oracle_listener` resource as follows:

- The `SUNW.oracle_listener` resource is named `oracle-lrs`.
- The `Type_version` property of this resource is set to *N*, which represents the version number of the resource type you migrated to.
- The timeout value in seconds that the fault monitor uses to probe an Oracle listener is set to 60 seconds.

Upgrading the `SUNW.oracle_server` Resource Type

The information that you require to complete the upgrade of the `SUNW.oracle_server` resource type is provided in the subsections that follow.

Information for Registering the New Resource Type Version

To determine the version of the resource type that is registered, use one command from the following list:

- `clresourcetype list`
- `clresourcetype show`

The resource type registration (RTR) file for this resource type is `/opt/SUNWscor/oracle_server/etc/SUNW.oracle_server`.

To determine the version of the latest installed version of the `SUNW.oracle_server` resource type, whether it is registered or not, use the following command:

```
# grep -i RT_VERSION /opt/cluster/lib/rgm/rtreg/SUNW.oracle_server
```

If the version of the latest installed resource type is later than the registered version, migrate to the newer version to ensure full functionality.

Note – In the global zone, Oracle Solaris Cluster Manager or the Resource Group option of `clsetup` locates for you the available resource type versions you can upgrade to.

Information for Migrating Existing Instances of the Resource Type

The information that you require to edit each instance of the `SUNW.oracle_server` resource type is as follows:

- You can perform the migration at any time.

- If you customized the behavior of the server fault monitor, set the `Custom_action_file` extension property. For more information, see "Customizing the HA for Oracle Server Fault Monitor" on page 65.

The following example shows a command for editing an instance of the `SUNW.oracle_server` resource type.

EXAMPLE 9   Editing an Instance of the `SUNW.oracle_server` Resource Type

```
# clresource set  -p Type_version=N \
  -p custom_action_file=/opt/SUNWscor/oracle_server/etc/srv_mon_cust_actions \
oracle-srs
```

This command edits a `SUNW.oracle_server` resource as follows:

- The `SUNW.oracle_server` resource is named `oracle-srs`.
- The `Type_version` property of this resource is set to N, which represents the version number of the resource type you migrated to.
- Custom behavior for the fault monitor of this resource is specified in the file `/opt/SUNWscor/oracle_server/etc/srv_mon_cust_actions`.

