Oracle® Solaris Cluster Data Service for Oracle Real Application Clusters Guide

Part No: E25683
March 2012, E25683–01

ORACLE

120307@25097

Contents

Figures

Tables

Examples

Preface

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

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

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

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

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

Using UNIX Commands

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

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

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface	Description	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.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://www.oracle.com/technetwork/indexes/documentation/index.html.

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

Access to Oracle Support

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

Getting Help

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

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

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

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

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

◆ ◆ ◆ CHAPTER 1

Installing Support for Oracle RAC

This chapter explains how to install Support for Oracle RAC on your Oracle Solaris Cluster nodes.

Overview of the Installation Process for Support for Oracle RAC

The following table summarizes the installation tasks and provides cross-references to detailed instructions for performing the tasks.

Perform these tasks in the order in which they are listed in the table.

TABLE 1–1 Tasks for Installing Support for Oracle RAC

Task	Instructions
Plan your installation	"Preinstallation Considerations" on page 20
Prepare the Oracle Solaris Cluster nodes	"Preparing the Oracle Solaris Cluster Nodes" on page 26
Install data service packages	"Installing the Support for Oracle RAC Package" on page 35

Preinstallation Considerations

This section contains the following preinstallation information:

General Requirements.

Oracle Real Application Clusters (Oracle RAC) is an application that can run on more than one machine concurrently. Oracle RAC can run either in the global-cluster voting nodes of the global cluster or in a zone cluster. An Oracle RAC installation is entirely contained within one cluster, either a global cluster or a specific zone cluster. Keeping the Oracle RAC installation in one cluster ensures the support of multiple independent Oracle RAC installations concurrently, where each Oracle RAC installation can be of a different version or use different options, such as storage. Support for Oracle RAC enables you to run Oracle RAC on Oracle Solaris Cluster nodes and to manage Oracle RAC by using Oracle Solaris Cluster commands.

Configuring this data service involves configuring resources for the following components of an Oracle RAC installation with Oracle Solaris Cluster software:

- **The Oracle RAC framework.** These resources enable Oracle RAC to run with Oracle Solaris Cluster software. The resources also enable reconfiguration parameters to be set by using Oracle Solaris Cluster commands. You *must* configure resources for the Oracle RAC framework. For more information, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.
- **Storage for Oracle files.** These resources provide fault monitoring and automatic fault recovery for volume managers and file systems that store Oracle files. Configuring storage resources for Oracle files is optional. For more information, see "Registering and Configuring Storage Resources for Oracle Files" on page 65.
- **Oracle RAC database instances.** These resource types enable Oracle Grid Infrastructure and Oracle Solaris Cluster software to interoperate. These resource types do *not* provide fault monitoring and automatic fault recovery. The Oracle Grid Infrastructure software provides this functionality.

  Do not use Oracle Solaris project resource properties and resource-group properties with Oracle Solaris Cluster resource types. The proxy resource does not directly start the database instance. Instead, Oracle Grid Infrastructure starts the database instances, and Solaris Resource Manager abstractions do not work with these versions of Oracle RAC.

  Configuring resources, which enable Oracle Solaris Cluster software to administer Oracle RAC database instances, is optional. For more information, see "Configuring Resources for Oracle RAC Database Instances" on page 86.

Note – When you use Oracle RAC in a zone cluster, ensure that the zone cluster is not configured with the `/opt` directory as an inherited read-only directory. For an Oracle RAC configuration in a zone cluster, the `/opt` file system must be writable and unique to each zone. If the zone cluster that you intended to use for Oracle RAC is configured with an `inherit-pkg-dir` resource for the `/opt` directory, destroy and recreate the zone cluster or create a new zone cluster to meet this requirement.

If your application deployment requires the zone cluster nodes to be accessible from the public network at their host names or to have concurrent outbound traffic from each node, you must have a fixed public network address for each zone cluster node. Examples of such deployments include running Oracle RAC in zone clusters or applications using the scalable services (the `SharedAddress` resource) in the zone clusters.

Hardware and Software Requirements

Before you begin the installation, note the hardware and software requirements in the subsections that follow.

Oracle Solaris Cluster Framework Requirements

Support for Oracle RAC requires a functioning cluster with the initial cluster framework already installed. See *Oracle Solaris Cluster Software Installation Guide* for details about initial installation of cluster software.

Oracle Grid Infrastructure Software Requirements

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

Software License Requirements

Verify that you have obtained and installed the appropriate licenses for your software. If you install your licenses incorrectly or incompletely, the nodes might fail to boot correctly.

Supported Topology Requirements

Check with an Oracle service representative for the current supported topologies for Support for Oracle RAC, cluster interconnect, storage management scheme, and hardware configurations.

Software Update Installation Requirements

Ensure that you have installed all the applicable software updates for the Oracle Solaris OS, Oracle Solaris Cluster, Oracle Database, and volume manager software. If you need to install any Support for Oracle RAC software updates, you must apply these updates after you install the data service packages.

Storage Management Requirements

This section provides the following information about storage management for Oracle RAC:

Storage Management Requirements for Oracle Files

Support for Oracle RAC enables you to use the storage management schemes for Oracle files that are listed in the following tables. The tables summarize the types of Oracle files or Oracle Grid Infrastructure files that each storage management scheme can store. Ensure that you choose a combination of storage management schemes that can store all types of Oracle files.

The meaning of each symbol in the tables is as follows:

+
: Indicates that the storage management scheme can store the type of Oracle file.

-
Indicates that the storage management scheme *cannot* store the type of Oracle file.

TABLE 1–2 Storage Management Schemes for Oracle DBMS Files

Oracle File Type	Solaris Volume Manager for Sun Cluster Scheme	Hardware RAID Scheme	Qualified NAS Devices Scheme	Oracle ASM Scheme	Cluster File System Scheme	Local Disks Scheme
Installation binary files	-	-	+	-	+	+
Configuration files	-	-	+	-	+	+
System parameter file (`SPFILE`)	-	-	+	+	+	-
Alert files	-	-	+	-	+	+
Trace files	-	-	+	-	+	+
Data files	+	+	+	+	-	-
Control files	+	+	+	+	-	-
Online redo log files	+	+	+	+	-	-
Archived redo log files	-	-	+	+	+	-
Flashback log files	-	-	+	+	+	-
Recovery files[1]	-	-	+	+	-	-

[1] The fast recovery area cannot reside on a cluster file system because this set of files includes online redo logs.

TABLE 1–3 Storage Management Schemes for Oracle Grid Infrastructure Files

Oracle File Type	Solaris Volume Manager for Sun Cluster Scheme	Hardware RAID Scheme	Qualified NAS Devices Scheme	Oracle ASM Scheme	Cluster File System Scheme	Local Disks Scheme
Installation binary files	-	-	+	-	-	+
OCR files	+	+	+	+	+	-
Voting disk	+	+	+	+	+	-

For more information, see "Database Storage Options" in *Oracle Database Installation Guide 11g Release 2 (11.2) for Oracle Solaris*.

Storage Management Requirements for Oracle Grid Infrastructure

Oracle Grid Infrastructure binary installation files are supported on the following storage management schemes:

- Qualified network-attached storage (NAS) devices

- Local file systems

Oracle Grid Infrastructure Oracle cluster registry (OCR) and voting disks are supported on the following storage management schemes:

- Solaris Volume Manager for Sun Cluster
- Hardware redundant array of independent disks (RAID) support
- Qualified network-attached storage (NAS) devices
- Cluster file systems
- Oracle ASM

Storage Management Requirements for the Oracle RAC Database

You can use the following storage management schemes for the Oracle RAC database:

- Solaris Volume Manager for Sun Cluster
- Hardware redundant array of independent disks (RAID) support
- Qualified network-attached storage (NAS) devices
- Oracle Automatic Storage Management (Oracle ASM)

Storage Management Requirements for Oracle Binary Files and Oracle Configuration Files

You can install the Oracle binary files and Oracle configuration files on one of the following locations.

- The local disks of each cluster node. See "Using Local Disks for Oracle Binary Files and Oracle Configuration Files" on page 24 for additional information.
- A shared file system from the following list:

Note – Oracle Grid Infrastructure binaries cannot reside on a cluster file system.

  - A PxFS-based cluster file system
  - A file system on a qualified NAS device

  See "Using a Shared File System for Oracle Binary Files and Oracle Configuration Files" on page 25 for additional information.

Using Local Disks for Oracle Binary Files and Oracle Configuration Files

Placing the Oracle binary files and Oracle configuration files on the individual cluster nodes enables you to upgrade the Oracle application later without shutting down the data service.

Note – Some versions of Oracle Database software require you to shut down the data service during an upgrade. To determine whether you can upgrade the Oracle application without shutting down the data service, see your Oracle Database documentation.

The disadvantage is that you then have several copies of the Oracle application binary files and Oracle configuration files to maintain and administer.

Using a Shared File System for Oracle Binary Files and Oracle Configuration Files

To simplify the maintenance of your Oracle RAC installation, you can install the Oracle binary files and Oracle configuration files on a shared file system.

Note – Oracle Grid Infrastructure binaries cannot reside on a cluster file system.

The following shared file systems are supported:

- A PxFS-based cluster file system

  If you use a PxFS-based cluster file system, use Solaris Volume Manager.
- A file system on a qualified NAS device

If you put the Oracle binary files and Oracle configuration files on a shared file system, you have only one copy to maintain and manage. However, you must shut down the data service in the entire cluster to upgrade the Oracle application. If a short period of downtime for upgrades is acceptable, place a single copy of the Oracle binary files and Oracle configuration files on a shared file system.

Storage Management Schemes Supported by Zone Clusters

You can use the following storage management schemes for running Oracle RAC in a zone cluster, depending on the version of Oracle RAC you are running.

- Solaris Volume Manager for Sun Cluster
- Oracle ASM
- A file system on a qualified NAS device with fencing

SPARC: Processor Architecture Requirements for Oracle Components

Before you decide which architecture to use for the Oracle relational database management system (RDBMS), note the following points.

- The architecture of both Oracle components must match.

- If you have 32-bit architecture for your Oracle components, you can boot the node on which the components reside in either 32-bit mode or 64-bit mode. However, if you have 64-bit architecture for your Oracle components, you must boot the node on which the components reside in 64-bit mode.
- You must use the same architecture when you boot all the nodes. For example, if you boot one node to use 32-bit architecture, you must boot all the nodes to use 32-bit architecture.

Using Oracle Data Guard With Support for Oracle RAC

You can use Support for Oracle RAC with Oracle Data Guard. To configure Support for Oracle RAC with Oracle Data Guard, perform the tasks in this guide. The tasks for clusters that are to be used in an Oracle Data Guard configuration are identical to the tasks for a standalone cluster.

For information about the installation, administration, and operation of Oracle Data Guard, see your Oracle documentation.

Preparing the Oracle Solaris Cluster Nodes

Preparing the Oracle Solaris Cluster nodes modifies the configuration of the operating system to enable Oracle RAC to run on Oracle Solaris Cluster nodes. Preparing the Oracle Solaris Cluster nodes and disks involves the following tasks:

- Bypassing the NIS name service
- Creating the database administrator (DBA) group and the DBA user accounts
- Configuring shared memory for the Oracle RAC software

Caution – Perform these tasks on all nodes where Support for Oracle RAC can run. If you do not perform these tasks on all nodes, the Oracle installation is incomplete. An incomplete Oracle installation causes Support for Oracle RAC to fail during startup.

To enable the Oracle RAC to run in a zone cluster, you need to perform the following additional tasks:

- Configuring shared memory for Oracle RAC software in a zone cluster
- Setting necessary privileges for Oracle RAC software in a zone cluster
- Configuring logical hostname resources for Oracle RAC software in a zone cluster

This section contains the following information:

- "Before You Begin" on page 27
- "How to Bypass the NIS Name Service" on page 27
- "How to Create the DBA Group and the DBA User Accounts" on page 28

Before You Begin

Before you prepare the Oracle Solaris Cluster nodes, ensure that all preinstallation tasks for Oracle RAC are completed. For more information, see your Oracle RAC documentation.

▼ How to Bypass the NIS Name Service

If Support for Oracle RAC refers to the NIS name service, unavailability of the name service might cause the Support for Oracle RAC data service to fail.

Bypassing the NIS name service ensures that the Support for Oracle RAC data service does not refer to the NIS name service when the data service sets the user identifier (ID). The Support for Oracle RAC data service sets the user ID when the data service starts or stops the database.

1 Become superuser on all nodes where Support for Oracle RAC can run.

2 On each node, ensure that the following entries in the `/etc/nsswitch.conf` file list `files` before `nis`.

```
passwd:    files nis
publickey: files nis
project:   files nis
group:     files nis
```

- Use the following command to display each lookup:

  ```
  # svccfg -s svc:/system/name-service/switch listprop config/lookupname
  ```

- To change a lookup entry, use the following command:

  ```
  # svccfg -s svc:/system/name-service/switch \
  setprop config/lookupname = astring: \"lookup-entry\"
  ```

For more information , see the `svccfg(1M)` and `nsswitch.conf(4)` man pages.

Example 1–1 Setting a Name Service Lookup Entry

The following example sets the lookup order for the `passwd` database to have `files` before `nis` and displays the current setting.

```
# svccfg -s svc:/sysgtem/name-service/switch \
setprop config/password = astring: \"files nis\"

# svccfg -s svc:/system/name-service/switch listprop config/password
config/password astring "files nis"
```

Next Steps Go to "How to Create the DBA Group and the DBA User Accounts" on page 28.

▼ How to Create the DBA Group and the DBA User Accounts

In installations of Oracle RAC with Oracle Solaris Cluster software, the DBA group is normally named `dba`. This group normally contains the `root` user and the `oracle` user.

Note – This configuration of users and groups differs from the configuration that is described in the Oracle documentation for a standalone installation of Oracle RAC. A standalone installation of Oracle RAC uses a primary DBA group that is named `oinstall` and a secondary group that is named `dba`. Some applications also require a secondary group that is named `oper`. For more information, see your Oracle documentation.

Perform this task on each cluster node.

1 On the cluster node where you are performing this task, become superuser.

2 Add an entry for the DBA group and potential users in the group to the `/etc/group` file.

`# groupadd -g` *group-id group-name*

group-name
Specifies the name of the group for which you are adding an entry. This group is normally named `dba`.

group-id
Specifies the group's unique numerical ID (GID) within the system.

Ensure that the command is identical on each node that can run Support for Oracle RAC.

You can create the name service entries in a network name service, such as the Network Information Service (NIS) or NIS+, so that the information is available to the data service clients. You can also create entries in the local `/etc` files to eliminate dependency on the network name service.

3 Create the home directory of each potential user in the DBA group that you defined in Step 2.

You are not required to create a home directory for the `root` user.

For each potential user whose home directory you are creating, type the following command:

```
# mkdir -p user-home
```

user-home
Specifies the full path of the home directory that you are creating.

4 Add each potential user in the DBA group that you defined in Step 2 to the system.

You are not required to add the `root` user.

Use the `useradd` command to add each user. Adding a user to the system adds an entry for the user to the following files:

- `/etc/passwd`
- `/etc/shadow`

```
# useradd -u user-id -g group-name -d user-home \
[ -s user-shell] user-name
```

`-u` *user-id*
Specifies the user's unique numerical ID (UID) within the system.

`-g` *group-name*
Specifies the name of the user group of which the user is a member. You must specify the DBA group that you defined in Step 2.

`-d` *user-home*
Specifies the full path of the user's home directory. You must specify the home directory that you created for the user in Step 3.

`-s` *user-shell*
Optionally specifies the full path name of the program that is to be used as the user's shell when the user logs in. If you omit the `-s` option, the system uses the `/bin/sh` program by default. If you specify the `-s` option, *user-shell* must specify a valid executable file.

user-name
Specifies the user name of the user that you are adding. You must specify the name of a potential user in the DBA group that you defined in Step 2.

Ensure that each user is identical on each node that can run Support for Oracle RAC.

5 Set the password of each user that you added in Step 4.

Use the `passwd` command to set the password of each user.

a. Type the following command:

```
# password user-name
```

user-name
Specifies the user name of the user whose password you are setting. You must specify the name of a user in the DBA group that you added in Step 4.

The `passwd` command prompts you for the password.

b. In response to the prompt, type the password and press Return.

The `passwd` command prompts you to retype the password.

c. In response to the prompt, retype the password and press Return.

6 Change the ownership of each home directory that you created in Step 3 as follows:

- Owner: the user for which you created the home directory
- Group: the DBA group that you defined in Step 2

For each home directory for which you are changing ownership, type the following command:

```
# chown user-name:group-name user-home
```

user-name
Specifies the user name of the user for whose home directory you are changing ownership. You must specify the name of a user in the DBA group that you added in Step 4.

group-name
Specifies the name of the user group of which the user is a member. You must specify the DBA group that you defined in Step 2.

user-home
Specifies the full path of the user's home directory. You must specify the home directory that you created for the user in Step 3.

7 Create a subdirectory of the `/var/opt` directory for each user in the DBA group that you added in Step 4.

For each subdirectory that you are creating, type the following command:

```
# mkdir /var/opt/user-name
```

user-name
Specifies the user name of the user whose subdirectory of the `/var/opt` directory you are creating. You must specify the name of a user in the DBA group that you added in Step 4.

8 Change the ownership of each directory that you created in Step 7 as follows:

- Owner: the user for which you created the directory
- Group: the DBA group that you defined in Step 2

For each directory for which you are changing ownership, type the following command:

```
# chown user-name:group-name /var/opt/user-name
```

user-name
Specifies the user name of the user for whose home directory you are changing ownership. You must specify the name of a user in the DBA group that you added in Step 4.

group-name
Specifies the name of the user group of which the user is a member. You must specify the DBA group that you defined in Step 2.

Example 1–2 Creating the DBA Group and the DBA User Accounts

This example shows the sequence of commands for creating the DBA group `dba`, which is to contain the users `root` and `oracle`.

The `dba` group and the `oracle` user are created as follows:

- The GID of the `dba` group is 520.
- The home directory of the `oracle` user is `/Oracle-home`.
- The UID of the `oracle` user is 120.
- The `oracle` user's login shell is the Bash shell.

```
# groupadd -g 520 dba
# mkdir /Oracle-home
# useradd -u 120 -g dba -d /Oracle-home -s /bin/bash oracle
# passwd oracle
New Password:oracle
Re-enter new Password:oracle
passwd: password successfully changed for oracle
# chown oracle:dba /Oracle-home
# mkdir /var/opt/oracle
# chown oracle:dba /var/opt/oracle
```

See Also The following man pages:

- `passwd(1)`
- `useradd(1M)`
- `group(4)`
- `passwd(4)`
- `shadow(4)`

Next Steps Go to "How to Configure Shared Memory for Oracle RAC Software in the Global Cluster" on page 31.