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 or zone 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
  ```

◆ ◆ ◆

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 81
- "SUNW.oracle_listener Extension Properties" on page 85

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

Alert_log_file (string)
Specifies the absolute path of the Oracle alert log file. The Oracle software logs alerts in this file. The Oracle 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 Preset Actions for DBMS Errors and Logged Alerts. 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 65.

Default: None

Range: Minimum = 1

Tunable: Any time

`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

`Connect_string` (string)

The Oracle database user ID and password that the server fault monitor uses to connect to the database.

Specify this property as follows:

userid/*password*

As part of the setup of HA for Oracle, you must define the database user ID and password before enabling the server resource and its fault monitor. To use Solaris authentication, type a slash (/) instead of a user ID and password.

Default: None

Range: Minimum = 1

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

Introduced in release: 3.1 10/03

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

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

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

Introduced in release: 3.1 4/04

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 1.
- Logged alerts for which an action is preset are listed in Table 2.

TABLE 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 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 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 progresss
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	*	*	osnfpm: cannot create shared memory segment
2811	SWITCH	*	*	Unable to attach shared memory segment

TABLE 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 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 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 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 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 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 help you to choose the appropriate Oracle ASM instance and provides information about the dependencies between HA for Oracle resources and Oracle ASM services.

Choosing the Appropriate Oracle ASM Instance

Once you decide to use Oracle ASM with Oracle, you should choose the appropriate Oracle ASM instance. You can use the following diagram to find out if you require single instance Oracle ASM or clustered instance Oracle ASM.

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

FIGURE 1   Choosing the Appropriate Oracle ASM Instance
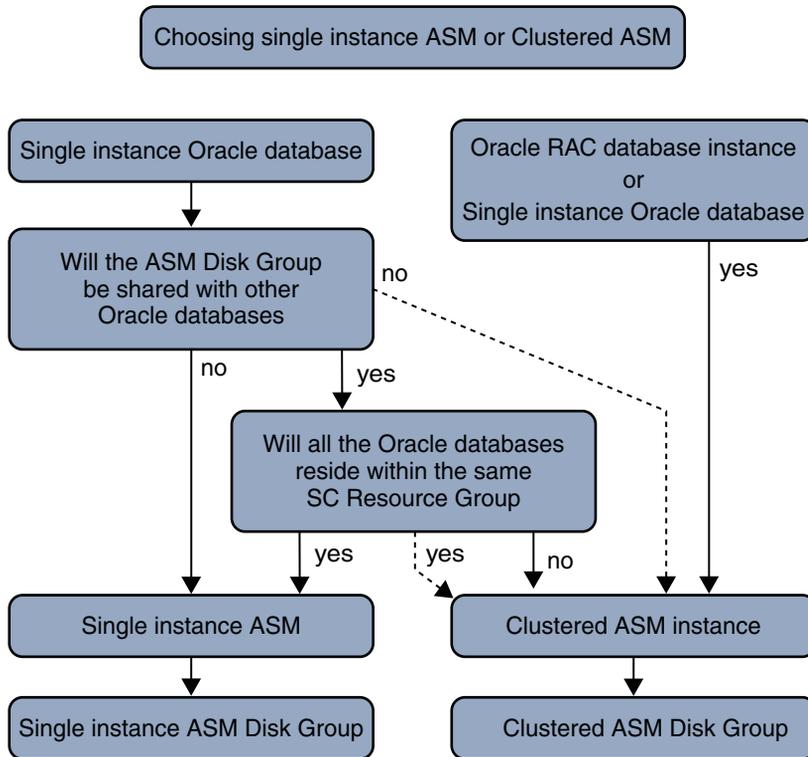
Single Instance Oracle ASM with Separate Disk Groups

This section contains two sample configuration diagrams for single instance Oracle databases with single instance Oracle ASM and single instance Oracle ASM disk groups. The second diagram is a continuation of the first diagram.

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

The following diagram represents a single instance Oracle ASM that serves two single instance Oracle databases DB1 and DB2 on a node Node1. The database DB1 exclusively uses an Oracle ASM disk group DG1 and the database DB2 exclusively uses an Oracle ASM disk group DG2. The upper part of the diagram shows the