▼ How to Configure Shared Memory for Oracle RAC Software in the Global Cluster

To enable the Oracle RAC software to run correctly, you must ensure that sufficient shared memory is available on all the cluster nodes. Perform this task on each cluster node.

1 Become superuser on a cluster node.

2 **Update the shared memory configuration information.**

See "Configuring Kernel Parameters in Oracle Solaris 10" in *Oracle Database Installation Guide 11g Release 2 (11.2) for Oracle Solaris*. This information is valid for Oracle Solaris Cluster 4.0 software on the Oracle Solaris 11 OS.

You must configure these parameters on the basis of the resources that are available in the cluster. However, the value of each parameter must be sufficient to enable the Oracle RAC software to create a shared memory segment that conforms to its configuration requirements.

3 **Shut down and reboot each node whose shared memory configuration information you updated in Step 2.**

For detailed instructions, see "Shutting Down and Booting a Single Node in a Cluster" in *Oracle Solaris Cluster System Administration Guide*.

Next Steps If you are using zone clusters, go to "How to Configure Shared Memory for Oracle RAC Software in a Zone Cluster" on page 32.

If not, go to "Installing the Support for Oracle RAC Package" on page 35.

▼ How to Configure Shared Memory for Oracle RAC Software in a Zone Cluster

To configure shared memory for the Oracle RAC software in a zone cluster, perform the following task.

Before You Begin Ensure that shared memory is configured in the global cluster. See "How to Configure Shared Memory for Oracle RAC Software in the Global Cluster" on page 31.

1 **Configure shared memory in each zone cluster.**

For the minimum required value of each parameter and the procedures to set the values, see your Oracle Clusterware and Oracle Database documentation.

Note – These steps do not affect the actual shared memory control for the zone cluster. You perform these steps to help the Oracle `dbca` utility enable you to set the database memory allocation. If the Oracle `dbca` utility is not used for Oracle RAC database creation, you can skip these steps in the zone cluster.

---

2 **Perform the following steps, if you want to limit the memory used for the zone cluster.**

a. **Become superuser on the global cluster node that hosts the zone cluster.**

b. Configure the `capped-memory` property attributes `physical`, `swap`, and `locked` by using the `clzonecluster` command.

```
#clzonecluster configure zcname
clzonecluster:zcname> add capped-memory
clzonecluster:cz1-2n:capped-memory> set physical=memsize
clzonecluster:cz1-2n:capped-memory> set swap=memsize
clzonecluster:cz1-2n:capped-memory> set locked=memsize
clzonecluster:cz1-2n:capped-memory> end
clzonecluster:cz1-2n>commit
```

`Physical`=*memsize*
Specifies the physical memory size.

`swap`=*memsize*
Specifies the swap memory size.

`locked`=*memsize*
Specifies the limit of the shared memory segment size that the Oracle RAC database processes can request to lock in memory.

Note – In addition to the `locked` attribute of the `capped-memory` property, you can use the `max-shm-memory` property to directly configure the limit of the shared memory segment in a zone cluster. See also the `zonecfg(1M)` man page .

---

c. Reboot the zone cluster.

```
#clzonecluster reboot zcname
```

Note – You can perform the step to configure the `capped-memory` property attributes as part of the zone cluster creation. If you configure the `capped-memory` property attributes as part of the zone cluster creation, the memory-related properties immediately take effect after the first zone cluster boot. See "How to Create a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

Next Steps Go to "How to Set the Necessary Privileges for Oracle RAC Software in a Zone Cluster" on page 33.

▼ How to Set the Necessary Privileges for Oracle RAC Software in a Zone Cluster

You should set the necessary privileges to enable Oracle RAC to run in a zone cluster configuration. You can use the `clzonecluster` command to include the necessary privileges in a zone cluster configuration by setting the `limitpriv` property. Perform the following steps to set the necessary privileges in a zone cluster to run Oracle RAC.

1. **Become superuser on the global cluster node that hosts the zone cluster.**

2. **Configure the `limitpriv` property by using the `clzonecluster` command.**

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

3. **Reboot the zone cluster.**

```
# clzonecluster reboot zcname
```

Note – You can perform the step to configure the `limitpriv` property as part of the zone cluster creation. For information on creating a zone cluster, see "How to Create a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

Next Steps Go to "How to Configure the Logical Hostname Resources or Virtual IP Addresses for Oracle RAC Software in a Zone Cluster" on page 34.

▼ How to Configure the Logical Hostname Resources or Virtual IP Addresses for Oracle RAC Software in a Zone Cluster

To support Oracle Grid Infrastructure virtual IP resources in Oracle RAC configurations in zone clusters, you should configure the failover-capable hostnames or IP addresses used by those resources in a given zone cluster using the `clzonecluster` command.

Perform the following steps to configure the virtual IP addresses in a zone cluster configuration for Oracle RAC.

1. **Become superuser on the global cluster node that hosts the zone cluster.**

2. **Configure the virtual IP addresses using the `clzonecluster` command.**

```
# clzonecluster configure zcname
clzonecluster:zcname>add net
clzonecluster:zcname:net>set address=racnode1-vip
clzonecluster:zcname:net>end
clzonecluster:zcname>add net
clzonecluster:zcname:net>set address=racnode2-vip
clzonecluster:zcname:net>end
clzonecluster:zcname>commit
```

Next Steps Go to "Installing the Support for Oracle RAC Package" on page 35.

Installing the Support for Oracle RAC Package

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

▼ How to Install the Support for Oracle RAC Package

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

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

2 Ensure that the `solaris` and `ha-cluster` publishers are valid.

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

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

3 Install the Support for Oracle RAC software package.

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

4 Verify that the package installed successfully.

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

Installation is successful if output shows that `State` is `Installed`.

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

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

◆ ◆ ◆ CHAPTER 2

Configuring Storage for Oracle Files

This chapter explains how to configure storage for Oracle files.

Summary of Configuration Tasks for Storage for Oracle Files

This section summarizes the following tasks for configuring each storage management scheme for Oracle files:

Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Files

The following tables summarize the tasks for configuring Solaris Volume Manager for Sun Cluster and provides cross-references to detailed instructions for performing the tasks.

Perform these tasks in the order in which they are listed in the table.

TABLE 2–1 Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Files in the Global Cluster

Task	Instructions
Configure Solaris Volume Manager for Sun Cluster	"Using Solaris Volume Manager for Sun Cluster" on page 42
Register and configure the multiple-owner volume manager resource group	If you are using the `clsetup` utility for this task, see "How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group by Using `clsetup`" on page 57. If you are using Oracle Solaris Cluster maintenance commands for this task, see "How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands" on page 173.
Create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Oracle RAC database	"How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60
Register and configure storage resources for Oracle files	If you are using the `clsetup` utility for this task, see "Registering and Configuring Storage Resources for Oracle Files" on page 65. If you are using the Oracle Solaris Cluster maintenance commands for this task, see "Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands" on page 179.

TABLE 2–2 Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Files in a Zone Cluster

Task	Instructions
Configure Solaris Volume Manager for Sun Cluster in the global cluster	"Using Solaris Volume Manager for Sun Cluster" on page 42
Register and configure the multiple-owner volume manage framework resource group in the global cluster	If you are using the `clsetup` utility for this task, see "How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group by Using `clsetup`" on page 57. If you are using Oracle Solaris Cluster maintenance commands for this task, see "How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands" on page 173.

TABLE 2–2 Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Files in a Zone Cluster *(Continued)*

Task	Instructions
Create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Oracle RAC database in the global cluster	"How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60
Configure Solaris Volume Manager devices in a zone cluster	See "How to Add a Disk Set to a Zone Cluster (Solaris Volume Manager)" in *Oracle Solaris Cluster Software Installation Guide*
Register and configure storage resources for Oracle files in the zone cluster	If you are using the `clsetup` utility for this task, see "Registering and Configuring Storage Resources for Oracle Files" on page 65. If you are using the Oracle Solaris Cluster maintenance commands for this task, see "Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands" on page 179.

Tasks for Configuring Hardware RAID Support for Oracle Files

The following table summarizes the tasks for configuring hardware RAID support and provides cross-references to detailed instructions for performing the tasks.

TABLE 2–3 Tasks for Configuring Hardware RAID Support for Oracle Files

Task	Instructions
Configure hardware RAID support	"Using Hardware RAID Support" on page 44

Note – For information configuring hardware RAID for a zone cluster, see "Adding Storage Devices to a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

Tasks for Configuring Oracle ASM for Oracle Files

The following table summarizes the tasks for configuring Oracle ASM and provides cross-references to detailed instructions for performing the tasks.

TABLE 2–4 Tasks for Configuring Oracle ASM for Oracle Files

Task	Instructions
Configure devices for Oracle ASM	"Using Oracle ASM" on page 46

Note – For information about configuring Oracle ASM for a zone cluster, see "Adding Storage Devices to a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

Tasks for Configuring Qualified NAS Devices for Oracle Files

The following table summarizes the tasks for configuring qualified NAS devices and provides cross-references to detailed instructions for performing the tasks. NAS devices are supported in both global and zone clusters.

Perform these tasks in the order in which they are listed in the table.

TABLE 2–5 Tasks for Configuring Qualified NAS Devices for Oracle Files

Task	Instructions
Install and configure the qualified NAS device	See the *Oracle Solaris Cluster 4.0 With Network-Attached Storage Device Manual*.
Register and configure the Oracle RAC framework resource group in a global cluster or zone cluster	If you are using the `clsetup` utility for this task, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51. If you are using the Oracle Solaris Cluster maintenance commands for this task, see "How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands" on page 173.
Register and configure storage resources for Oracle files, including Oracle RAC to support NAS NFS	If you are using the `clsetup` utility for this task, see "Registering and Configuring Storage Resources for Oracle Files" on page 65. If you are using the Oracle Solaris Cluster maintenance commands for this task, see "Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands" on page 179.

Tasks for Configuring a Cluster File System for Oracle Files

The following table summarizes the tasks for configuring a PxFS-based cluster file system, and provides cross-references to detailed instructions for performing the tasks.

Perform these tasks in the order in which they are listed in the table.

TABLE 2–6 Tasks for Configuring a PxFS-Based Cluster File System for Oracle Files

Task	Instructions
Install and configure the cluster file system	"Using a Cluster File System" on page 48
Register and configure the Oracle RAC framework resource group	If you are using the `clsetup` utility for this task, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51. If you are using the Oracle Solaris Cluster maintenance commands for this task, see "How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands" on page 173.

Note – A PxFS-based cluster file system is currently not supported for Oracle RAC in zone clusters.

Installing Storage Management Software With Support for Oracle RAC

Install the software for the storage management schemes that you are using for Oracle files. For more information, see "Storage Management Requirements" on page 22.

Note – For information about how to install and configure qualified NAS devices with Support for Oracle RAC, see *Oracle Solaris Cluster 4.0 With Network-Attached Storage Device Manual*.

This section contains the following information:

Using Solaris Volume Manager for Sun Cluster

Always install Solaris Volume Manager software, which includes the Solaris Volume Manager for Sun Cluster feature, in the global cluster, even when supporting zone clusters. Solaris

Volume Manager software is not automatically installed as part of an Oracle Solaris 11 software installation. You must install it manually by using the following command:

```
# pkg install system/svm
```

The `clzonecluster` command configures Solaris Volume Manager for Sun Cluster devices from the global-cluster voting node into the zone cluster. All administration tasks for Solaris Volume Manager for Sun Cluster are performed in the global-cluster voting node, even when the Solaris Volume Manager for Sun Cluster volume is used in a zone cluster.

When an Oracle RAC installation inside a zone cluster uses a file system that exists on top of a Solaris Volume Manager for Sun Cluster volume, you should still configure the Solaris Volume Manager for Sun Cluster volume in the global cluster. In this case, the scalable device group resource belongs to this zone cluster.

When an Oracle RAC installation inside a zone cluster runs directly on the Solaris Volume Manager for Sun Cluster volume, you must first configure the Solaris Volume Manager for Sun Cluster in the global cluster and then configure the Solaris Volume Manager for Sun Cluster volume into the zone cluster. In this case, the scalable device group belongs to this zone cluster.

For information about the types of Oracle files that you can store by using Solaris Volume Manager for Sun Cluster, see "Storage Management Requirements" on page 22.

How to Use Solaris Volume Manager for Sun Cluster

To use the Solaris Volume Manager for Sun Cluster software with Support for Oracle RAC, perform the following tasks. Solaris Volume Manager for Sun Cluster is installed during the installation of the Solaris Operating System.

1. **Configure the Solaris Volume Manager for Sun Cluster software on the global-cluster nodes.**

   For information about configuring Solaris Volume Manager for Sun Cluster in the global cluster, see "Configuring Solaris Volume Manager Software" in *Oracle Solaris Cluster Software Installation Guide*.

2. **If you are using a zone cluster, configure the Solaris Volume Manager for Sun Cluster volume into the zone cluster.**

   For information on configuring Solaris Volume Manager for Sun Cluster volume into a zone cluster, see "How to Add a Disk Set to a Zone Cluster (Solaris Volume Manager)" in *Oracle Solaris Cluster Software Installation Guide*.

Next Steps Ensure that all other storage management schemes that you are using for Oracle files are installed.

After all storage management schemes that you are using for Oracle files are installed, go to Chapter 3, "Registering and Configuring the Resource Groups."

Using Hardware RAID Support

For information about the types of Oracle files that you can store by using hardware RAID support, see "Storage Management Requirements" on page 22.

Oracle Solaris Cluster software provides hardware RAID support for several storage devices. To use this combination, configure raw device identities (`/dev/did/rdsk*`) on top of the disk arrays' logical unit numbers (LUNs). To set up the raw devices for Oracle RAC on a cluster that uses StorEdge SE9960 disk arrays with hardware RAID, perform the following task.

▼ How to Use Hardware RAID Support

1 Create LUNs on the disk arrays.

See the Oracle Solaris Cluster hardware documentation for information about how to create LUNs.

2 After you create the LUNs, run the `format(1M)` command to partition the disk arrays' LUNs into as many slices as you need.

The following example lists output from the `format` command.

```
# format

0. c0t2d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
   /sbus@3,0/SUNW,fas@3,8800000/sd@2,0
1. c0t3d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
   /sbus@3,0/SUNW,fas@3,8800000/sd@3,0
2. c1t5d0 <Symbios-StorEDGEA3000-0301 cyl 21541 alt 2 hd 64 sec 64>
   /pseudo/rdnexus@1/rdriver@5,0
3. c1t5d1 <Symbios-StorEDGEA3000-0301 cyl 21541 alt 2 hd 64 sec 64>
   /pseudo/rdnexus@1/rdriver@5,1
4. c2t5d0 <Symbios-StorEDGEA3000-0301 cyl 21541 alt 2 hd 64 sec 64>
   /pseudo/rdnexus@2/rdriver@5,0
5. c2t5d1 <Symbios-StorEDGEA3000-0301 cyl 21541 alt 2 hd 64 sec 64>
   /pseudo/rdnexus@2/rdriver@5,1
6. c3t4d2 <Symbios-StorEDGEA3000-0301 cyl 21541 alt 2 hd 64 sec 64>
   /pseudo/rdnexus@3/rdriver@4,2
```

Note – To prevent a loss of disk partition information, do not start the partition at cylinder 0 for any disk slice that is used for raw data. The disk partition table is stored in cylinder 0 of the disk.

3 Determine the raw device identity (DID) that corresponds to the LUNs that you created in Step 1.

Use the `cldevice(1CL)` command for this purpose.

The following example lists output from the `cldevice list -v` command.

```
# cldevice list -v

DID Device     Full Device Path
----------     ----------------
```

```
d1              phys-schost-1:/dev/rdsk/c0t2d0
d2              phys-schost-1:/dev/rdsk/c0t3d0
d3              phys-schost-2:/dev/rdsk/c4t4d0
d3              phys-schost-1:/dev/rdsk/c1t5d0
d4              phys-schost-2:/dev/rdsk/c3t5d0
d4              phys-schost-1:/dev/rdsk/c2t5d0
d5              phys-schost-2:/dev/rdsk/c4t4d1
d5              phys-schost-1:/dev/rdsk/c1t5d1
d6              phys-schost-2:/dev/rdsk/c3t5d1
d6              phys-schost-1:/dev/rdsk/c2t5d1
d7              phys-schost-2:/dev/rdsk/c0t2d0
d8              phys-schost-2:/dev/rdsk/c0t3d0
```

In this example, the `cldevice` output identifies that the raw DID that corresponds to the disk arrays' shared LUNs is `d4`.

4 Obtain the full DID device name that corresponds to the DID device that you identified in Step 3.

The following example shows the output from the `cldevice show` for the DID device that was identified in the example in Step 3. The command is run from node `phys-schost-1`.

```
# cldevice show d4

=== DID Device Instances ===

DID Device Name:                                /dev/did/rdsk/d4
  Full Device Path:                               phys-schost-1:/dev/rdsk/c2t5d0
  Replication:                                    none
  default_fencing:                                global
```

5 If you are using a zone cluster configure the DID devices into the zone cluster. Otherwise, proceed to Step 6.

For information about configuring DID devices into a zone cluster, see "How to Add a DID Device to a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

6 Create or modify a slice on each DID device to contain the disk-space allocation for the raw device.

Use the `format(1M)` command, `fmthard(1M)` command, or `prtvtoc(1M)` for this purpose. Specify the full device path from the node where you are running the command to create or modify the slice.

For example, if you choose to use slice `s0`, you might choose to allocate 100 GB of disk space in slice `s0`.

7 Change the ownership and permissions of the raw devices that you are using to allow access to these devices.

To specify the raw device, append `s`*N* to the DID device name that you obtained in Step 4, where *N* is the slice number.

For example, the `cldevice` output in Step 4 identifies that the raw DID that corresponds to the disk is `/dev/did/rdsk/d4`. If you choose to use slice `s0` on these devices, specify the raw device `/dev/did/rdsk/d4s0`.

Next Steps Ensure that all other storage management schemes that you are using for Oracle files are installed.

After all storage management schemes that you are using for Oracle files are installed, go to Chapter 3, "Registering and Configuring the Resource Groups."

Using Oracle ASM

Use Oracle ASM with one storage management scheme from the following list:

- **Hardware RAID.** For more information, see "How to Use Oracle ASM With Hardware RAID" on page 46.
- **Solaris Volume Manager for Sun Cluster.** For more information, see "How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60.