relationships that the Oracle database instances with the single Oracle ASM instances on Node1 and Node2. The single instance Oracle ASM manages two Oracle ASM disk groups DG1 and DG2 on Node1. The lower part of the diagram represents the existing Oracle Solaris Cluster resource group and resources for single instance Oracle databases and their functional requirements for single instance Oracle ASM services.

In the diagram the dotted boxes represent existing HA for Oracle resources with new Oracle ASM services. Inside the nodes, new dependencies between HA for Oracle resources and single instance Oracle ASM services are represented using arrows.

FIGURE 2   Single Instance Oracle ASM with Separate Disk Groups [1]
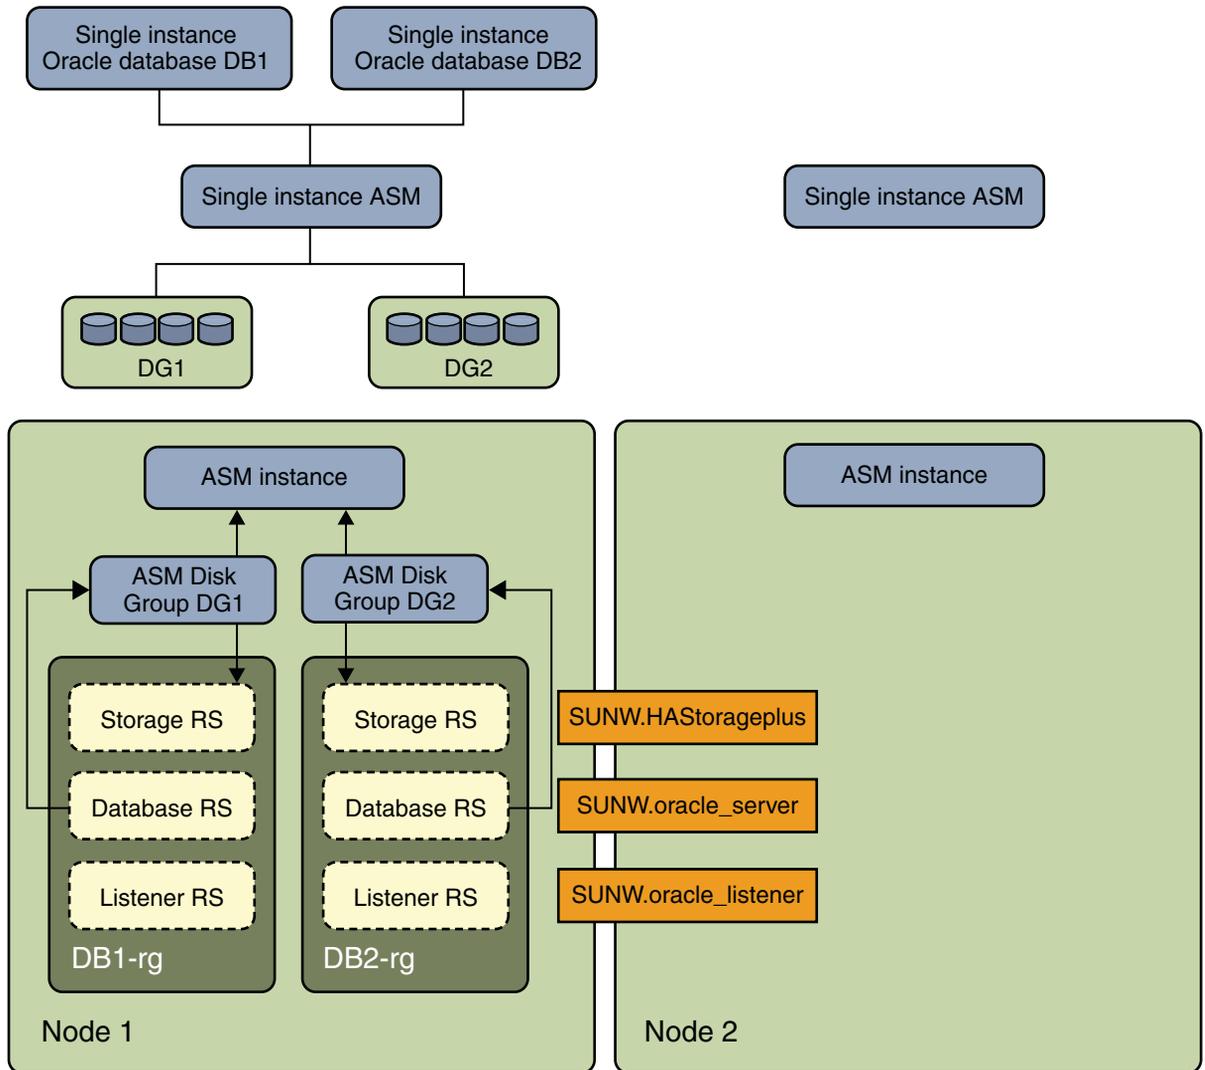
The following diagram is a continuation of the previous diagram in this section. In this diagram, the single instance Oracle database DB2 that exclusively uses Oracle ASM disk group DG2 is running on Node2 after a failover of resource group DB2–rg to Node2. The upper part of