For information about the types of Oracle files that you can store by using Oracle ASM, see "Storage Management Requirements" on page 22.

Note – When an Oracle RAC installation in a zone cluster uses Oracle ASM, you must configure all the devices needed by that Oracle RAC installation into that zone cluster by using the `clzonecluster` command. When Oracle ASM runs inside a zone cluster, the administration of Oracle ASM occurs entirely within the same zone cluster.

▼ How to Use Oracle ASM With Hardware RAID

1 On a cluster member, log in as root or become superuser.

2 Determine the identities of device identity (DID) devices that correspond to shared disks that are available in the cluster.

Use the `cldevice(1CL)` command for this purpose.

The following example shows an extract from output from the `cldevice list -v` command.

```
# cldevice list -v
DID Device          Full Device Path
----------          ----------------
...
d5                  phys-schost-3:/dev/rdsk/c3t216000C0FF084E77d0
d5                  phys-schost-1:/dev/rdsk/c5t216000C0FF084E77d0
d5                  phys-schost-2:/dev/rdsk/c4t216000C0FF084E77d0
d5                  phys-schost-4:/dev/rdsk/c2t216000C0FF084E77d0
d6                  phys-schost-3:/dev/rdsk/c4t216000C0FF284E44d0
d6                  phys-schost-1:/dev/rdsk/c6t216000C0FF284E44d0
d6                  phys-schost-2:/dev/rdsk/c5t216000C0FF284E44d0
d6                  phys-schost-4:/dev/rdsk/c3t216000C0FF284E44d0
...
```

In this example, DID devices `d5` and `d6` correspond to shared disks that are available in the cluster.

3 **Obtain the full DID device name for each DID device that you are using for the Oracle ASM disk group.**

The following example shows the output from the `cldevice show` for the DID devices that were identified in the example in Step 2. The command is run from node `phys-schost-1`.

```
# cldevice show d5 d6

=== DID Device Instances ===

DID Device Name:                                /dev/did/rdsk/d5
  Full Device Path:                             phys-schost-1:/dev/rdsk/c5t216000C0FF084E77d0
  Replication:                                      none
  default_fencing:                                global

DID Device Name:                                 /dev/did/rdsk/d6
  Full Device Path:                              phys-schost-1:/dev/rdsk/c6t216000C0FF284E44d0
  Replication:                                      none
  default_fencing:                                  global
```

4 **If you are using a zone cluster, configure the DID devices into the zone cluster. Otherwise, proceed to Step 5.**

For information about configuring DID devices in a zone cluster, see "How to Add a DID Device to a Zone Cluster" in *Oracle Solaris Cluster Software Installation Guide*.

5 **Create or modify a slice on each DID device to contain the disk-space allocation for the Oracle ASM disk group.**

Use the `format`(1M) command, `fmthard`(1M) command, or `prtvtoc`(1M) for this purpose. Specify the full device path from the node where you are running the command to create or modify the slice.

For example, if you choose to use slice `s0` for the Oracle ASM disk group, you might choose to allocate 100 Gbytes of disk space in slice `s0`.

6 **Prepare the raw devices that you are using for Oracle ASM.**

a. **Change the ownership and permissions of each raw device that you are using for Oracle ASM, to allow access by Oracle ASM to these devices.**

Note – If Oracle ASM on hardware RAID is configured for a zone cluster, perform this step in that zone cluster.

To specify the raw device, append s*X* to the DID device name that you obtained in Step 3, where *X* is the slice number.

```
# chown oraasm:oinstall /dev/did/rdsk/dNsX
# chmod 660 /dev/disk/rdsk/dNsX
# ls -lhL /dev/did/rdsk/dNsX
crw-rw----  1 oraasm  oinstall  239, 128 Jun 15 04:38 /dev/did/rdsk/dNsX
```

For more information about changing the ownership and permissions of raw devices for use by Oracle ASM, see your Oracle documentation.

b. Clean out the disk headers for each raw device that you are using for Oracle ASM.

```
# dd if=/dev/zero of=/dev/did/rdsk/dNsX bs=1024k count=200
2000+0 records in
2000+0 records out
```

7 Modify the `ASM_DISKSTRING` Oracle ASM instance-initialization parameter to specify the devices that you are using for the Oracle ASM disk group.

For example, to use the `/dev/did/` path for the Oracle ASM disk group, add the value `/dev/did/rdsk/d*` to the `ASM_DISKSTRING` parameter. If you are modifying this parameter by editing the Oracle initialization parameter file, edit the parameter as follows:

```
ASM_DISKSTRING = '/dev/did/rdsk/*'
```

For more information, see your Oracle documentation.

Next Steps Ensure that all other storage management schemes that you are using for Oracle files are installed.

After all storage management schemes that you are using for Oracle files are installed, go to Chapter 3, "Registering and Configuring the Resource Groups."

Using a Cluster File System

Oracle RAC is supported on cluster file systems:

- Cluster file systems use the Oracle Solaris Cluster Proxy File System (PxFS)

  For general information about how to create and mount PxFS-based cluster file systems, see the following documentation:

  - "Planning the Global Devices, Device Groups, and Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide*
  - "Creating Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide*

For information that is specific to the use of cluster file systems with Support for Oracle RAC, see the subsections that follow.

- "Types of Oracle Files That You Can Store on a PxFS-Based Cluster File System" on page 49

- "Optimizing Performance and Availability When Using a PxFS-Based Cluster File System" on page 49
- "How to Use a PxFS-Based Cluster File System" on page 50

Types of Oracle Files That You Can Store on a PxFS-Based Cluster File System

You can store only these files that are associated with Oracle RAC on a PxFS-based cluster file system:

- Oracle RDBMS binary files
- Oracle Grid Infrastructure binary files

Note – Oracle Grid Infrastructure binaries cannot reside on a cluster file system.

- Oracle configuration files (for example, `init.ora`, `tnsnames.ora`, `listener.ora`, and `sqlnet.ora`)
- System parameter file (`SPFILE`)
- Alert files (for example, `alert_`*sid*`.log`)
- Trace files (`*.trc`)
- Archived redo log files
- Flashback log files
- Oracle cluster registry (OCR) files
- Oracle Grid Infrastructure voting disk

Note – You *must not* store data files, control files, online redo log files, or Oracle recovery files on a PxFS-based cluster file system.

Optimizing Performance and Availability When Using a PxFS-Based Cluster File System

The I/O performance during the writing of archived redo log files is affected by the location of the device group for archived redo log files. For optimum performance, ensure that the primary of the device group for archived redo log files is located on the same node as the Oracle RAC database instance. This device group contains the file system that holds archived redo log files of the database instance.

To improve the availability of your cluster, consider increasing the desired number of secondary nodes for device groups. However, increasing the desired number of secondary nodes for device groups might also impair performance. To increase the desired number of

secondary nodes for device groups, change the `numsecondaries` property. For more information, see "Multiported Device Groups" in *Oracle Solaris Cluster Concepts Guide*.

▼ How to Use a PxFS-Based Cluster File System

1 Create and mount the cluster file system.

See "Creating Cluster File Systems" in *Oracle Solaris Cluster Software Installation Guide* for information about how to create and mount the cluster file system.

Note – Oracle Grid Infrastructure binaries cannot reside on a cluster file system.

2 If you are using the UNIX file system (UFS), ensure that you specify the correct mount options for various types of Oracle files.

For the correct options, see the table that follows. You set these options when you add an entry to the `/etc/vfstab` file for the mount point.

File Type	Options
Oracle RDBMS binary files	`global,logging`
Oracle Grid Infrastructure binary files	`global,logging`
Oracle configuration files	`global,logging`
System parameter file (`SPFILE`)	`global,logging`
Alert files	`global,logging`
Trace files	`global,logging`
Archived redo log files	`global,logging,forcedirectio`
Flashback log files	`global,logging,forcedirectio`
OCR files	`global,logging,forcedirectio`
Oracle Grid Infrastructure voting disk	`global,logging,forcedirectio`

Next Steps Ensure that all other storage management schemes that you are using for Oracle files are installed.

After all storage management schemes that you are using for Oracle files are installed, go to Chapter 3, "Registering and Configuring the Resource Groups."

◆ ◆ ◆ CHAPTER 3

Registering and Configuring the Resource Groups

This chapter explains how to register and configure the resource groups that are used in an Oracle RAC configuration.

Registering and Configuring the Oracle RAC Framework Resource Group

Registering and configuring the Oracle RAC framework resource group enables Oracle RAC to run with Oracle Solaris Cluster software.

Note – You *must* register and configure the Oracle RAC framework resource group. Otherwise, Oracle RAC cannot run with Oracle Solaris Cluster software.

The Oracle RAC framework resource in the global-cluster voting node can support any Oracle RAC installation running in the global cluster. The Oracle RAC framework resource in a zone cluster supports the Oracle RAC installation running in that specific zone cluster. Multiple Oracle RAC framework resource groups can exist in a single Oracle Solaris Cluster configuration.

This section contains the following information about registering the Oracle RAC framework resource group:

- "Tools for Registering and Configuring the Oracle RAC Framework Resource Group" on page 52
- "How to Register and Configure the Oracle RAC Framework Resource Group by Using `clsetup`" on page 52

Tools for Registering and Configuring the Oracle RAC Framework Resource Group

Oracle Solaris Cluster software provides the following tools for registering and configuring the Oracle RAC framework resource group in the global cluster or in a zone cluster:

- **The `clsetup` utility.** For more information, see "How to Register and Configure the Oracle RAC Framework Resource Group by Using `clsetup`" on page 52.
- **Oracle Solaris Cluster maintenance commands.** For more information, see Appendix D, "Command-Line Alternatives."

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

Note – The `clsetup` utility runs only in a voting node of the global cluster.

▼ How to Register and Configure the Oracle RAC Framework Resource Group by Using clsetup

When you register and configure the Oracle RAC framework resource group for a cluster, the Oracle RAC framework resource group is created.

Perform this procedure during your initial setup of Support for Oracle RAC. Perform this procedure from one node only.

Before You Begin Ensure that the following prerequisites are met:

- All preinstallation tasks for Oracle RAC are completed.
- The Oracle Solaris Cluster nodes are prepared.
- The data services packages are installed.

Ensure that you have the following information:

- The names of the nodes where you require Support for Oracle RAC to run.

1 **Become superuser on any cluster node.**

2 **Start the `clsetup` utility.**

```
# clsetup
```

The `clsetup` main menu is displayed.

3 **Choose the menu item, Data Services.**

The Data Services menu is displayed.

4 **Choose the menu item, Oracle Real Application Clusters.**

The `clsetup` utility displays information about Support for Oracle RAC.

5 **Press Return to continue.**

The `clsetup` utility prompts you to select whether you are performing the initial configuration of Support for Oracle RAC or administering an existing configuration.

Note – The `clsetup` utility currently allows ongoing administration of an Oracle RAC framework running only in the global cluster. For ongoing administration of an Oracle RAC framework configured in a zone cluster, you need to instead use the Oracle Solaris Cluster maintenance commands.

---

6 **Choose the menu item, Oracle RAC Create Configuration.**

The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7 **Type the option number for the location of the Oracle RAC cluster and press Return.**

- **If you select the global cluster option, the `clsetup` utility displays the list of components of Oracle RAC to configure. Skip to Step 9.**
- **If you select the zone cluster option, the `clsetup` utility prompts you to select the required zone cluster. Proceed to Step 8.**

8 **Type the option number for the required zone cluster and press Return.**

The `clsetup` utility displays a list of components of Oracle RAC to configure.

9 **Choose the menu item, RAC Framework Resource Group .**

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

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

The `clsetup` utility displays a list of the cluster nodes on which the Support for Oracle RAC packages are installed.

11 Select the nodes where you require Support for Oracle RAC to run.

- **To accept the default selection of all listed nodes in an arbitrary order, type `a` and press Return.**
- **To select a subset of the listed nodes, type a comma-separated or space-separated list of the option numbers for the nodes you choose and press Return.**

  Ensure that the nodes are listed in the order in which the nodes are to appear in the Oracle RAC framework resource group's node list.
- **To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the option numbers for the nodes you choose and press Return.**

  Ensure that the nodes are listed in the order in which the nodes are to appear in the Oracle RAC framework resource group's node list.

12 If needed to end the selection of nodes, type `d` and press Return.

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

13 If you require a different name for any Oracle Solaris Cluster objects, change the name.

a. Type the option number for the object name to change 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 names of the Oracle Solaris Cluster objects that the utility will create.

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

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

15 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.

16 **Press Return to continue.**

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

17 **(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. If you choose to quit `clsetup`, the utility recognizes your existing Oracle RAC framework resource group when you restart the utility.

18 **Determine whether the Oracle RAC framework resource group and its resources are online.**

Use the `clresourcegroup(1CL)` utility for this purpose. By default, the `clsetup` utility assigns the name `rac-framework-rg` to the Oracle RAC framework resource group.

- **In the global cluster, type the following command.**

```
# clresourcegroup status rac-framework-rg
```

- **In a zone cluster, type the following command.**

```
# clresourcegroup status -Z zcname rac-framework-rg
```

19 **If the Oracle RAC framework resource group and its resources are *not* online, bring them online.**

- **In the global cluster, type the following command.**

```
# clresourcegroup online -emM rac-framework-rg
```

- **In a zone cluster, type the following command.**

```
# clresourcegroup online -emM -Z zcname rac-framework-rg
```

More Information

Resource Configuration

The following table lists the default resource configuration that the `clsetup` utility creates when you complete this task.

Resource Name, Resource Type, and Resource Group	Dependencies	Description
Resource type: `SUNW.rac_framework`	None.	Oracle RAC framework resource.
Resource name: `rac-framework-rs`		
Resource group: `rac-framework-rg`		

Next Steps The next step depends on the volume manager that you are using, as shown in the following table.

Volume Manager	Next Step
Solaris Volume Manager for Sun Cluster	"Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group" on page 56
None	"Registering and Configuring Storage Resources for Oracle Files" on page 65

Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group

Registering and configuring the multiple-owner volume-manager framework resource group enables Oracle RAC to manage multiple-owner volume-manager resources with Oracle Solaris Cluster software.

The multiple-owner volume manager framework resource in the global-cluster voting node supports any volume manager used by Oracle RAC anywhere on the machine, including the global cluster and all zone clusters.

This section contains the following information about registering the multiple-owner volume-manager framework resource group:

Tools for Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group

Oracle Solaris Cluster software provides the following tools for registering and configuring the multiple-owner volume manager framework resource group in the global cluster or in a zone cluster:

- **The `clsetup` utility.** For more information, see "How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group by Using `clsetup`" on page 57.
- **Oracle Solaris Cluster maintenance commands.** For more information, see Appendix D, "Command-Line Alternatives."

The `clsetup` utility provides a wizard for configuring resources for the multiple-owner volume manager framework resource group. The wizard reduces the possibility of configuration errors

that might result from command syntax errors or omissions. This wizard also ensures that all required resources are created and that all required dependencies between resources are set.

Note – The `clsetup` utility runs only in a voting node of the global cluster.

How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group by Using clsetup

When you register and configure the multiple-owner volume manager framework resource group for a cluster, the multiple-owner volume manager framework resource group is created.

Perform this procedure during your initial setup of Support for Oracle RAC. Perform this procedure from one node only.

Before You Begin

- Ensure that all storage management software that you intend to use is installed and configured on all nodes where Oracle RAC is to run.
- Ensure that you have the list of storage management schemes that you are using for Oracle files.

1. **Become superuser on any cluster node.**

2. **Start the `clsetup` utility.**

   ```
   # clsetup
   ```

   The `clsetup` main menu is displayed.

3. **Type the option number for Data Services and press Return.**

   The Data Services menu is displayed.

4. **Type the option number for Oracle Real Application Clusters and press Return.**

   The `clsetup` utility displays information about Support for Oracle RAC.

5. **Press Return to continue.**

   The `clsetup` utility prompts you to select whether you are performing the initial configuration of Support for Oracle RAC or administering an existing configuration.

Note – The `clsetup` utility currently allows ongoing administration of an Oracle RAC framework running only in the global cluster. For ongoing administration of an Oracle RAC framework configured in a zone cluster, use instead the Oracle Solaris Cluster maintenance commands.

---

6 Type the option number for Oracle RAC Create Configuration and press Return.

The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7 Type the option number for Global Cluster and press Return.

Note – Configure a multiple-owner volume-manager framework resource group only in the global cluster.

The `clsetup` utility displays the list of components of Oracle RAC to configure.

8 Type the option number for Multiple-Owner Volume-Manager Framework Resource Group and press Return.

The `clsetup` utility displays an overview about this task.

9 Press Return to continue.

The `clsetup` utility prompts you to select the multiple-owner volume managers to use. Only installed volume managers are listed.

10 Type the option numbers for the multiple-owner volume managers to use and press Return.

11 To confirm your selection of multiple-owner volume managers, type `d` and press Return.

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

12 If you require a different name for any Oracle Solaris Cluster objects, change the name.

a. Type the option number for the name to change 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 names of the Oracle Solaris Cluster objects that the utility will create.

Note – If, after you configure the multiple-owner volume manager framework resource group, you again run this wizard to configure another volume manager, the names of the multiple-owner volume manager framework resource group and its existing resources are not available to rename.

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

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

14 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.

15 Press Return to continue.

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

16 (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. If you choose to quit `clsetup`, the utility recognizes your existing multiple-owner volume-manager framework resource group when you restart the utility.

17 Determine whether the multiple-owner volume-manager framework resource group and its resources are online.

Use the `clresourcegroup(1CL)` utility for this purpose. By default, the `clsetup` utility assigns the name `vucmm_framework_rg` to the multiple-owner volume-manager framework resource group.

```
# clresourcegroup status vucmm_framework_rg
```

18 If the multiple-owner volume-manager framework resource group and its resources are *not* online, bring them online.

```
# clresourcegroup online vucmm_framework_rg
```

More Information

Resource Configuration

The following table lists the default resource configuration that the `clsetup` utility creates when you complete this task.

Resource Name, Resource Type, and Resource Group	Dependencies	Description
Resource type: `SUNW.vucmm_framework` Resource name: `vucmm_framework_rs` Resource group: `vucmm_framework_rg`	None.	Multiple-owner volume-manager framework resource.
Resource type: `SUNW.vucmm_svm` Resource name: `vucmm_svm_rs` Resource group: `vucmm_framework_rg`	Strong dependency on the multiple-owner volume-manager framework resource.	Solaris Volume Manager for Sun Cluster resource. Created only if Solaris Volume Manager for Sun Cluster was selected.

Next Steps The next step depends on the volume manager that you are using, as shown in the following table.

Volume Manager	Next Step
Solaris Volume Manager for Sun Cluster	"How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60
None	"Registering and Configuring Storage Resources for Oracle Files" on page 65

Creating a Global Device Group for the Oracle RAC Database

If you are using a volume manager for Oracle database files, the volume manager requires a global device group for the Oracle RAC database to use.

The type of global device group to create depends on the volume manager that you are using:

- If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set. See "How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60.

▼ How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database

Note – Perform this task only if you are using Solaris Volume Manager for Sun Cluster.

If you are using Solaris Volume Manager for Sun Cluster, Solaris Volume Manager requires a multi-owner disk set for the Oracle RAC database or Oracle ASM to use. For information about Solaris Volume Manager for Sun Cluster multi–owner disk sets, see "Multi-Owner Disk Set Concepts" in *Solaris Volume Manager Administration Guide*.

Before You Begin Note the following points.

- Ensure that the required Support for Oracle RAC software packages are installed on each node. For more information, see "Installing the Support for Oracle RAC Package" on page 35.
- Ensure that the multiple-owner volume-manager framework resource group is created and is online. For more information, see "Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group" on page 56.

Caution – Do not attempt to manually create the `/var/run/nodelist` file. Doing so can put the cluster at risk of a cluster-wide panic.

- Do not create any file systems in the multi-owner disk set.
- Disk devices that you add to the multi-owner disk set must be directly attached to all the cluster nodes.

1 Create a multi-owner disk set.

Use the `metaset(1M)` command for this purpose.

```
# metaset -s setname -M -a -h nodelist
```

`-s` *setname*
Specifies the name of the disk set that you are creating.

`-M`
Specifies that the disk set that you are creating is a multi-owner disk set.

`-a`
Specifies that the nodes that the `-h` option specifies are to be added to the disk set.

`-h` *nodelist*
Specifies a space-separated list of nodes that are to be added to the disk set. The Support for Oracle RAC software packages *must* be installed on each node in the list.

2 Add global devices to the disk set that you created in Step 1.

```
# metaset -s setname -a devicelist
```

`-s` *setname*
Specifies that you are modifying the disk set that you created in Step 1.

`-a`
Specifies that the devices that *devicelist* specifies are to be added to the disk set.

devicelist
Specifies a space-separated list of full device ID path names for the global devices that are to be added to the disk set. To enable consistent access to each device from any node in the cluster, ensure that each device ID path name is of the form `/dev/did/dsk/d`*N*, where *N* is the device number.

3 For the disk set that you created in Step 1, create the volumes that the Oracle RAC database will use.

Tip – If you are creating many volumes for Oracle data files, you can simplify this step by using soft partitions. For information about soft partitions, see Chapter 12, "Soft Partitions (Overview)," in *Solaris Volume Manager Administration Guide* and Chapter 13, "Soft Partitions (Tasks)," in *Solaris Volume Manager Administration Guide*.

Create each volume by concatenating slices on global devices that you added in Step 2. Use the `metainit(1M)` command for this purpose.

```
# metainit -s setname volume-abbrev numstripes width slicelist
```

`-s` *setname*
Specifies that you are creating a volume for the disk set that you created in Step 1.

volume-abbrev
Specifies the abbreviated name of the volume that you are creating. An abbreviated volume name has the format `d`*V*, where *V* is the volume number.

numstripes
Specifies the number of stripes in the volume.

width
Specifies the number of slices in each stripe. If you set *width* to greater than 1, the slices are striped.

slicelist
Specifies a space-separated list of slices that the volume contains. Each slice must reside on a global device that you added in Step 2.

4 If you are using mirrored devices, create the mirrors by using volumes that you created in Step 3 as submirrors.

If you are not using mirrored devices, omit this step.

Use the `metainit` command to create each mirror as follows:

```
# metainit -s setname mirror -m submirror-list
```

`-s` *setname*
: Specifies that you are creating a mirror for the disk set that you created in Step 1.

mirror
: Specifies the name of the mirror that you are creating in the form of an abbreviated volume name. An abbreviated volume name has the format `d`*V*, where *V* is the volume number.

submirror-list
: Specifies a space-separated list of submirrors that the mirror is to contain. Each submirror must be a volume that you created in Step 3. Specify the name of each submirror in the form of an abbreviated volume name.

Note – For information on configuring a Solaris Volume Manager disk set in a zone cluster, see "How to Add a Disk Set to a Zone Cluster (Solaris Volume Manager)" in *Oracle Solaris Cluster Software Installation Guide*.

---

5 Verify that each node is correctly added to the multi-owner disk set.

Use the `metaset` command for this purpose.

```
# metaset -s setname
```

`-s` *setname*
: Specifies that you are verifying the disk set that you created in Step 1.

This command displays a table that contains the following information for each node that is correctly added to the disk set:

- The `Host` column contains the node name.
- The `Owner` column contains the text `multi-owner`.
- The `Member` column contains the text `Yes`.

6 Verify that the multi-owner disk set is correctly configured.

```
# cldevicegroup show setname
```

setname
: Specifies that configuration information only for the disk set that you created in Step 1 is displayed.

This command displays the device group information for the disk set. For a multi-owner disk set, the device group type is `Multi-owner_SVM`.

7 Verify the online status of the multi-owner disk set.

```
# cldevicegroup status setname
```

This command displays the status of the multi-owner disk set on each node in the multi-owner disk set.

8 On each node that can own the disk set, change the ownership of each volume that you created in Step 3.

Note – For a zone cluster, perform this step in the zone cluster.

Change the volume ownership as follows:

- Owner: the DBA user
- Group: the DBA group

The DBA user and the DBA group are created as explained in "How to Create the DBA Group and the DBA User Accounts" on page 28.

Ensure that you change ownership only of volumes that the Oracle RAC database will use.

```
# chown user-name:group-name volume-list
```

user-name
Specifies the user name of the DBA user. This user is normally named `oracle`.

group-name
Specifies the name of the DBA group. This group is normally named `dba`.

volume-list
Specifies a space-separated list of the logical names of the volumes that you created for the disk set. The format of these names depends on the type of device where the volume resides, as follows:

- For block devices: `/dev/md/`*setname*`/dsk/d`*V*
- For raw devices: `/dev/md/`*setname*`/rdsk/d`*V*

The replaceable items in these names are as follows:

setname
Specifies the name of the multi-owner disk set that you created in Step 1.

V
Specifies the volume number of a volume that you created in Step 3.

Ensure that this list specifies each volume that you created in Step 3.

9 Grant read access and write access to the owner of each volume whose ownership you changed in Step 8.

Note – For a zone cluster, perform this step in the zone cluster.

Grant access to the volume on each node that can own the disk set. Ensure that you change access permissions only of volumes that the Oracle RAC database will use.

```
# chmod u+rw volume-list
```

volume-list
Specifies a space-separated list of the logical names of the volumes to whose owners you are granting read access and write access. Ensure that this list contains the volumes that you specified in Step 8.

10 If you are using Oracle ASM, specify the raw devices that you are using for the Oracle ASM disk group.

To specify the devices, modify the `ASM_DISKSTRING` Oracle ASM instance-initialization parameter.

For example, to use the `/dev/md/`*setname*`/rdsk/d` path for the Oracle ASM disk group, add the value `/dev/md/*/rdsk/d*` to the `ASM_DISKSTRING` parameter. If you are modifying this parameter by editing the Oracle initialization parameter file, edit the parameter as follows:

```
ASM_DISKSTRING = '/dev/md/*/rdsk/d*'
```

If you are using mirrored devices, specify external redundancy in the Oracle ASM configuration.

For more information, see your Oracle documentation.

Next Steps Go to "Registering and Configuring Storage Resources for Oracle Files" on page 65.

Registering and Configuring Storage Resources for Oracle Files

Storage resources provide fault monitoring and automatic fault recovery for global device groups and file systems.

If you are using global device groups or shared file systems for Oracle files, configure storage resources to manage the availability of the storage on which the Oracle software depends.

Configure the following types of storage resources:

- Global device groups:
  - Solaris Volume Manager for Sun Cluster multi-owner disk sets
- Shared file systems:
  - A file system on a qualified NAS device with Oracle RAC:
    - Oracle's Sun ZFS Storage Appliance NAS device

Note – You can configure a NAS NFS in a zone cluster with the existing tools. See "Tasks for Configuring Qualified NAS Devices for Oracle Files" on page 40.

This section contains the following information about registering and configuring storage resources for Oracle files:

- "Tools for Registering and Configuring Storage Resources for Oracle Files" on page 66
- "How to Register and Configure Storage Resources for Oracle Files by Using `clsetup`" on page 66

Tools for Registering and Configuring Storage Resources for Oracle Files

Oracle Solaris Cluster provides the following tools for registering and configuring storage resources for Oracle files in a global cluster or a zone cluster:

- **The `clsetup`(1CL) utility.** For more information, see "How to Register and Configure Storage Resources for Oracle Files by Using `clsetup`" on page 66.
- **Oracle Solaris Cluster maintenance commands.** For more information, see "Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands" on page 179.

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

▼ How to Register and Configure Storage Resources for Oracle Files by Using clsetup

Perform this procedure from only one node of the cluster.

Before You Begin Ensure that the following prerequisites are met:

- The Oracle RAC framework resource group is created and is online. For more information, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.
- The multiple-owner volume-manager framework resource group is created and is online. For more information, see "Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group" on page 56.

Caution – Do not attempt to manually create the `/var/run/nodelist` file. Doing so can put the cluster at risk of a cluster-wide panic.

- Required volumes, global device groups, and file systems are created. For more information, see the following sections:
  - "Installing Storage Management Software With Support for Oracle RAC" on page 42
  - "Creating a Global Device Group for the Oracle RAC Database" on page 60
- Required file systems are mounted.

Ensure that you have the following information:

- The name of each scalable device group that you are using for Oracle files, if any
- The mount point of each shared file system that you are using for Oracle files, if any

1 On one node of the cluster, become superuser.

2 Start the `clsetup` utility.

```
# clsetup
```

The `clsetup` main menu is displayed.

3 Choose the menu item, Data Services.

The Data Services menu is displayed.

4 Choose the menu item, Oracle Real Application Clusters.

The `clsetup` utility displays information about Support for Oracle RAC.

5 Press Return to continue.

The `clsetup` utility prompts you to select whether you are performing the initial configuration of Support for Oracle RAC or administering an existing configuration.

6 Choose the menu item, Oracle RAC Create Configuration.

The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7 Type the option number for the location of the Oracle RAC cluster and press Return.

- **If you select the global cluster option, the `clsetup` utility displays the list of components to configure. Skip to Step 9.**
- **If you select the zone cluster option, the `clsetup` utility prompts you to select the required zone cluster. Proceed to Step 8.**

8 **Type the option number for the required zone cluster and press Return.**

The `clsetup` utility displays the list of components of Oracle RAC to configure.

9 **Choose the menu item, Storage Resources for Oracle Files.**

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

10 **Verify that the prerequisites are met, and press Return.**

If you are prompted for resources for scalable device groups, omit this step.

11 **If you are prompted to select a Storage Management Scheme for your Oracle files, select the appropriate scheme.**

- NAS Device

12 **If no suitable resources exist, or if no resource exists for a device group that you are using, add a resource to the list.**

The `clsetup` utility displays a list of the resources for scalable device groups that are configured on the cluster. If no suitable resources exist, this list is empty.

If resources exist for all the device groups that you are using, omit this step.

For each resource that you are adding, perform the following steps:

a. **Press Return.**

The `clsetup` utility displays a list of the scalable device groups that are configured on the cluster.

b. **Type the option number for the device group to use and press Return.**

Once you select the device group, you can either select the entire disk group or choose to specify logical devices, or disks, in the disk group.

c. **Choose whether you want to specify logical devices.**

- **To specify logical devices, type `yes`. Proceed to Step d.**
- **To select the entire disk group, type `no`. Skip to Step e.**

d. **Type a comma-separated list of numbers that corresponds to the logical devices or disks you choose or type a for all.**

The `clsetup` utility returns you to the list of resources for scalable device groups that are configured on the cluster.

e. To confirm your selection of device groups, type d and press Return.

The `clsetup` utility returns you to the list of the resources for scalable device groups that are configured on the cluster. The resource that you are creating is added to the list.

13 Type the numbers that correspond to the resources that you require, if not already selected.

You can select existing resources, resources that are not yet created, or a combination of existing resources and new resources. If you select more than one existing resource, the selected resources must be in the same resource group.

14 To confirm your selection of resources for device groups, type d and press Return.

The `clsetup` utility displays a list of the resources for shared file-system mount points that are configured on the cluster. If no suitable resources exist, this list is empty.

15 If no suitable resources exist, or if no resource exists for a file-system mount point that you are using, add a resource to the list.

If resources exist for all the file-system mount points that you are using, omit this step.

For each resource that you are adding, perform the following steps:

a. Press Return.

The `clsetup` utility displays a list of the shared file systems that are configured on the cluster.

b. Type a comma-separated or space-separated list of numbers that correspond to the file systems that you are using for Oracle files and press Return.

c. To confirm your selection of file systems, type d and press Return.

The `clsetup` utility returns you to the list of the resources for file-system mount points that are configured on the cluster. The resource that you are creating is added to the list.

16 Type the option numbers for the resources that you require, if not already selected.

You can select existing resources, resources that are not yet created, or a combination of existing resources and new resources. If you select more than one existing resource, the selected resources must be in the same resource group.

17 To confirm your selection of resources for file-system mount points, type d and press Return.

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects that the utility will create or add to your configuration.

18 If you need to modify an Oracle Solaris Cluster object that the utility will create, modify the object.

a. Type the option number for the Oracle Solaris Cluster object that you are modifying and press Return.

The `clsetup` utility displays a list of properties that are set for the object.

b. Modify each property that you are changing as follows:

i. Type the option number for the property that you are changing and press Return.

The `clsetup` utility prompts you for the new value.

ii. At the prompt, type the new value and press Return.

The `clsetup` utility returns you to the list of properties that are set for the object.

c. When you have modified all the properties that you need to change, type `d` 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 or add to your configuration.

19 When you have modified all the Oracle Solaris Cluster objects that you need to change, type `d` and press Return.

The `clsetup` utility displays information about the Oracle RAC framework resource group for which storage resources will be configured.

20 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.

21 Press Return to continue.

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

22 (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. If you choose to quit `clsetup`, the utility recognizes your existing Oracle RAC framework resource group when you restart the utility.

23 Determine whether the resource groups that the wizard created are online.

```
# clresourcegroup status
```

24 If a resource group that the wizard created is *not* online, bring the resource group online.

For each resource group that you are bringing online, type the following command:

```
# clresourcegroup online -emM rac-storage-rg
```

rac-storage-rg
Specifies the name of the resource group that you are bringing online.

More Information

Resource Configuration

The following table lists the default resource configuration that the `clsetup` utility creates when you complete this task.

Resource Type, Resource Name, and Resource Group	Dependencies	Description
Resource type: `SUNW.ScalDeviceGroup` Resource name: `scal`*dg-name*`-rs`, where *dg-name* is the name of the device group that the resource represents Resource group: `scaldg-rg`	Strong dependency on the resource in the multiple-owner volume-manager framework resource group for the volume manager that is associated with the device group: either the Solaris Volume Manager for Sun Cluster resource.	Scalable device-group resource. One resource is created for each scalable device group that you are using for Oracle files.
Resource type: `SUNW.ScalMountPoint` Resource name: `scal-`*mp-dir*`-rs`, where *mp-dir* is the mount point of the file system, with `/` replaced by `–` Resource group: `scalmnt-rg`	Offline-restart dependency on the scalable device-group resource, if any. If you are using a file system on a qualified NAS device without a volume manager, this resource does not depend on any other resources.	Scalable file system mount point resource. One resource is created for each shared file system that you are using for Oracle files.

Note – For detailed information for the resource configuration for zone clusters, see the figures in Appendix A, "Sample Configurations of This Data Service."

Next Steps

If you are using Oracle ASM, go to "Registering and Configuring the Oracle ASM Resource Group" on page 72.

Otherwise, go to Chapter 4, "Enabling Oracle RAC to Run in a Cluster."

Registering and Configuring the Oracle ASM Resource Group

Oracle Automatic Storage Management (Oracle ASM) manages the storage that used by the Oracle database. This wizard creates an Oracle ASM instance resource for the Oracle database.

This section contains the following information about registering the Oracle ASM resource group:

Tools for Registering and Configuring the Oracle ASM Resource Group

Oracle Solaris Cluster software provides the following tools for registering and configuring the Oracle ASM resource group in the global cluster or in a zone cluster:

- **The `clsetup` utility.** For more information, see "How to Register and Configure the Oracle ASM Resource Group by Using `clsetup`" on page 72.
- **Oracle Solaris Cluster maintenance commands.** For more information, see Appendix D, "Command-Line Alternatives."

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

Note – The `clsetup` utility runs only in a voting node of the global cluster.

▼ How to Register and Configure the Oracle ASM Resource Group by Using clsetup

When you register and configure the Oracle ASM resource group for a cluster, the Oracle ASM resource group is created.

Perform this procedure from one node only.

Before You Begin

- Ensure that Oracle ASM disk groups are configured. For more information, see "Using Oracle ASM" on page 46.

Ensure that you have the following information:

- The name of the Oracle Grid Infrastructure home directory.
- The list of Oracle ASM systems identifiers (SID).
- The names of the Oracle ASM disk groups to use.

1 **Become superuser on any cluster node.**

2 **Start the `clsetup` utility.**

```
# clsetup
```

The `clsetup` main menu is displayed.

3 **Type the option number for Data Services and press Return.**

The Data Services menu is displayed.

4 **Type the option number for Oracle Real Application Clusters and press Return.**

The `clsetup` utility displays information about Support for Oracle RAC.

5 **Press Return to continue.**

The `clsetup` utility prompts you to select whether to create the configuration of Support for Oracle RAC or administer an existing configuration.

Note – The `clsetup` utility currently allows ongoing administration of an Oracle RAC framework running only in the global cluster. For ongoing administration of an Oracle RAC framework configured in a zone cluster, use instead the Oracle Solaris Cluster maintenance commands.