the diagram shows the relationships of the Oracle database instances with the single Oracle ASM instances on Node1 and Node2. The Oracle ASM disk group DG2 is now served from the single instance Oracle ASM on Node2 while Oracle ASM disk group DG1 is served by the single instance Oracle ASM on Node1.

The lower part of the diagram represents the existing Oracle Solaris Cluster resource groups and resources for single instance Oracle databases and their requirement for single instance Oracle ASM services.

FIGURE 3   Single Instance Oracle ASM with Separate Disk Groups [2]
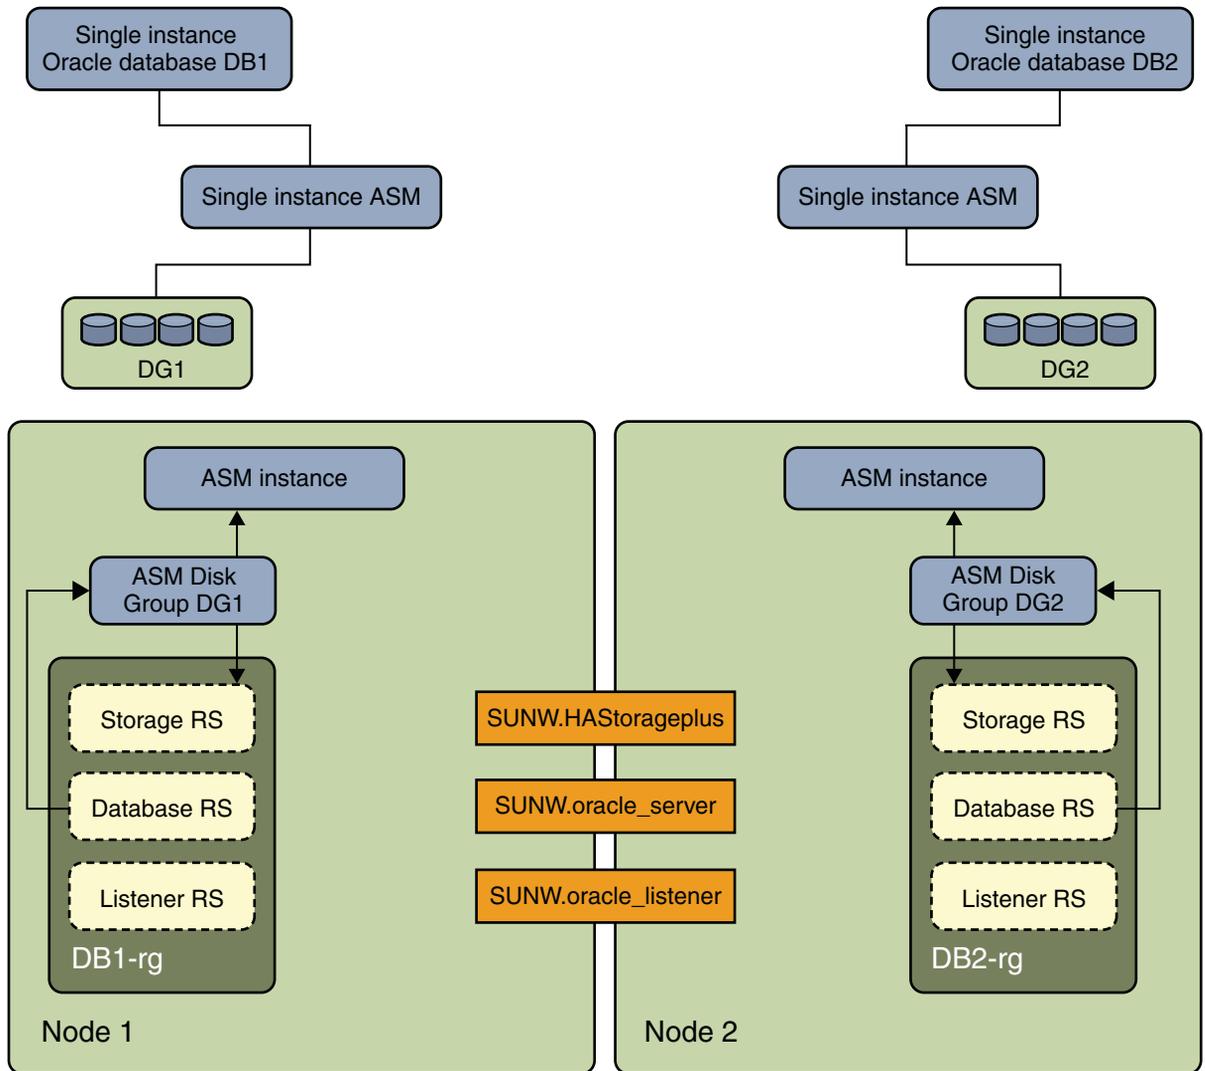
Clustered Oracle ASM with Clustered Disk Groups

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 4   Clustered Oracle ASM with Clustered Disk Groups [1]
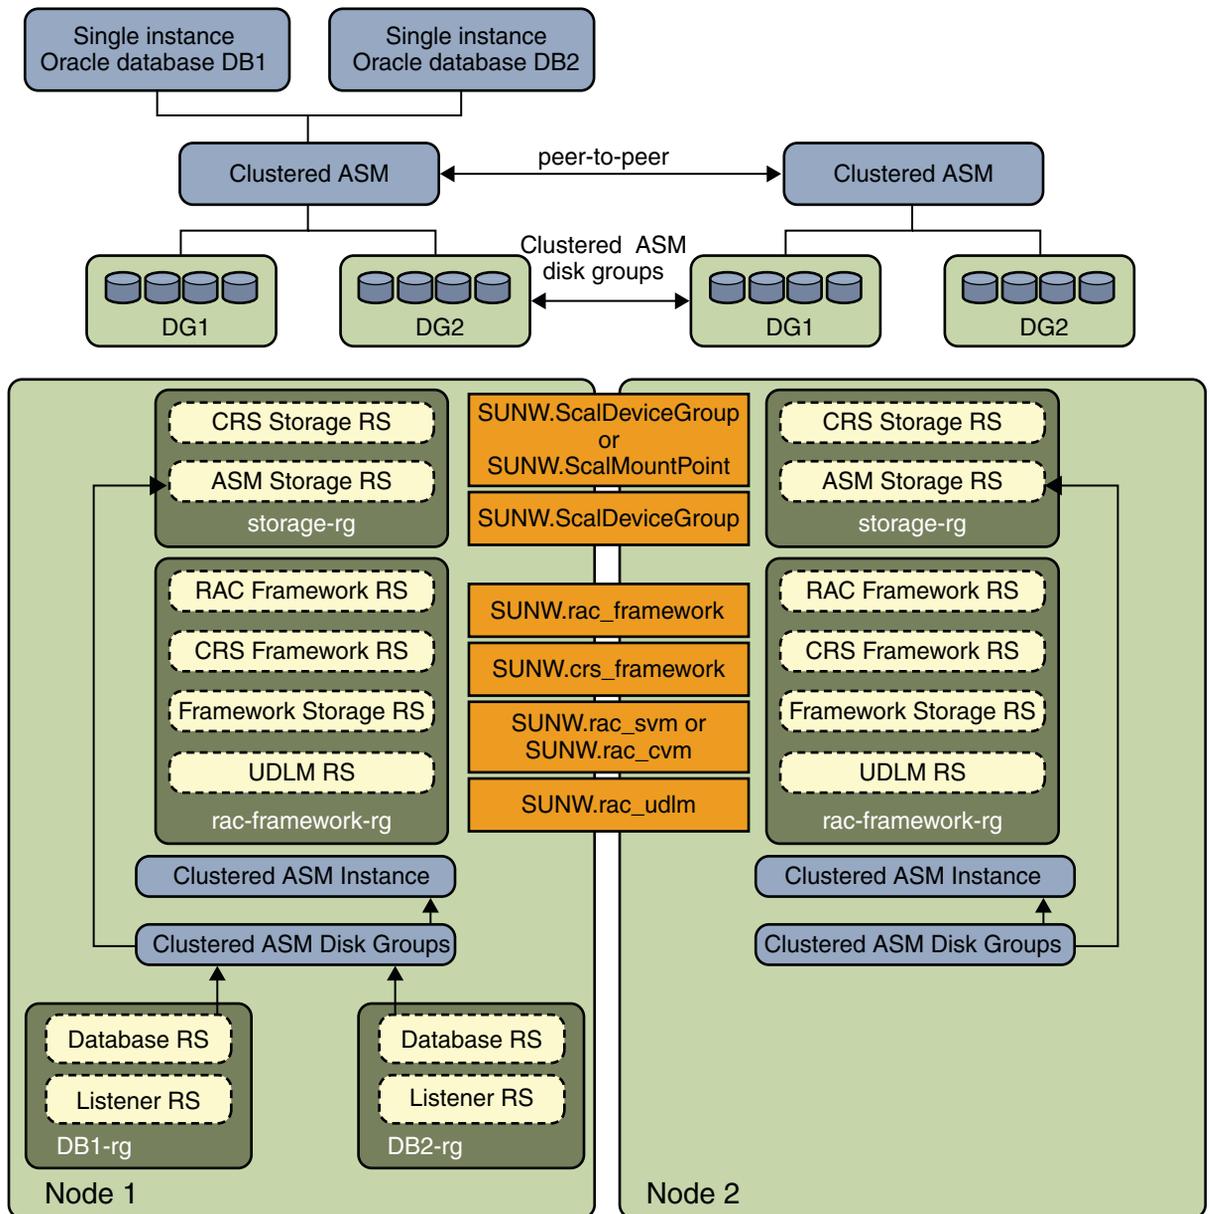
In the following diagram, a single instance Oracle database DB2 that shares Oracle ASM disk groups DG1 and DG2 is now running on Node2 after a failover of resource groups DB2-rg to Node2. The upper part of the diagram shows the relationships of the Oracle database instances with the clustered Oracle 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`, `SUNW.rac_svm` and `SUNW.rac_cvm` is not required.

FIGURE 5 Clustered Oracle ASM with Clustered Disk Groups [2]