---

6 **Type the option number for Oracle RAC Create Configuration and press Return.**

The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7 **Type the option number for the location of the Oracle RAC cluster and press Return.**

- **If you select the global cluster option, the `clsetup` utility displays the list of components of Oracle RAC to configure. Skip to Step 9.**
- **If you select the zone cluster option, the `clsetup` utility prompts you to select the required zone cluster. Proceed to Step 8.**

8 **Type the option number for the required zone cluster and press Return.**

The `clsetup` utility displays a list of components of Oracle RAC to configure.

9. **Type the option number for Automatic Storage Management (ASM) and press Return.**

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

10. **Select the nodes where you require Oracle ASM to run.**

    - **To accept the default selection of all listed nodes in an arbitrary order, type `a` and press Return.**
    - **To select a subset of the listed nodes, type a comma-separated or space-separated list of the option numbers for the nodes you choose and press Return.**

      Ensure that the nodes are listed in the order in which the nodes are to appear in the Oracle RAC framework resource group's node list.
    - **To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the option numbers for the nodes you choose and press Return.**

      Ensure that the nodes are listed in the order in which the nodes are to appear in the Oracle RAC framework resource group's node list.

11. **If needed to end the selection of nodes, type `d` and press Return.**

    The `clsetup` utility displays a list of Oracle ASM instance resources.

12. **Type the option number that corresponds to the Oracle ASM instance resource to use.**

    If no Oracle ASM instance resource is available and you are prompted to create a resource, press Return. Skip to Step 14.

13. **To confirm your selection of an Oracle ASM instance resource, type `d` and press Return.**

    The `clsetup` utility displays the selection screen for the Oracle Grid Infrastructure home directory.

14. **Type the option number to choose a listed directory or to specify a directory explicitly and press Return.**

    The `clsetup` utility displays a list of Oracle ASM system identifiers (SID) that are discovered in the cluster.

15. **Review the list of SIDs.**

    - **If the list is correct, type `d` and press Return.**
    - **If the list is not correct, type the option number for the SID to change and press Return.**

    The `clsetup` utility displays information about Oracle ASM disk-group resources.

16 Respond to the prompt that asks whether to create a disk-group resource.

- **To create a new disk-group resource, type `y` and press Return.**

  The `clsetup` utility displays a list of existing Oracle ASM disk groups. Proceed to Step 17.

- **If you do not want to create a new disk-group resource, type `n` and press Return.**

  The `clsetup` utility displays a list of discovered storage resources. Skip to Step 19.

17 Specify the Oracle ASM disk groups.

Type the option number for each disk group to use and press Return. When all disk groups are selected, type `d` and press Return.

The selected Oracle ASM disk groups are added to the Oracle ASM disk-group resources selection panel.

18 Review the list of Oracle ASM disk-group resources.

- **If the list is correct, type `d` and press Return.**

- **If the list is not correct, type the option number or letter to edit the list of resources and press Return.**

  When the list of disk-group resources is correct, type `d` and press Return.

The `clsetup` utility displays a list of discovered storage resources.

19 Review the list of storage resources to manage the file-system mount point where the Oracle Grid Infrastructure home is installed.

- **If the list is correct, type `d` and press Return.**

- **If no storage resource is listed, type `d` and press Return.**

  The `clsetup` utility will create a new resource when you complete Oracle ASM configuration.

- **If the list is not correct, type the option number for the correct storage resource and press Return.**

The `clsetup` utility displays a list of Oracle ASM disk-group resources that manage Oracle ASM disk groups.

20 If no suitable disk-group resources exist, or if no resource exists for an Oracle ASM disk group that you are using, add a resource to the list.

a. Type `y` and press Return.

The `clsetup` utility discovers Oracle ASM disk groups.

b. Type a comma-separated or space-separated list of the option numbers for the Oracle ASM disk groups to use and press Return.

c. To confirm your selection of disk groups, type `d` and press Return.

The `clsetup` utility returns you to the list of Oracle ASM disk-group resources. The resource that you are creating is added to the list.

21 Type the option numbers for the resources that you require, if not already selected.

You can select existing resources, resources that are not yet created, or a combination of existing resources and new resources. If you select more than one existing resource, the selected resources must be in the same resource group.

22 To confirm your selection of resources for Oracle ASM disk groups, type `d` and press Return.

The `clsetup` utility displays a list of discovered underlying disk sets or disk groups.

23 Type a comma-separated list of the option numbers for the Oracle ASM disk groups to use and press Return.

24 To confirm your selection of Oracle ASM disk groups, type `d` and press Return.

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects for Oracle ASM that the utility will create or add to your configuration.

25 If you need to modify an Oracle Solaris Cluster object, modify the object.

a. Type the option number for the object to modify and press Return.

The `clsetup` utility prompts you for the new value.

b. At the prompt, type the new value and press Return.

The `clsetup` utility returns you to the list of properties that are set for the object.

26 When you have modified all the Oracle Solaris Cluster objects that you need to change, type `d` and press Return.

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

27 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.

28 Press Return to continue.

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

29 (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.

More Information

Resource Configuration

The following table lists the default resource configuration that the `clsetup` utility creates when you complete this task.

Resource Name, Resource Type, and Resource Group	Dependencies	Description
Resource type: `SUNW.scalable_rac_server_proxy` Resource name: `rac_server_proxy-rs` Resource group: `rac_server_proxy-rg`	Strong dependency on the Oracle RAC framework resource. Offline-restart dependency on the clustered Oracle ASM disk-group resource. Offline-restart dependency on the Oracle Grid Infrastructure framework resource. When configured with Solaris Volume Manager for Sun Cluster, strong positive affinity by the Oracle RAC instance proxy resource group for the Oracle RAC framework resource group. When configured with hardware RAID, strong positive affinity by the Oracle RAC instance proxy resource group for the clustered Oracle ASM disk-group resource group.	Oracle RAC instance proxy resource

Resource Name, Resource Type, and Resource Group	Dependencies	Description
Resource type: `SUNW.oracle_asm_diskgroup` Resource name: `asm-dg-rs` Resource group: `asm-dg-rg`	Strong positive affinity by the clustered Oracle ASM disk-group resource group for the clustered Oracle ASM instance resource group. When configured with Solaris Volume Manager for Sun Cluster: ■ Strong dependency on the clustered Oracle ASM instance resource. ■ Offline-restart dependency on the scalable device-group resource for database files. When configured with hardware RAID, offline-restart dependency on the clustered Oracle ASM disk-group resource.	Clustered Oracle ASM disk-group resource
SPARC: Resource type: `SUNW.scalable_oracle_asm_instance_proxy` Resource name: `asm-inst-rs` Resource group: `asm-inst-rg`	Offline-restart dependency on the Oracle Grid Infrastructure resource. When configured with hardware RAID, strong positive affinity by the clustered Oracle ASM instance resource group for the Oracle Clusterware framework resource group.	Clustered Oracle ASM instance resource
Resource type: `SUNW.ScalMountPoint` Resource name: `asm-mp-rs` Resource group: `asm-mp-rg`	Offline-restart dependency on the scalable device-group resource, if any. If you are using a file system on a qualified NAS device without a volume manager, this resource does not depend on any other resources.	Resource for the Oracle Grid Infrastructure home on a scalable file system mount point. One resource is created for each shared file system that you are using for Oracle files.
Resource type: `SUNW.ScalDeviceGroup` Resource name: `scal`*dg-name*`-rs`, where *dg-name* is the name of the device group that the resource represents Resource group: `scaldg-rg`	Strong dependency on the resource in the multiple-owner volume-manager framework resource group for the volume manager that is associated with the device group: either the Solaris Volume Manager for Sun Cluster resource.	Scalable device-group resource. One resource is created for each scalable device group that you are using for Oracle files.

Next Steps Go to Chapter 4, "Enabling Oracle RAC to Run in a Cluster."

◆ ◆ ◆ CHAPTER 4

Enabling Oracle RAC to Run in a Cluster

This chapter explains how to enable Oracle RAC to run on your Oracle Solaris Cluster nodes.

Overview of Tasks for Enabling Oracle RAC to Run in a Cluster

Table 4–1 summarizes the tasks for enabling Oracle RAC to run in a cluster.

Perform these tasks in the order in which they are listed in the table.

TABLE 4–1 Tasks for Enabling Oracle RAC to Run in a Cluster

Task	Instructions
Install the Oracle RAC software	"Installing Oracle RAC Software" on page 80
Verify the installation of the Oracle RAC software	"Verifying the Installation of Oracle RAC" on page 81
Create an Oracle ASM instance	"Creating an Oracle ASM Instance and Disk Groups" on page 81
Create an Oracle Clusterware framework resource.	"Creating an Oracle Grid Infrastructure Framework Resource" on page 82
Create your Oracle database	"Creating an Oracle Database" on page 85

TABLE 4–1 Tasks for Enabling Oracle RAC to Run in a Cluster *(Continued)*

Task	Instructions
Verify the Support for Oracle RAC installation and configuration	"Verifying the Installation and Configuration of Support for Oracle RAC" on page 93

Installing Oracle RAC Software

This section contains the following information:

For detailed instructions for installing Oracle RAC, see your Oracle documentation.

Installing Binary Files and Configuration Files on a Shared File System

To simplify the maintenance of your Oracle installation, you can install the Oracle binary files and Oracle configuration files on a shared file system. The following shared file systems are supported:

- The PxFS-based cluster file system
- A file system on a qualified NAS device

If you are installing the Oracle binary files and Oracle configuration files on a shared file system, specify the absolute paths to the file system when the Oracle installation tool requests this information. Do *not* use a symbolic link whose target is the shared file system.

To install Oracle binary files and configuration files on a local file system, follow the usual procedures as described in your Oracle Database documentation.

Overriding Networking Defaults for Oracle Grid Infrastructure

By default, the network interfaces panel of the Universal Installer shows all interfaces as private. If you are installing Oracle RAC for use with Oracle Solaris Cluster software, override these defaults as follows:

- Ensure that `clprivnet0` is the only private interface.
- Set the public network interfaces to public.

- Ensure that all other interfaces are not used. These interfaces represent the underlying network interfaces for the cluster interconnects.

Next Steps

Go to "Verifying the Installation of Oracle RAC" on page 81.

Verifying the Installation of Oracle RAC

After you have installed Oracle RAC, verify that the installation is correct. Perform this verification before you attempt to create your Oracle database. This verification does *not* verify that Oracle RAC database instances can be started and stopped automatically.

This section contains the following procedures:

- "How to Verify the Installation of Oracle RAC" on page 81

▼ How to Verify the Installation of Oracle RAC

- **Confirm that the tests that the Oracle installer runs to verify the cluster were passed.**

  If the results of these tests are no longer available for review, run the Oracle utility `cluvfy` to repeat the tests.

  For more information, see your Oracle documentation.

Creating an Oracle ASM Instance and Disk Groups

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.

▼ How to Create an Oracle ASM Instance and Disk Groups

Before You Begin

- Ensure that the Oracle Grid Infrastructure software is already installed.
- Ensure that the Oracle RAC framework is running on all nodes where the Oracle ASM instance is to be created.

- Ensure that the file system where the Oracle ASM `$ORACLE_HOME` directory will reside is already created.

1 **Become superuser on a cluster node.**

2 **Install and configure an Oracle ASM instance and create disk groups.**

Use Oracle ASM Configuration Assistant (ASMCA). See the appropriate Oracle ASM documentation for instructions.

Creating an Oracle Grid Infrastructure Framework Resource

The Oracle Grid Infrastructure framework resource is a resource of the type `SUNW.crs_framework`. This resource is created within the Oracle RAC framework resource group whenever Oracle Grid Infrastructure is present on the cluster. This resource enables Oracle Solaris Cluster software to control the starting and stopping of Oracle Grid Infrastructure through the resource dependencies in Oracle Solaris Cluster. This ensures that Oracle Grid Infrastructure is only started when the resources it depends on are also available, and that it is stopped cleanly when any resource it depends on is no longer available.

Note – When the Oracle Grid Infrastructure framework resource is created in an Oracle Solaris Cluster configuration, it disables the automatic startup of Oracle Clusterware itself. The removal of the Oracle Grid Infrastructure framework resource does not by itself re-enable automatic startup of Oracle Grid Infrastructure. To re-enable automatic startup of Oracle Grid Infrastructure, see the Oracle Grid Infrastructure documentation for your release of Oracle Grid Infrastructure software.

- "How to Create an Oracle Grid Infrastructure Framework Resource" on page 82

▼ How to Create an Oracle Grid Infrastructure Framework Resource

Before You Begin Ensure that the following prerequisites are met:

- The Oracle RAC framework is created and is online. See "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.
- Storage resources for Oracle files are configured. See "Registering and Configuring Storage Resources for Oracle Files" on page 65.
- The Oracle RAC software is installed. See "Installing Oracle RAC Software" on page 80.
- The installation of the Oracle RAC software is verified. See "Verifying the Installation of Oracle RAC" on page 81.

Ensure that you have the following information:

- The full path to the Oracle Grid Infrastructure home directory

1 **On one node of the cluster, become superuser.**

2 **Start the `clsetup` utility.**

```
# clsetup
```

The `clsetup` main menu is displayed.

3 **Choose the menu item, Data Services.**

The Data Services menu is displayed.

4 **Choose the menu item, Oracle Real Application Clusters.**

The `clsetup` utility displays information about Support for Oracle RAC.

5 **Press Return to continue.**

The `clsetup` utility prompts you to select whether you are performing the initial configuration of Support for Oracle RAC or administering an existing configuration.

6 **Choose the menu item, Oracle RAC Create Configuration.**

The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7 **Type the option number for the location of the Oracle RAC cluster and press Return.**

- **If you select the global cluster option, the `clsetup` utility displays the list of components of Oracle RAC to configure. Skip to Step 9.**
- **If you select the zone cluster option, the `clsetup` utility prompts you to select the required zone cluster. Proceed to Step 8.**

8 **Type the option number for the required zone cluster and press Return.**

The `clsetup` utility displays the list of components of Oracle RAC to configure.

9 **Choose the menu item, Oracle Clusterware Framework Resource.**

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

10 **Verify that the prerequisites are met, and press Return.**

The `clsetup` utility displays a list of Oracle Grid Infrastructure home directories that exist on the cluster.

11 Specify the Oracle Grid Infrastructure home directory for your installation of the Oracle Grid Infrastructure software.

- **If the directory is listed, type the option number for the directory that you are selecting and press Return.**
- **If the directory is not listed, type the full path to the Oracle Grid Infrastructure home directory and press Return.**

12 Specify the Oracle Clusterware OCR and voting disks.

- **If a desired name is listed, type the number that corresponds to the name that you are selecting and press Return.**

  The `clsetup` utility displays a list of Oracle home directories that exist on the cluster.

- **If a desired name is not listed, specify the name explicitly:**

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

     The `clsetup` utility prompts you for the full path for the OCR and voting disk.

  b. **Type the full path of the file-system mount point or disk group and press Return.**

     The `clsetup` utility prompts you whether to enter more values. Type `yes` to specify another path name or `no` if you have no more path names to add and press Return.

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

13 If you require a different name for any Oracle Solaris Cluster objects, change the name.

a. **Type the option number for the name to change 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 names of the Oracle Solaris Cluster objects that the utility will create.

14 To confirm your selection of Oracle Solaris Cluster object names, press Return.

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

15 **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.

16 **Press Return to continue.**

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

17 **(Optional) Type `q` and press Return repeatedly until you quit the `clsetup` utility.**

Creating an Oracle Database

Perform this task to configure and create an Oracle database in an Oracle Solaris Cluster environment.

To create the database, use one of the commands from the following list:

- The Oracle `dbca` command
- The Oracle `sqlplus` command

Note – To use raw devices such as for hardware RAID or Solaris Volume Manager for Sun Cluster, manually create the database by using a `CREATE DATABASE` statement. For more information, see "Creating a Database with the CREATE DATABASE Statement" (`http://docs.oracle.com/cd/E11882_01/server.112/e25494/create003.htm#CIAEJDBE`) in *Oracle Database Administrator's Guide* and *Oracle Database SQL Language Reference* (`http://docs.oracle.com/cd/E11882_01/server.112/e26088/toc.htm`).

You must create the Oracle database to be administrator managed, not policy managed. This ensures that Oracle Solaris Cluster software controls on which servers the database is started.

For detailed instructions for creating an Oracle database, see your Oracle documentation.

The way to specify to the `dbca` command the location of data files on a shared file system depends on your release of Oracle.

- "How to Specify the Location of Data Files on a Shared File System" on page 86

▼ How to Specify the Location of Data Files on a Shared File System

1. **When `dbca` prompts you for a storage option, select Cluster File System.**
2. **When `dbca` prompts you for the location, select the option for common location.**
3. **After selecting the option for a common location, type the name of the directory on the shared file system and press Return.**

Next Steps After you have created the Oracle database, go to "Configuring Resources for Oracle RAC Database Instances" on page 86.

Configuring Resources for Oracle RAC Database Instances

Resources for Oracle RAC database instances enable the database instances to be administered from Oracle Solaris Cluster. These resources also enable Oracle Solaris Cluster and Oracle Grid Infrastructure to interoperate.

This section contains the following information about configuring resources for Oracle RAC database instances:

Tools for Registering and Configuring Resources for Oracle RAC Database Instances

Oracle Solaris Cluster provides the following tools for registering and configuring resources for Oracle RAC database instances in the global cluster or in a zone cluster:

- **The clsetup utility.** For more information, see the following sections:
  - "How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate" on page 87
- **Oracle Solaris Cluster maintenance commands.** For more information, see the following sections:
  - "Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands" on page 183

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

How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate

When you perform this task, the `clsetup` utility creates the following resources for interoperation with Oracle Grid Infrastructure:

- An Oracle Solaris Cluster resource to act as a proxy for the Oracle RAC database
- An Oracle Solaris Cluster resource to represent the Oracle Grid Infrastructure framework
- One of the following storage choices:
  - Oracle Grid Infrastructure resources to represent scalable device groups and scalable file-system mount points
  - Oracle Grid Infrastructure resources to represent Oracle ASM disk groups

Before You Begin Ensure that the following prerequisites are met:

- The Oracle RAC framework resource group is created and is online. See "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.
- Storage resources for Oracle files are configured. See "Registering and Configuring Storage Resources for Oracle Files" on page 65.
- The Oracle RAC software is installed. See "Installing Oracle RAC Software" on page 80.
- The installation of the Oracle RAC software is verified. See "Verifying the Installation of Oracle RAC" on page 81.
- The Oracle database is created. See "Creating an Oracle Database" on page 85.
- The Oracle database is running.

To determine whether the Oracle database is running, use the Oracle `srvctl` command as follows:

```
$ srvctl status database -d db-name
```

db-name specifies the name of the Oracle database.

Ensure that you have the following information:

- The full path to the Oracle Grid Infrastructure home directory
- The full path to the Oracle home directory
- The Oracle system identifier (SID) for each instance of the database for which you are configuring resources

- The names of all Oracle Solaris Cluster storage resources for Oracle files that you are using

1. **On one node of the cluster, become superuser.**

2. **Start the `clsetup` utility.**

   ```
   # clsetup
   ```

   The `clsetup` main menu is displayed.

3. **Choose the menu item, Data Services.**

   The Data Services menu is displayed.

4. **Choose the menu item, Oracle Real Application Clusters.**

   The `clsetup` utility displays information about Support for Oracle RAC.

5. **Press Return to continue.**

   The `clsetup` utility prompts you to select whether you are performing the initial configuration of Support for Oracle RAC or administering an existing configuration.

6. **Choose the menu item, Oracle RAC Create Configuration.**

   The `clsetup` utility prompts you to select the Oracle RAC cluster location. This location can be the global cluster or a zone cluster.

7. **Type the option number for the location of the Oracle RAC cluster and press Return.**

   - **If you select the global cluster option, the `clsetup` utility displays the list of components of Oracle RAC to configure. Skip to Step 9.**

   - **If you select the zone cluster option, the `clsetup` utility prompts you to select the required zone cluster. Proceed to Step 8.**

8. **Type the option number for the required zone cluster and press Return.**

   The `clsetup` utility displays the list of components of Oracle RAC to configure.

9. **Choose the menu item, Resources for Oracle Real Application Clusters Database Instances.**

   The `clsetup` utility displays a list of names of Oracle RAC databases that are configured on the cluster.

10. **Specify the name of the Oracle RAC database that you are configuring.**

    - **If the name is listed, type the number that corresponds to the name that you are selecting and press Return.**

      The `clsetup` utility displays a list of Oracle home directories that exist on the cluster.

- **If the name is not listed, specify the name explicitly:**

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

     The `clsetup` utility prompts you for the name of the Oracle RACdatabase that you are configuring.

  b. **Type the name of the Oracle RAC database that you are configuring and press Return.**

     The `clsetup` utility displays a list of Oracle home directories that exist on the cluster.

11 Specify the Oracle home directory for your installation of the Oracle RAC software.

- **If the directory is listed, type the number that corresponds to the directory that you are selecting 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 the first node in the list of nodes where Oracle RAC is to run.

- **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 the first node in the list of nodes where Oracle RAC is to run.

12 Specify the Oracle SID of the Oracle RAC database instance for the node that you are configuring.

Perform this step for each node in the list of nodes where Oracle RAC is to run.

Note – You must specify a unique Oracle SID for each node. If you specify a duplicate SID, the `clsetup` utility warns you that the SID is a duplicate and prompts you to specify the SID again.

- **If the SID is listed, type the number that corresponds to the SID that you are selecting and press Return.**

The response of the `clsetup` utility depends on whether you have specified the Oracle SID for all nodes where Oracle RAC is to run.

- If you have *not* specified the Oracle SID for all nodes, the `clsetup` utility prompts you for the SID of the next node in the node list.
- If you have specified the Oracle SID for all nodes, the `clsetup` utility displays a list of Oracle Solaris Cluster storage resources for Oracle files that are configured. For information about these resources, see "Registering and Configuring Storage Resources for Oracle Files" on page 65.

- **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 response of the `clsetup` utility depends on whether you have specified the Oracle SID for all nodes where Oracle RAC is to run.

  - If you have *not* specified the Oracle SID for all nodes, the `clsetup` utility prompts you for the SID of the next node in the node list.
  - If you have specified the Oracle SID for all nodes, the `clsetup` utility displays a list of Oracle Solaris Cluster storage resources for Oracle files that are configured. For information about these resources, see "Registering and Configuring Storage Resources for Oracle Files" on page 65.

13 Type a comma-separated or space-separated list of the numbers that correspond to the storage resources that you are using, and press Return.

- **If you choose Oracle Automatic Storage Management (Oracle ASM), the `clsetup` utility displays the names of the Oracle ASM disk group resources.**
- **If you choose Other, the `clsetup` utility displays available scalable device group and scalable mount point resources.**

14 Type the option number for the storage resources to use.

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

15 If you require a different name for any Oracle Solaris Cluster objects, change the name.

a. **Type the option number for the name to change 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 names of the Oracle Solaris Cluster objects that the utility will create.

16 To confirm your selection of Oracle Solaris Cluster object names, press Return.

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

17 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.

18 Press Return to continue.

The `clsetup` utility returns you to the list of options for configuring Support for Oracle RAC.

19 (Optional) Type `q` and press Return repeatedly until you quit the `clsetup` utility.

More Information

Resource Configuration

The following table lists the default configuration of Oracle Solaris Cluster resources that the `clsetup` utility creates when you complete this task.

Resource Type, Resource Name, and Resource Group	Dependencies	Description
Resource type: `SUNW.crs_framework` Resource name: `crs_framework-rs` Resource group: `rac-framework-rg`	Strong dependency on the Oracle RAC framework resource. Offline-restart dependencies on all scalable device-group resources for Oracle files, if any. If you are using a shared file system without a volume manager, this resource has an offline-restart dependency on all scalable file-system mount-point resources for Oracle files. If you are not using a shared file system or a volume manager, this resource has no offline-restart dependencies on any other resources.	Oracle Grid Infrastructure framework resource.

Resource Type, Resource Name, and Resource Group	Dependencies	Description
Resource type: `SUNW.scalable_rac_server_proxy` Resource name: `rac_server_proxy-rs` Resource group: `rac_server_proxy-rg`	Strong dependency on the Oracle RAC framework resource. Offline-restart dependency on the Oracle Grid Infrastructure framework resource.	Proxy resource for the Oracle RAC database server.

Note – For detailed information for the resource configuration for zone clusters, see the figures in Appendix A, "Sample Configurations of This Data Service."

The `clsetup` utility also creates an Oracle Grid Infrastructure resource for each Oracle Solaris Cluster resource for scalable device groups and scalable file-system mount points on which Oracle components depend.

The name of each Oracle Grid Infrastructure resource that the `clsetup` utility creates is as follows:

`sun.`*node*`.`*sc-rs*

The replaceable items in this name are as follows:

- *node* specifies the name of the node where the Oracle Grid Infrastructure resource is to run.
- *sc-rs* specifies the name of the Oracle Solaris Cluster resource that the Oracle Grid Infrastructure resource represents. This resource is created when the task in "Registering and Configuring Storage Resources for Oracle Files" on page 65 is performed.

The Oracle Grid Infrastructure resource for the database instance on *node* depends on the Oracle Grid Infrastructure resource that the `clsetup` utility creates for that node. Oracle Grid Infrastructure resources for database instances are created during the installation and configuration of the Oracle RAC

The name of the Oracle Grid Infrastructure resource for the database instance on a node is as follows:

`ora.`*dbname*`.db`

The *dbname* portion of the name specifies the database name of the database instance that the Oracle Grid Infrastructure resource represents.

Next Steps Go to "Verifying the Installation and Configuration of Support for Oracle RAC" on page 93.

Verifying the Installation and Configuration of Support for Oracle RAC

After you install, register, and configure Support for Oracle RAC, verify the installation and configuration. Verifying the Support for Oracle RAC installation and configuration determines whether resources and resource groups for Oracle RAC behave as required.

Offline-restart dependencies between resources ensure that, if the independent resource is about to be taken offline, the dependent resource is taken offline first. The dependent resource remains offline until the independent resource is restarted. The procedures in this section explain how to verify that these dependencies are correctly set. For detailed information about offline-restart dependencies, see the description of the `resource_dependencies_offline_restart` resource property in the `r_properties(5)` man page.

Verifying the installation and configuration of Support for Oracle RAC involves the following tasks:

1. Verifying the configuration of the Oracle RAC framework resource group and, if used, verifying the configuration of the multiple-owner volume-manager framework resource group. See "How to Verify the Configuration of the Oracle RAC Framework Resource Group" on page 93 and "How to Verify the Configuration of the Multiple-Owner Volume-Manager Framework Resource Group" on page 94.
2. Verifying the configuration of storage resources for Oracle files. See "How to Verify the Configuration of Storage Resources for Oracle Files" on page 95.
3. Verifying the configuration of resources for Oracle RAC database instances. See "How to Verify the Configuration of Resources for Oracle RAC Database Instances" on page 96.
4. Verifying the correct behavior for shutdown and booting of the cluster. See "How to Verify the Correct Behavior for Shutdown and Booting of the Cluster" on page 98.

▼ How to Verify the Configuration of the Oracle RAC Framework Resource Group

The Oracle RAC framework resource group is created when the task in "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51 is performed.

1 Become superuser on a cluster node or assume a role that provides `solaris.cluster.read` and `solaris.cluster.admin` RBAC authorizations.

2 Verify that the Oracle RAC framework resource group is correctly configured.

```
# clresourcegroup show rac-fmk-rg
```

rac-fmk-rg
Specifies the name of the Oracle RAC framework resource group.

3 If the Oracle RAC framework resource group is not already online, bring online the Oracle RAC framework resource group.

```
# clresourcegroup online rac-fmwk-rg
```

rac-fmk-rg
Specifies the name of the Oracle RAC framework resource group.

4 Verify that the Oracle RAC framework resource group is online.

```
# clresourcegroup status
```

How to Verify the Configuration of the Multiple-Owner Volume-Manager Framework Resource Group

Perform this procedure if you use a multiple-owner volume-manager framework resource group to contain the volume manager resources that are used by your RAC configuration.

1 Become superuser on a cluster node or assume a role that provides `solaris.cluster.read` and `solaris.cluster.admin` RBAC authorizations.

2 Verify that the multiple-owner volume-manager framework resource group is correctly configured.

```
# clresourcegroup show vucmm-fmk-rg
```

vucmm-fmk-rg
Specifies the name of the multiple-owner volume-manager framework resource group.

3 If the multiple-owner volume-manager framework resource group is not already online, bring online the multiple-owner volume-manager framework resource group.

```
# clresourcegroup online vucmm-fmwk-rg
```

vucmm-fmk-rg
Specifies the name of the multiple-owner volume-manager framework resource group.

4 Verify that the RAC resource group is online.

```
# clresourcegroup status
```

How to Verify the Configuration of Storage Resources for Oracle Files

Perform this task only if your configuration of storage management schemes requires storage resources for Oracle files.

1 Become superuser on a cluster node.

2 Verify that all resource groups that contain storage resources for Oracle files are correctly configured.

```
# clresourcegroup show rac-storage-rg-list
```

rac-storage-rg-list
Specifies a comma-separated list of resource groups that contain storage resources for Oracle files.

3 Take offline the resource group that contains the resources on which the RAC database resource depends.

```
# clresourcegroup offline rac-storage-rg
```

rac-storage-rg
Specifies the name of the resource group that contains the resources on which the RAC database resource depends. The resource group to take offline depends on your choice of storage management schemes for Oracle files.

- If you are using a shared file system for Oracle files, take offline the resource group that contains scalable file-system mount-point resources.
- If you are using a volume manager without a file system for Oracle files, take offline the resource group that contains scalable device-group resources.

This step might require several minutes to complete.

4 Verify that the resource groups for Oracle RAC behave as follows:

- The status of the resource group that you took offline in Step 3 is offline.
- The status of the Oracle RAC database resource group is offline.
- The status of the Oracle RAC framework resource group is pending online blocked.

```
# clresourcegroup status
```

5 Verify that the Oracle Grid Infrastructure resources are offline.

To verify that the Oracle Grid Infrastructure resources are offline, consult the system messages file to find messages to indicate that Oracle Grid Infrastructure resources are being shut down.

6 **Bring online the resource group that you took offline in Step 3.**

```
# clresourcegroup online rac-storage-rg
```

rac-storage-rg
Specifies the name of the resource group that you took offline in Step 3.

This step might require several minutes to complete.

7 **On each node, restart the Oracle Grid Infrastructure.**

```
# Grid_home/bin/crsctl start crs
Startup will be queued to init within 30 seconds.
```

Grid_home
Specifies the Oracle Grid Infrastructure home directory. This directory contains the Oracle Grid Infrastructure binary files and Oracle Grid Infrastructure configuration files.

This step might require several minutes to complete.

8 **On each node, verify that the Oracle Grid Infrastructure resources are online.**

Use the Oracle command `crstat` for this purpose.

```
# Grid_home/bin/crsctl stat res -t
```

Note – The startup of Oracle Grid Infrastructure on all nodes might require several minutes to complete. If you perform this step before the Oracle Grid Infrastructure resources are online on all nodes, the status on some nodes might be offline. In this situation, repeat this step until the status of the Oracle Grid Infrastructure is online on all nodes.

---

9 **Verify that all resource groups for Oracle RAC are online.**

```
# clresourcegroup status
```

Note – Bringing online the resource group that contains the resources on which the RAC database resource depends also brings online the RAC database resource group. If you perform this step before the RAC database resource group is online on all nodes, the status on some nodes might be online faulted. In this situation, repeat this step until the status of the RAC database resource group is online on all nodes.

▼ How to Verify the Configuration of Resources for Oracle RAC Database Instances

Resources for Oracle RAC database instances are created when the tasks in one of the following sections are performed:

- "How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate" on page 87

- "Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands" on page 183

1 Become superuser on a cluster node.

2 Verify that the database resource group is correctly configured.

```
# clresourcegroup show rac-db-rg
```

rac-db-rg
Specifies the name of the database resource group.

3 If required, verify that the Oracle Grid Infrastructure resources to represent Oracle Solaris Cluster storage resources are correctly configured.

```
# Grid_home/bin/crsctl stat res -t
```

Grid_home
Specifies the Oracle Grid Infrastructure home directory. This directory contains the Oracle Grid Infrastructure binary files and Oracle Grid Infrastructure configuration files.

4 Disable the Oracle Grid Infrastructure framework resource.

```
# clresource disable -t SUNW.crs_framework +
```

This step might require several minutes to complete.

5 Verify that the Oracle Grid Infrastructure resources are offline.

a. Attempt to obtain the status of Oracle Grid Infrastructure resources.

Use the Oracle command `crstat` for this purpose.

```
# Grid_home/bin/crsctl stat res -t
CRS-0184: Cannot communicate with the CRS daemon.
```

Grid_home
Specifies the Oracle Grid Infrastructure home directory. This directory contains the Oracle Grid Infrastructure binary files and Oracle Grid Infrastructure configuration files.

b. Consult the system messages file to find messages to indicate that Oracle Grid Infrastructure resources are being shut down.

6 Verify that the following resources are offline on all nodes:

- The Oracle Grid Infrastructure framework resource
- The RAC database resource

```
# clresource status -t SUNW.crs_framework,SUNW.scalable_rac_server_proxy +
```

7 **On each node, restart the Oracle Grid Infrastructure.**

```
# Grid_home/bin/crsctl start crs
Startup will be queued to init within 30 seconds.
```

Grid_home
Specifies the Oracle Grid Infrastructure home directory. This directory contains the Oracle Grid Infrastructure binary files and Oracle Grid Infrastructure configuration files.

This step might require several minutes to complete.

8 **On each node, verify that the Oracle Grid Infrastructure resources are online.**

Use the Oracle command `crstat` for this purpose.

```
# Grid_home/bin/crsctl stat res -t
```

Note – The startup of Oracle Grid Infrastructure on all nodes might require several minutes to complete. If you perform this step before the Oracle Grid Infrastructure resources are online on all nodes, the status on some nodes might be offline. In this situation, repeat this step until the status of the Oracle Grid Infrastructure is online on all nodes.

---

9 **Enable the Oracle Grid Infrastructure framework resource.**

```
# clresource enable -t SUNW.crs_framework +
```

This step might require several minutes to complete.

10 **Verify that the following resources are online on all nodes:**

- The Oracle Grid Infrastructure framework resource
- The RAC database resource

```
# clresource status -t SUNW.crs_framework,SUNW.scalable_rac_server_proxy +
```

Note – Enabling the Oracle Grid Infrastructure framework resource also enable the RAC database resource. If you perform this step before the RAC database resource is enabled on all nodes, the status on some nodes might be offline. In this situation, repeat this step until the status of the RAC database resource is online on all nodes.

▼ How to Verify the Correct Behavior for Shutdown and Booting of the Cluster

If your configuration of Support for Oracle RAC is correct, Oracle Solaris Cluster ensures that Oracle RAC is stopped and started correctly when you shut down and boot the cluster.

Caution – This task requires downtime. If you are verifying the installation and configuration of Support for Oracle RAC on a production cluster where other data services are running, omit this task.

1 **Shut down the cluster.**

For instructions for performing this task, see "How to Shut Down a Cluster" in *Oracle Solaris Cluster System Administration Guide*.

2 **Confirm that the cluster shuts down correctly.**

3 **Boot the cluster.**

For instructions for performing this task, see "How to Boot a Cluster" in *Oracle Solaris Cluster System Administration Guide*.

4 **Confirm that the cluster starts correctly.**

5 **Verify that all Oracle RAC resource groups and their resources are online.**

```
# clresourcegroup status
```

Note – The bringing online of all Oracle RAC resource groups might require several minutes to complete. If you perform this step before the resource groups are enabled on all nodes, the status on some nodes might be offline. In this situation, repeat this step until the status of all Oracle RAC resource groups is online on all nodes.

◆ ◆ ◆ CHAPTER 5

Administering Support for Oracle RAC

This chapter explains how to administer Support for Oracle RAC on your Oracle Solaris Cluster nodes.

Overview of Administration Tasks for Support for Oracle RAC

Table 5–1 summarizes the administration tasks for Support for Oracle RAC.

Perform these tasks whenever they are required.

TABLE 5–1 Administration Tasks for Support for Oracle RAC

Task	Instructions
AdministerOracle RAC databases from Oracle Solaris Cluster	"Administering Oracle RAC Databases From Oracle Solaris Cluster Software" on page 102
Tune Support for Oracle RAC extension properties	"Tuning Support for Oracle RAC" on page 105
Tune Support for Oracle RAC fault monitors	"Tuning the Support for Oracle RAC Fault Monitors" on page 107
Troubleshoot Support for Oracle RAC	Chapter 6, "Troubleshooting Support for Oracle RAC"

Automatically Generated Names for Oracle Solaris Cluster Objects

When the `clsetup` utility is used to create resources, the tool assigns preset names to the resources. If you are administering resources that were created by using the `clsetup` utility, see the following table for these names.