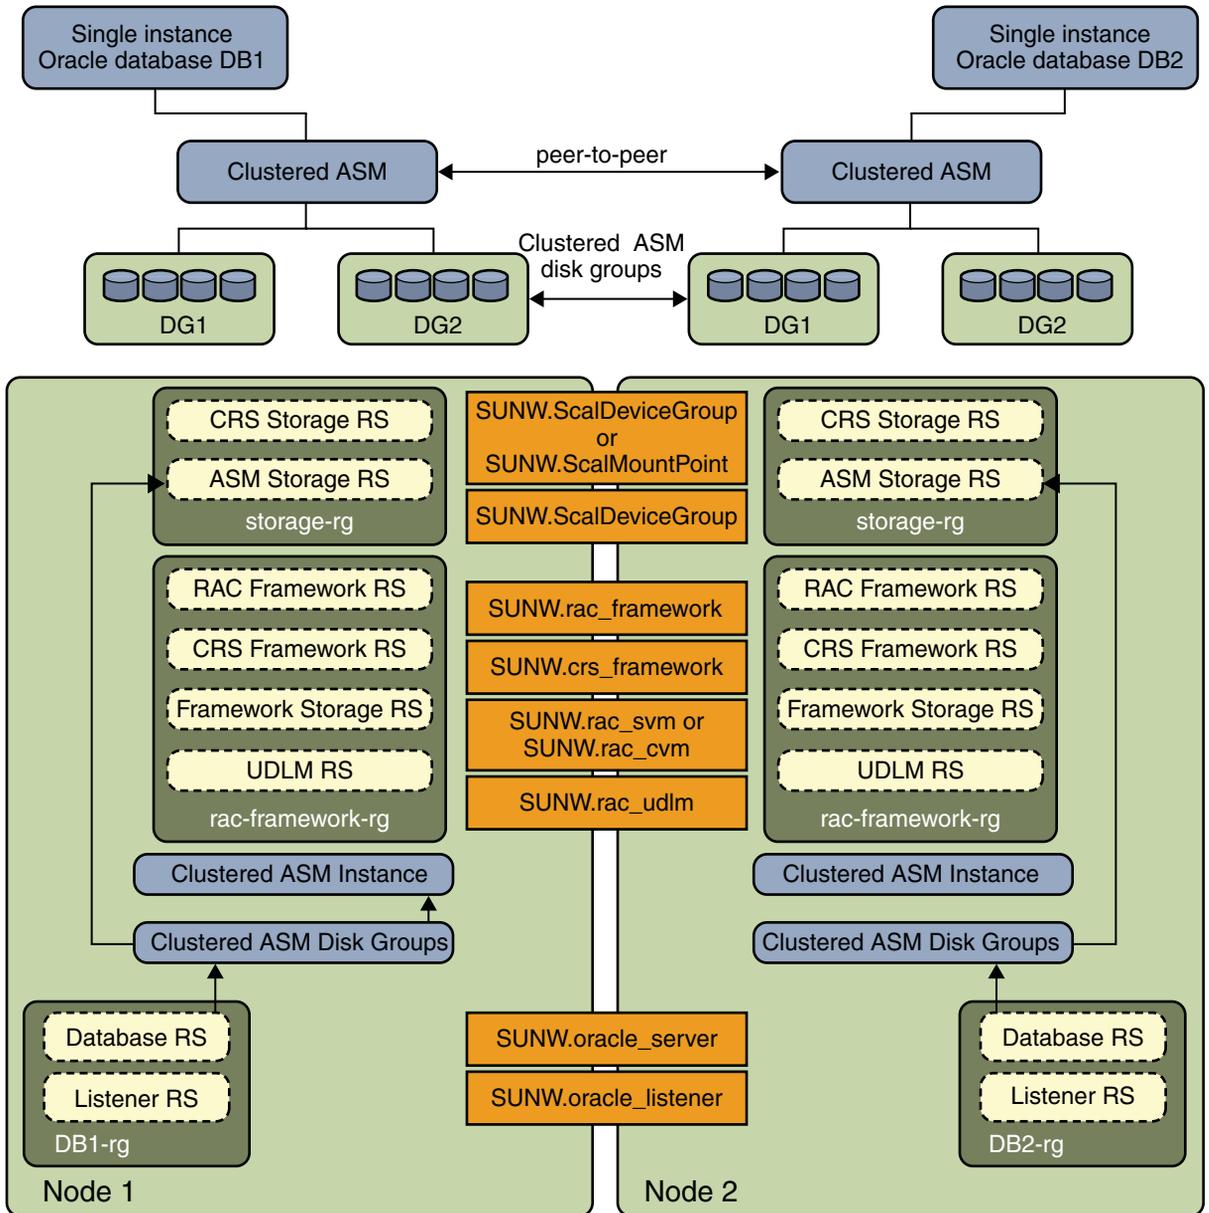
Single Instance Oracle ASM with Separate Disk Groups in a Non-Global Zone

This section provides information on the dependency requirements for single instance Oracle ASM with single instance Oracle databases and single instance Oracle ASM disk groups in non-global zones. There are two diagrams, where the second diagram is a continuation of the first diagram.

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

The following diagram represents a single instance Oracle ASM that serves two single instance Oracle databases DB1 and DB2 running within a non-global zone ZoneA on Node1. The database DB1 exclusively uses Oracle ASM disk group DG1 and database DB2 exclusively uses Oracle ASM disk group DG2.

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 single instance Oracle ASM manages two disk groups DG1 and DG2 on Node1. The lower part of the diagram represents the existing Oracle Solaris Cluster resource groups and resources to independently failover DB1, DG1 or DB2, DG2 to ZoneA on Node2.