Resource Type	Resource Name
`SUNW.rac_framework`	`rac-framework-rs`
`SUNW.scalable_rac_server_proxy`	`rac_server_proxy-rs`
`SUNW.crs_framework`	`crs_framework-rs`
`SUNW.ScalDeviceGroup`	`scal`*dg-name*`-rs`, where *dg-name* is the name of the device group that the resource represents
`SUNW.ScalMountPoint`	`scal-`*mp-dir*`-rs`, where *mp-dir* is the mount point of the file system, with `/` replaced by –
	`asm-mp-rs`
`SUNW.scalable_rac_server_proxy`	`rac_server_proxy_rs`
`SUNW.scalable_asm_diskgroup_proxy`	`asm-dg-rs`
`SUNW.scalable_oracle_asm_instance_proxy`	`asm-inst-rs`
`SUNW.LogicalHostname`	*lh-name*, where *lh-name* is the logical hostname that you specified when you created the resource

Administering Oracle RAC Databases From Oracle Solaris Cluster Software

Administering Oracle RAC databases from Oracle Solaris Cluster software involves using Oracle Solaris Cluster administration tools to modify the states of Oracle Solaris Cluster resources for Oracle RAC database instances. For information about how to create these resources, see "Configuring Resources for Oracle RAC Database Instances" on page 86.

In an Oracle Solaris Cluster configuration, with one exception, there are no general restrictions on the use of the Oracle database software command line interface (CLI) to manage Oracle Grid Infrastructure or on the Oracle `srvctl` command to manage the Oracle database and its services. The exception is that `autostart` must remain disabled in an Oracle Solaris Cluster configuration. Otherwise, you can use Oracle database software commands as you would in a configuration that does not include Oracle Solaris Cluster. The Oracle Solaris Cluster software detects the changes made by the Oracle database CLI and reacts appropriately.

Each cluster, such as the global cluster or a specific zone cluster, constitutes an independent name space for resource groups and resources. Thus, there is no name collision between names in these different name spaces. You register resource types independently for each cluster.

You can administer the resource groups and resources belonging to the global cluster only from the global-cluster voting node. You can administer the resource groups and resources belonging to a zone cluster from that zone cluster. In the global-cluster voting node, you can also administer the resource groups and resources in a zone cluster by using the `-Z` option to specify the zone cluster. You can set up inter-cluster resource dependencies or affinities only from the global-cluster voting node. The examples in the following sections typically explain the situations when the command is issued in the cluster where the resource group or resource resides.

The effects of changes to the states of Oracle Solaris Cluster resources on Oracle database components are explained in the subsections that follow:

- "Effects of State Changes to Oracle Solaris Cluster Resources for Oracle RAC Database Instances" on page 103

Effects of State Changes to Oracle Solaris Cluster Resources for Oracle RAC Database Instances

Oracle Grid Infrastructure manages the startup and shutdown of Oracle database instances, listeners, and other components that are configured in the Oracle Grid Infrastructure. Oracle Grid Infrastructure is a mandatory component. Oracle Grid Infrastructure also monitors components that are started by Oracle Grid Infrastructure and, if failures are detected, perform actions to recover from failures.

Because Oracle Grid Infrastructure manages the startup and shutdown of Oracle database components, these components cannot be stopped and started exclusively under the control of the Oracle Solaris Cluster RGM. Instead, Oracle Grid Infrastructure and the RGM interoperate so that when Oracle RAC database instances are started and stopped by Oracle Grid Infrastructure, the state of the database instances is propagated to Oracle Solaris Cluster resources.

The following table describes the state changes that occur between Oracle Solaris Cluster resources and Oracle Grid Infrastructure resources.

TABLE 5–2 Propagation of State Changes Between Oracle Solaris Cluster Resources and Oracle Grid Infrastructure Resources

Trigger	Initial State: Oracle Solaris Cluster Resource	Initial State: Oracle Grid Infrastructure Resource	Resulting State: Oracle Solaris Cluster Resource	Resulting State: Oracle Grid Infrastructure Resource
Oracle Solaris Cluster command to take offline a resource	Enabled and online	Enabled and online	Enabled and offline	Enabled and offline
Oracle Grid Infrastructure command to stop a resource	Enabled and online	Enabled and online	Enabled and offline	Enabled and offline
Oracle Solaris Cluster command to bring online a resource	Enabled and offline	Enabled and offline	Enabled and online	Enabled and online
Oracle Grid Infrastructure command to start a resource	Enabled and offline	Enabled and offline	Enabled and online	Enabled and online
Oracle Solaris Cluster command to disable a resource	Enabled and online	Enabled and online	Disabled and offline	Disabled and offline
Oracle Grid Infrastructure command to disable a resource	Enabled and online	Enabled and online	Enabled and online	Disabled and online
Oracle `SQLPLUS` command to shut down the database	Enabled and online	Enabled and online	Enabled and offline	Enabled and offline
Oracle Solaris Cluster command to enable a resource	Disabled and offline	Disabled and offline	Enabled and online or offline	Enabled and online or offline
Oracle Grid Infrastructure command to enable a resource	Disabled and offline	Disabled and offline	Disabled and offline	Enabled and offline

The names of the states of Oracle Solaris Cluster resources and Oracle Grid Infrastructure resources are identical. However, the meaning of each state name is different for Oracle Solaris Cluster resources and Oracle Grid Infrastructure resources. For more information, see the following table.

TABLE 5–3 Comparisons of States for Oracle Solaris Cluster Resources and Oracle Grid Infrastructure Resources

State	Meaning for Oracle Solaris Cluster Resources	Meaning for Oracle Grid Infrastructure Resources
Enabled	The resource is available to the Oracle Solaris Cluster RGM for automatic startup, failover, or restart. A resource that is enabled can also be in either the online state or the offline state.	The resource is available to run under Oracle Grid Infrastructure for automatic startup, failover, or restart. A resource that is enabled can also be in either the online state or the offline state.
Disabled	The resource is unavailable to the Oracle Solaris Cluster RGM for automatic startup, failover, or restart. A resource that is disabled is also offline.	The resource is unavailable to run under Oracle Grid Infrastructure for automatic startup, failover, or restart. A resource that is disabled can also be in either the online state or the offline state.
Online	The resource is running and providing service.	The resource is running and providing service. A resource that is online must also be enabled.
Offline	The resource is stopped and not providing service.	The resource is stopped and not providing service. A resource that is offline can also be in either the disabled state or the enabled state.

For detailed information about the state of Oracle Solaris Cluster resources, see "Resource and Resource Group States and Settings" in *Oracle Solaris Cluster Concepts Guide*.

For detailed information about the state of Oracle Grid Infrastructure resources, see your Oracle documentation.

Tuning Support for Oracle RAC

To tune the Support for Oracle RAC data service, you modify the extension properties of the resources for this data service. For details about these extension properties, see Appendix C, "Support for Oracle RAC Extension Properties." Typically, you use the option **-p** *property*=*value* of the `clresource` command to set extension properties of Support for Oracle RAC resources. You can also use the procedures in Chapter 2, "Administering Data Service Resources," in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure the resources later.

This section contains the following information about tuning the Support for Oracle RAC data service:

Guidelines for Setting Timeouts

Many of the extension properties for Support for Oracle RAC specify timeouts for steps in reconfiguration processes. The optimum values for most of these timeouts are independent of your cluster configuration. Therefore, you should not need to change the timeouts from their default values.

Timeouts that depend on your cluster configuration are described in the subsections that follow. If timeouts occur during reconfiguration processes, increase the values of these timeout properties to accommodate your cluster configuration.

This section describes the following timeouts:

Reservation Step Timeout

The time that is required for reservation commands to run is affected by the following factors:

The time that is required for reservation commands to run is affected by the following factors:

- The number of shared physical disks in the cluster
- The load on the cluster

If the number of shared physical disks in the cluster is large, or if your cluster is heavily loaded, the reconfiguration of Support for Oracle RAC might time out. If such a timeout occurs, increase the reservation step timeout.

To increase the reservation step timeout, increase the `Reservation_timeout` extension property of the `SUNW.rac_framework` resource.

For more information, see "`SUNW.rac_framework` Extension Properties" on page 156.

EXAMPLE 5–1 Setting the Reservation Step Timeout

```
# clresource set -p reservation_timeout=350 rac-framework-rs
```

This example sets the timeout for the reservation step of a reconfiguration of Support for Oracle RAC to 350 seconds. This example assumes that the Oracle RAC framework component is represented by an instance of the `SUNW.rac_framework` resource type that is named `rac-framework-rs`.

Tuning the Support for Oracle RAC Fault Monitors

Fault monitoring for the Support for Oracle RAC data service is provided by fault monitors for the following resources:

- Scalable device group resource
- Scalable file-system mount-point resource

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

TABLE 5–4 Resource Types for Support for Oracle RAC Fault Monitors

Fault Monitor	Resource Type
Scalable device group	`SUNW.ScalDeviceGroup`
Scalable file-system mount point	`SUNW.ScalMountPoint`

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 Support for Oracle RAC fault monitors *only* if you need to modify this preset behavior.

Tuning the Support for Oracle RAC 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 Support for Oracle RAC fault monitors that you need to perform these tasks is provided in the subsections that follow:

- "Operation of the Fault Monitor for a Scalable Device Group" on page 107
- "Operation of the Fault Monitor for Scalable File-System Mount Points" on page 108
- "Obtaining Core Files for Troubleshooting DBMS Timeouts" on page 108

Operation of the Fault Monitor for a Scalable Device Group

By default, the fault monitor monitors all logical volumes in the device group that the resource represents. If you require only a subset of the logical volumes in a device group to be monitored, set the `LogicalDeviceList` extension property.

The status of the device group is derived from the statuses of the individual logical volumes that are monitored. If all monitored logical volumes are healthy, the device group is healthy. If any monitored logical volume is faulty, the device group is faulty. If a device group is discovered to be faulty, monitoring of the resource that represents the group is stopped and the resource is put into the disabled state.

The status of an individual logical volume is obtained by querying the volume's volume manager. If the status of a Solaris Volume Manager for Sun Cluster volume cannot be determined from a query, the fault monitor performs file input/output (I/O) operations to determine the status.

Note – For mirrored disks, if one submirror is faulty, the device group is still considered to be healthy.

If reconfiguration of userland cluster membership causes an I/O error, the monitoring of device group resources by fault monitors is suspended while userland cluster membership monitor (UCMM) reconfigurations are in progress.

Operation of the Fault Monitor for Scalable File-System Mount Points

To determine if the mounted file system is available, the fault monitor performs I/O operations such as opening, reading, and writing to a test file on the file system. If an I/O operation is not completed within the timeout period, the fault monitor reports an error. To specify the timeout for I/O operations, set the `IOTimeout` extension property.

The response to an error depends on the type of the file system, as follows:

- If the file system is an NFS file system on a qualified NAS device, the response is as follows:
  - Monitoring of the resource is stopped on the current node.
  - The resource is placed into the disabled state on the current node, causing the file system to be unmounted from that node.

Obtaining Core Files for Troubleshooting DBMS Timeouts

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

◆ ◆ ◆ CHAPTER 6

Troubleshooting Support for Oracle RAC

If you encounter a problem with Support for Oracle RAC, troubleshoot the problem by using the techniques that are described in the following sections.

Verifying the Status of Support for Oracle RAC

The status of resource groups and resources for Support for Oracle RAC indicates the status of Oracle RAC in your cluster. Use Oracle Solaris Cluster maintenance commands to obtain this status information.

- To obtain status information for resource groups, use the `clresourcegroup`(1CL) command.
- To obtain status information for resources, use the `clresource`(1CL) command.

▼ How to Verify the Status of Support for Oracle RAC

1. **Become superuser or assume a role that provides `solaris.cluster.read` RBAC authorization.**
2. **Display status information for the Oracle Solaris Cluster objects in which you are interested.**

   For example:

   - **To display status information for all resource groups in your cluster, type the following command:**

     ```
     # clresourcegroup status +
     ```

- **To display status information for all resources in a resource group, type the following command:**

  ```
  # clresource status -g resource-group +
  ```

  resource-group
  Specifies the resource group that contains the resources whose status information you are displaying.

See Also For information about options that you can specify to filter the status information that is displayed, see the following man pages:

- `clresource(1CL)`
- `clresourcegroup(1CL)`

Sources of Diagnostic Information

If the state of a scalable device group resource or a file-system mount-point resource changes, the new state is logged through the `syslog(3C)` function.

The directories `/var/cluster/ucmm` and `/var/cluster/vucmm` contain the sources of diagnostic information that are shown in the following table.

Source	Location
Log files for previous multiple-owner volume-manager reconfigurations	`/var/cluster/vucmm/vucmm_reconf.log.0` (`0`,`1`,...)
Log file for the current userland cluster membership monitor (UCMM) reconfiguration	`/var/cluster/ucmm/ucmm_reconf.log`
Log files for previous UCMM reconfigurations	`/var/cluster/ucmm/ucmm_reconf.log.0` (`0`,`1`,...)

The directory `/var/opt/SUNWscor/oracle_server/proxy`*resource* contains log files for the resource that represents the Oracle RAC proxy server. Messages for server-side components and client-side components of the proxy server resource are written to separate files:

- Messages for server-side components are written to the file `message_log.`*resource*.
- Messages for client-side components are written to the file `message_log.client.`*resource*.

In these file names and directory names, *resource* is the name of the resource that represents the Oracle RAC server component.

The system messages file also contains diagnostic information.

If a problem occurs with Support for Oracle RAC, consult these files to obtain information about the cause of the problem.

Common Problems and Their Solutions

The subsections that follow describe problems that can affect Support for Oracle RAC. Each subsection provides information about the cause of the problem and a solution to the problem.

Failure of an Oracle RAC Framework Resource Group

This section describes problems that can affect the Oracle RAC framework resource group.

Node Panic During Initialization of Support for Oracle RAC

If a fatal problem occurs during the initialization of Support for Oracle RAC, the node panics with an error messages similar to the following error message:

```
panic[cpu0]/thread=40037e60: Failfast: Aborting because "ucmmd" died 30 seconds
ago
```

Description: A component that the UCMM controls returned an error to the UCMM during a reconfiguration.

Cause: The most common causes of this problem are as follows:

A node might also panic during the initialization of Support for Oracle RAC because a reconfiguration step has timed out. For more information, see "Node Panic Caused by a Timeout" on page 115.

Solution: For instructions to correct the problem, see "How to Recover From a Failure of the `ucmmd` Daemon or a Related Component" on page 112.

Note – When the node is a global-cluster voting node of the global cluster, the node panic brings down the entire machine. When the node is a zone-cluster node, the node panic brings down only that specific zone and other zones remain unaffected.

Failure of the ucmmd Daemon to Start

The UCMM daemon, `ucmmd`, manages the reconfiguration of Support for Oracle RAC. When a cluster is booted or rebooted, this daemon is started only after all components of Support for Oracle RAC are validated. If the validation of a component on a node fails, the `ucmmd` daemon fails to start on the node.

The most common causes of this problem are as follows:

- An error occurred during a previous reconfiguration of a component Support for Oracle RAC.
- A step in a previous reconfiguration of Support for Oracle RAC timed out, causing the node on which the timeout occurred to panic.

For instructions to correct the problem, see "How to Recover From a Failure of the `ucmmd` Daemon or a Related Component" on page 112.

▼ How to Recover From a Failure of the ucmmd Daemon or a Related Component

Perform this task to correct the problems that are described in the following sections:

- "Node Panic During Initialization of Support for Oracle RAC" on page 111
- "Failure of the `ucmmd` Daemon to Start" on page 112

1 To determine the cause of the problem, examine the log files for UCMM reconfigurations and the system messages file.

For the location of the log files for UCMM reconfigurations, see "Sources of Diagnostic Information" on page 110.

When you examine these files, start at the most recent message and work backward until you identify the cause of the problem.

For more information about error messages that might indicate the cause of reconfiguration errors, see *Oracle Solaris Cluster Error Messages Guide*.

2. **Correct the problem that caused the component to return an error to the UCMM.**

   For example:

   - **If a reconfiguration step has timed out, increase the value of the extension property that specifies the timeout for the step.**

     For more information, see "Node Panic Caused by a Timeout" on page 115.

3. **If the solution to the problem requires a reboot, reboot the node where the problem occurred.**

   The solution to only certain problems requires a reboot. For example, increasing the amount of shared memory requires a reboot. However, increasing the value of a step timeout does *not* require a reboot.

   For more information about how to reboot a node, see "Shutting Down and Booting a Single Node in a Cluster" in *Oracle Solaris Cluster System Administration Guide*.

4. **On the node where the problem occurred, take offline and bring online the Oracle RAC framework resource group.**

   This step refreshes the resource group with the configuration changes you made.

   a. **Become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.**

   b. **Type the command to take offline the Oracle RAC framework resource group and its resources.**

      ```
      # clresourcegroup offline -n node rac-fmwk-rg
      ```

      `-n` *node*
      Specifies the node name or node identifier (ID) of the node where the problem occurred.

      rac-fmwk-rg
      Specifies the name of the resource group that is to be taken offline.

   c. **Type the command to bring online and in a managed state the Oracle RAC framework resource group and its resources.**

      ```
      # clresourcegroup online -emM -n node rac-fmwk-rg
      ```

Failure of a Multiple-Owner Volume-Manager Framework Resource Group

This section describes problems that can affect the multiple-owner volume-manager framework resource group.

- "Node Panic During Initialization of the Multiple-Owner Volume-Manager Framework" on page 114

- "Failure of the vucmmd Daemon to Start" on page 114
- "How to Recover From a Failure of the vucmmd Daemon or a Related Component" on page 114

Node Panic During Initialization of the Multiple-Owner Volume-Manager Framework

If a fatal problem occurs during the initialization of the multiple-owner volume-manager framework, the node panics with an error messages similar to the following error message:

Note – When the node is a global-cluster voting node of the global cluster, the node panic brings down the entire machine.

Failure of the vucmmd Daemon to Start

The multiple-owner volume-manager framework daemon, vucmmd, manages the reconfiguration of the multiple-owner volume-manager framework. When a cluster is booted or rebooted, this daemon is started only after all components of the multiple-owner volume-manager framework are validated. If the validation of a component on a node fails, the vucmmd daemon fails to start on the node.

The most common causes of this problem are as follows:

- An error occurred during a previous reconfiguration of a component of the multiple-owner volume-manager framework.
- A step in a previous reconfiguration of the multiple-owner volume-manager framework timed out, causing the node on which the timeout occurred to panic.

For instructions to correct the problem, see "How to Recover From a Failure of the vucmmd Daemon or a Related Component" on page 114.