In the diagram the dotted boxes represent existing HA for Oracle resources with new Oracle ASM resources. Inside the nodes, new dependencies between HA for Oracle resources and single instance Oracle ASM resources are represented using arrows.

FIGURE 6 Single Instance Oracle ASM with Separate Disk Groups in a Non-Global Zone [1]
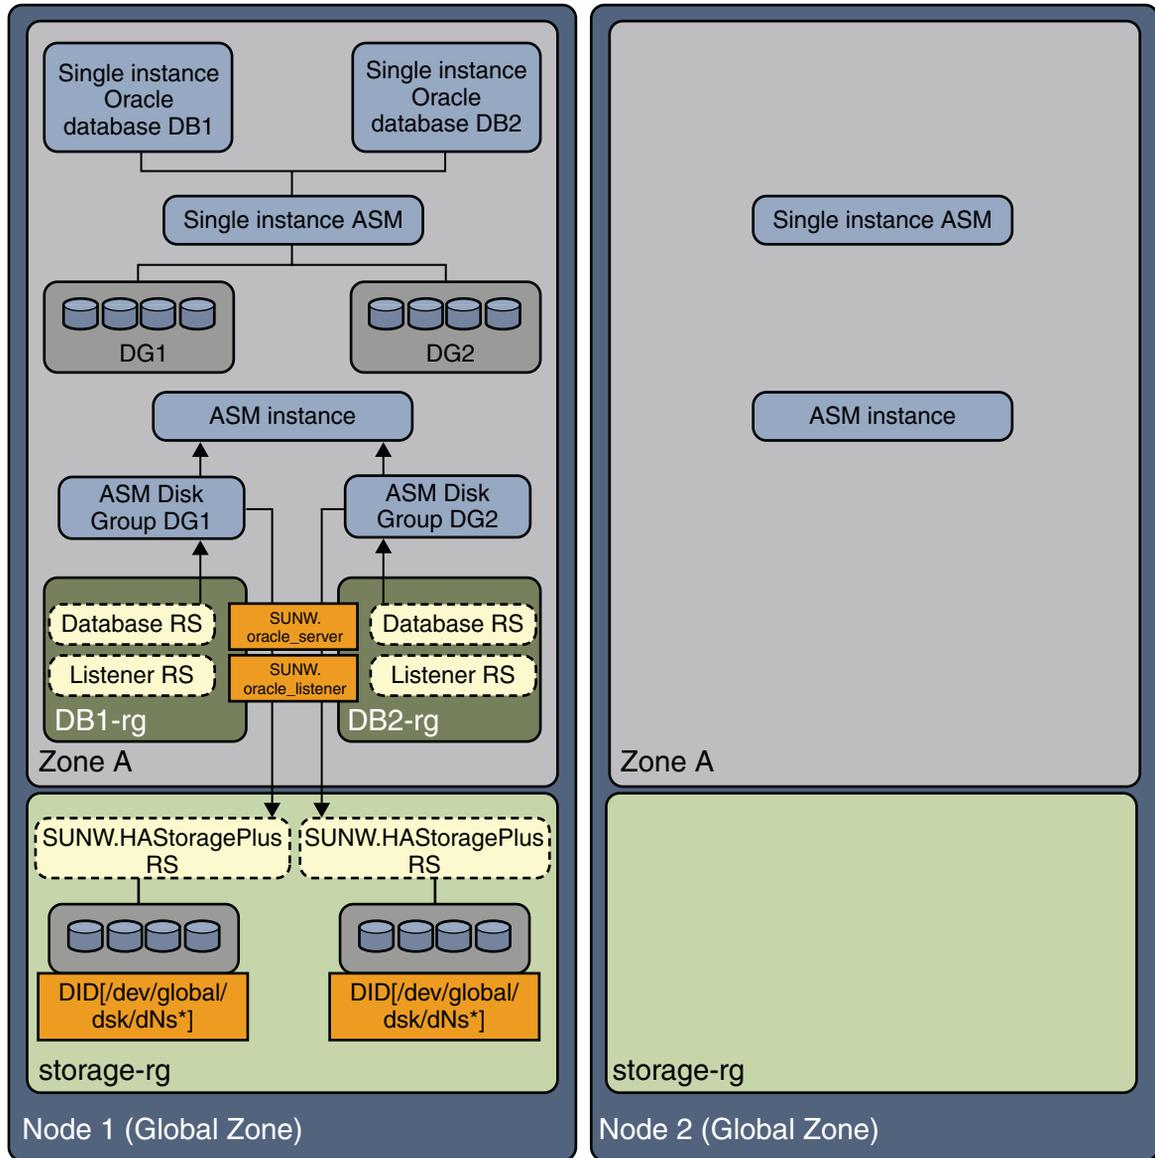
The following diagram is a continuation of the previous diagram. In this diagram, the single instance Oracle database DB2 that exclusively uses Oracle ASM disk group DG2 is running on Node2 after a failover of resource group DB2–rg to Node2. The upper part of the diagram within the non-global zone shows the relationships of the Oracle database instances with the single instance Oracle ASM instances on Node1 and Node2. The Oracle ASM disk group DG2 is now

served from the single instance Oracle ASM in ZoneA on Node2 while Oracle ASM disk group DG1 is served by the single instance Oracle ASM in ZoneA on Node1.

The lower part of the diagram represents the existing Oracle Solaris Cluster resource groups and resources to independently failover DB1, DG1 or DB2, DG2 to ZoneA on Node2.

FIGURE 7 Single Instance Oracle ASM with Separate Disk Groups in a Non-Global Zone [2]

Index