▼ How to Recover From a Failure of the vucmmd Daemon or a Related Component

Perform this task to correct the problems that are described in the following sections:

- "Node Panic During Initialization of the Multiple-Owner Volume-Manager Framework" on page 114
- "Failure of the vucmmd Daemon to Start" on page 114

1 To determine the cause of the problem, examine the log files for multiple-owner volume-manager framework reconfigurations and the system messages file.

For the location of the log files for multiple-owner volume-manager framework reconfigurations, see "Sources of Diagnostic Information" on page 110.

When you examine these files, start at the most recent message and work backward until you identify the cause of the problem.

For more information about error messages that might indicate the cause of reconfiguration errors, see *Oracle Solaris Cluster Error Messages Guide*.

2 **Correct the problem that caused the component to return an error to the multiple-owner volume-manager framework .**

3 **If the solution to the problem requires a reboot, reboot the node where the problem occurred.**

The solution to only certain problems requires a reboot. For example, increasing the amount of shared memory requires a reboot. However, increasing the value of a step timeout does *not* require a reboot.

For more information about how to reboot a node, see "Shutting Down and Booting a Single Node in a Cluster" in *Oracle Solaris Cluster System Administration Guide*.

4 **On the node where the problem occurred, take offline and bring online the multiple-owner volume-manager framework resource group.**

This step refreshes the resource group with the configuration changes you made.

a. **Become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.**

b. **Type the command to take offline the multiple-owner volume-manager framework resource group and its resources.**

```
# clresourcegroup offline -n node vucmm-fmwk-rg
```

`-n` *node*
Specifies the node name or node identifier (ID) of the node where the problem occurred.

vucmm-fmwk-rg
Specifies the name of the resource group that is to be taken offline.

c. **Type the command to bring online and in a managed state the multiple-owner volume-manager framework resource group and its resources.**

```
# clresourcegroup online -emM -n node vucmm-fmwk-rg
```

Node Panic Caused by a Timeout

The timing out of any step in the reconfiguration of Support for Oracle RAC causes the node on which the timeout occurred to panic.

To prevent reconfiguration steps from timing out, tune the timeouts that depend on your cluster configuration. For more information, see "Guidelines for Setting Timeouts" on page 106.

If a reconfiguration step times out, use the Oracle Solaris Cluster maintenance commands to increase the value of the extension property that specifies the timeout for the step. For more information, see Appendix C, "Support for Oracle RAC Extension Properties."

After you have increased the value of the extension property, bring online the Oracle RAC framework resource group on the node that panicked.

Failure of a SUNW.rac_framework or SUNW.vucmm_framework Resource to Start

If a `SUNW.rac_framework` or `SUNW.vucmm_framework`resource fails to start, verify the status of the resource to determine the cause of the failure. For more information, see "How to Verify the Status of Support for Oracle RAC" on page 109.

The state of a resource that failed to start is shown as `Start failed`. The associated status message indicates the cause of the failure to start.

This section contains the following information:

SUNW.rac_framework Failure-to-Start Status Messages

The following status messages are associated with the failure of a `SUNW.rac_framework` resource to start:

`Faulted - ucmmd is not running`

Description: The `ucmmd` daemon is not running on the node where the resource resides.

Solution: For information about how to correct this problem, see "Failure of the `ucmmd` Daemon to Start" on page 112.

`Degraded - reconfiguration in progress`

Description: The UCMM is undergoing a reconfiguration. This message indicates a problem only if the reconfiguration of the UCMM is not completed and the status of this resource persistently remains degraded.

Cause: If this message indicates a problem, the cause of the failure is a configuration error in one or more components of Support for Oracle RAC.

Solution: The solution to this problem depends on whether the message indicates a problem:

- If the message indicates a problem, correct the problem as explained in "How to Recover From a Failure of the ucmmd Daemon or a Related Component" on page 112.
- If the message does not indicate a problem, no action is required.

`Online`
Description: Reconfiguration of Oracle RAC was not completed until after the `START` method of the `SUNW.rac_framework` resource timed out.

Solution: For instructions to correct the problem, see "How to Recover From the Timing Out of the START Method" on page 118.

SUNW.vucmm_framework Failure-to-Start Status Messages

The following status messages are associated with the failure of a `SUNW.vucmm_framework` resource to start:

`Faulted - vucmmd is not running`
Description: The `vucmmd` daemon is not running on the node where the resource resides.

Solution: For information about how to correct this problem, see "Failure of the vucmmd Daemon to Start" on page 114.

`Degraded - reconfiguration in progress`
Description: The multiple-owner volume-manager framework is undergoing a reconfiguration. This message indicates a problem only if the reconfiguration of the multiple-owner volume-manager framework is not completed and the status of this resource persistently remains degraded.

Cause: If this message indicates a problem, the cause of the failure is a configuration error in one or more components of the volume manager reconfiguration framework.

Solution: The solution to this problem depends on whether the message indicates a problem:

- If the message indicates a problem, correct the problem as explained in "How to Recover From a Failure of the vucmmd Daemon or a Related Component" on page 114.
- If the message does not indicate a problem, no action is required.

`Online`
Description: Reconfiguration of Oracle RAC was not completed until after the `START` method of the `SUNW.vucmm_framework` resource timed out.

Solution: For instructions to correct the problem, see "How to Recover From the Timing Out of the START Method" on page 118.

How to Recover From the Timing Out of the START Method

1. **Become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.**

2. **On the node where the `START` method timed out, take offline the framework resource group that failed to start.**

   To perform this operation, switch the primary nodes of the resource group to the other nodes where the group is online.

   ```
   # clresourcegroup offline -n nodelist resource-group
   ```

   `-n` *nodelist*
   Specifies a comma-separated list of other cluster nodes on which *resource-group* is online. Omit from this list the node where the `START` method timed out.

   resource-group
   Specifies the name of the framework resource group.

   If your configuration uses both a multiple-owner volume-manager framework resource group and an Oracle RAC framework resource group, first take offline the multiple-owner volume-manager framework resource group. When the multiple-owner volume-manager framework resource group is offline, then take offline the Oracle RAC framework resource group.

   If the Oracle RAC framework resource group was created by using the `clsetup` utility, the name of the resource group is `rac-framework-rg`.

3. **On all cluster nodes that can run Support for Oracle RAC, bring online the framework resource group that failed to come online.**

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

   resource-group
   Specifies that the resource group that you brought offline in Step 2 is to be moved to the `MANAGED` state and brought online.

Failure of a Resource to Stop

If a resource fails to stop, correct this problem as explained in "Clearing the STOP_FAILED Error Flag on Resources" in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

◆ ◆ ◆ CHAPTER 7

Modifying an Existing Configuration of Support for Oracle RAC

This chapter explains how to modify an existing configuration of Support for Oracle RAC.

Overview of Tasks for Modifying an Existing Configuration of Support for Oracle RAC

Table 7–1 summarizes the administration tasks for Support for Oracle RAC.

Perform these tasks whenever they are required.

TABLE 7–1 Tasks for Modifying an Existing Configuration of Support for Oracle RAC

Task	Instructions
Modify online the resource for a scalable device group	"Modifying Online the Resource for a Scalable Device Group" on page 120
Extend an existing configuration of Support for Oracle RAC	"Extending an Existing Configuration of Support for Oracle RAC" on page 120
Remove Support for Oracle RAC	"Removing Support for Oracle RAC" on page 127

Modifying Online the Resource for a Scalable Device Group

Modifying online the resource for a scalable device group involves changing the list of logical volumes that are to be monitored. The `LogicalDeviceList` extension property of the `SUNW.ScalDeviceGroup` resource type specifies the list of logical volumes in a global device group that are to be monitored.

▼ How to Modify Online the Resource for a Scalable Device Group

1. **Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

2. **Modify the `LogicalDeviceList` extension property of the `ScalDeviceGroup` resource.**

   - **To add device groups to a `ScalDeviceGroup` resource, type the following command:**

     ```
     # clresource set -p LogicalDeviceList+=logical-device-listscal-mp-rs
     ```

     The addition of the logical volume is effective immediately.

   - **To remove device groups from a `ScalDeviceGroup` resource, type the following command:**

     ```
     # clresource set -p LogicalDeviceList-=logical-device-listscal-mp-rs
     ```

     The removal of the logical volume is effective immediately.

Extending an Existing Configuration of Support for Oracle RAC

Extend an existing configuration of Support for Oracle RAC in any of the following situations:

- You are adding nodes to a cluster and you require Support for Oracle RAC to run on the nodes. See "How to Add Support for Oracle RAC to Selected Nodes " on page 120.
- You are adding a volume manager. See "How to Add a Volume Manager Resource to the `SUNW.vucmm_framework` Resource Group" on page 124.

▼ How to Add Support for Oracle RAC to Selected Nodes

Perform this procedure if you are adding nodes to a cluster and you require Support for Oracle RAC to run on the nodes. Perform this procedure from only one node.

This task involves adding the selected nodes from the following resource groups in the following order:

- Resource groups for scalable file-system mount-point resources
- The multiple-owner volume-manager framework resource group, if used.
- The Oracle RAC framework resource group
- Resource groups for scalable device group resources
- Resource groups for logical hostname resources
- The resource group for the Oracle RAC database

Before You Begin

- Ensure that the required Support for Oracle RAC software packages are installed on each node to which you are adding Support for Oracle RAC. For more information, see "Installing the Support for Oracle RAC Package" on page 35.
- Ensure that the node that you add is connected to the shared storage that is used by the Oracle RAC configuration.

1 Become superuser on any cluster node.

2 Add the nodes to any resource groups that contain scalable file-system mount-point resources.

If no resource groups that contain scalable file-system mount-point resources are configured, omit this step.

For each resource group to which you are adding nodes, run the following command:

```
# clresourcegroup add-node -S -n nodelist scal-mp-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes to which you are adding Support for Oracle RAC.

scal-mp-rg
Specifies the name of the resource group to which you are adding nodes.

3 Add the nodes to the multiple-owner volume-manager framework resource group, if used.

```
# clresourcegroup add-node -S -n nodelist vucmm-fmwk-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes to which you are adding the resource group.

vucmm-fmwk-rg
Specifies the name of the resource group to which you are adding nodes.

4 Add the nodes to the Oracle RAC framework resource group.

```
# clresourcegroup add-node -S -n nodelist rac-fmwk-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes to which you are adding Support for Oracle RAC.

rac-fmwk-rg
Specifies the name of the resource group to which you are adding nodes.

5 **Add the nodes to any scalable device groups that you are using for Oracle files.**

If you are not using any scalable device groups for Oracle files, omit this step.

How to perform this step depends on the type of the scalable device group.

- **For each Solaris Volume Manager for Sun Cluster multi-owner disk set, type the following command:**

  ```
  # metaset -s set-name -M -a -h nodelist
  ```

  `-s` *set-name*
  Specifies the Solaris Volume Manager for Sun Cluster multi-owner disk set to which you are adding nodes.

  `-h` *nodelist*
  Specifies a *space-separated* list of cluster nodes that you are adding to the multi-owner disk set.

6 **Add the nodes to any resource groups that contain scalable device group resources.**

If no resource groups that contain scalable device group resources are configured, omit this step.

For each resource group to which you are adding nodes, run the following command:

```
# clresourcegroup add-node -S -n nodelist scal-dg-rg
```

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes to which you are adding Support for Oracle RAC.

scal-dg-rg
Specifies the name of the resource group to which you are adding nodes.

7 **Mount each shared file system that is to be accessed from the nodes that you are adding.**

If no shared file systems are to be accessed from the nodes that you are adding, omit this step.

For each file system that you are mounting, type the following command:

```
# mount mount-point
```

mount-point
Specifies the mount point of the file system that you are mounting.

8 **Bring online all resource groups to which you added nodes in Step 6.**

These resource groups contain scalable device group resources.

If no resource groups that contain scalable device group resources are configured, omit this step.

For each resource group that you are bringing online, type the following command:

```
# clresourcegroup online scal-dg-rg
```

scal-dg-rg
Specifies the name of the resource group that you are bringing online.

9 Start Oracle Grid Infrastructure.

```
# /etc/init.d/init.crs start
Startup will be queued to init within 30 seconds.
```

10 Add the nodes to the resource group for each Oracle RAC database that is to run on the nodes.

For each resource group to which you are adding nodes, run the following command:

```
# clresourcegroup add-node -S -n nodelist rac-db-rg
```

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes to which you are adding Support for Oracle RAC.

rac-db-rg
Specifies the name of the resource group to which you are adding nodes.

11 For each node that you are adding, create the Oracle Grid Infrastructure resources that are required to represent Oracle Solaris Cluster resources.

Create an Oracle Grid Infrastructure resource for each Oracle Solaris Cluster resource for scalable device groups and scalable file-system mount points on which Oracle components depend. For more information, see "How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster" on page 185.

12 Modify each resource for Oracle RAC databases to set a value of each per-node property for each node that you are adding.

For each resource that you are modifying, perform these steps:

a. Disable the resource.

```
# clresource disable rac-db-rs
```

rac-db-rs
Specifies the name of the Oracle RAC database resource that you are disabling.

b. Set a value of each per-node property for each node that you are adding.

The per-node properties of each resource type for Oracle RAC databases are shown in the following table.

Resource Type	Properties
SUNW.scalable_rac_server_proxy	oracle_sid

For information about extension properties of resource types for Oracle RAC databases, see the following sections:

- "SUNW.scalable_rac_server_proxy Extension Properties" on page 159

```
# clresource set \
-p property{node}=value[...] \
[-p property{node}=value[...]][...] \
rac-db-rs
```

property
Specifies the name of a per-node property that you are setting.

node
Specifies the node for which you are setting a value for *property*.

value
Specifies the value to which you are setting *property* for *node*.

rac-db-rs
Specifies the name of the Oracle RAC database resource whose per-node properties you are setting.

c. Enable the resource.

```
# clresource enable rac-db-rs
```

rac-db-rs
Specifies the name of the Oracle RAC database resource that you are enabling.

13 Bring online each resource group for Oracle RAC databases.

For each resource group that you are bringing online, run the following command:

```
# clresourcegroup online rac-db-rg
```

rac-db-rg
Specifies the name of the resource group that you are bringing online.

▼ How to Add a Volume Manager Resource to the SUNW.vucmm_framework Resource Group

Perform this task if you are adding a volume manager to an existing configuration of Support for Oracle RAC. The `SUNW.vucmm_framework` resource group must contain a resource that represents the volume manager that you are adding. You can add a volume manager resource only if the framework resource is disabled and if the framework daemon is stopped on all cluster nodes.

Caution – This task requires downtime because you must disable the framework resource and reboot the nodes where Oracle RAC is running.

Before You Begin Ensure that the volume manager for which you are adding a resource is installed and configured on all nodes where Oracle RAC is to run.

1. **Become superuser on any cluster node.**

2. **Disable the framework resource in the framework resource group and any other resources that depend on this resource.**

   ```
   # clresource disable -r fmwk-rs
   ```

   fmwk-rs
   Specifies the name of the resource of type `SUNW.vucmm_framework` that you are disabling.

3. **Reboot all the nodes that are in the node list of the framework resource group.**

4. **Register and add an instance of the resource type that represents the volume manager that you are adding.**

   - **If you are adding Solaris Volume Manager for Sun Cluster, register and add the instance as follows:**

     a. **Register the Solaris Volume Manager for Sun Cluster resource type.**

        ```
        # clresourcetype register SUNW.vucmm_svm
        ```

     b. **Add an instance of the Solaris Volume Manager for Sun Cluster resource type to the framework resource group.**

        Ensure that this instance depends on the resource that you disabled in Step 2.

        ```
        # clresource create -g fmwk-rg \
          -t svm-rt \
          -p resource_dependencies=fmwk-rs  svm-rs
        ```

        `-g` *fmwk-rg*
        Specifies the name of the framework resource group. This resource group contains the resource of type `SUNW.vucmm_framework` that you disabled in Step 2.

        svm-rt
        Specifies the name of the Solaris Volume Manager for Sun Cluster resource type.

        `-p resource_dependencies=`*fmwk-rs*
        Specifies that this instance depends on the resource that you disabled in Step 2.

        svm-rs
        Specifies the name that you are assigning to the resource of type `SUNW.vucmm_svm`.

5 **Bring online and in a managed state the framework resource group and its resources.**

```
# clresourcegroup online -emM fmwk-rg
```

fmwk-rg
Specifies that the framework resource group is to be moved to the `MANAGED` state and brought online. This resource group contains the resource that you disabled in Step 2.

Next Steps Go to "How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database" on page 60.

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource:

▼ How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1 **Become superuser.**

2 **Display the current start dependency that the database has on the Oracle Grid Infrastructure `storage_proxy` resource.**

```
# Grid_home/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=hard(sun.grid-storage-proxy-rs) weak(type:ora.listener.type,
global:type:ora.scan_listener.type,uniform:ora.ons,uniform:ora.eons)
# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs
=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs scal-dg1-rs
```

3 **Remove the offline-restart dependency on the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource from the Oracle RAC instance proxy resource.**

This command clears the dependencies that the Oracle Grid Infrastructure database resource has on the Oracle Grid Infrastructure `storage_proxy` resource. Note that the command includes the minus (-) symbol.

```
# clresource set -p resource_dependencies_offline_restart-=scal-dg1-rs rac-server-proxy-rs
```

4 **Verify that the start dependency on the Oracle Grid Infrastructure resource is removed.**

```
# Grid_home/bin/crsctl stat res ora.testdb.db -p | grep START_DEPENDENCIES
START_DEPENDENCIES=weak(type:ora.listener.type,global:type:ora.scan_listener.type,
```

```
uniform:ora.ons,uniform:ora.eons)
# clresource show -p resource_dependencies_offline_restart rac-server-proxy-rs

=== Resources ===

Resource: rac-server-proxy-rs
Resource_dependencies_offline_restart: crs-fw-rs
```

▼ How to Delete the sun.*resource* Resource

1. **Become superuser.**
2. **Ensure that the dependency is removed, as described in "How to Remove a Dependency" on page 126, and that `sun.`*resource* is stopped.**

   ```
   # Grid_home/bin/crsctl stop res sun.scal-dg1-rs
   CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx3'
   CRS-2673: Attempting to stop 'sun.scal-dg1-rs' on 'pnsx1'
   CRS-2673: Attempting to stop 'sun.scal-dg1-rss' on 'pnsx2'
   CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx3' succeeded
   CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx1' succeeded
   CRS-2677: Stop of 'sun.scal-dg1-rs' on 'pnsx2' succeeded
   ```
3. **Delete `sun.`*resource*.**

   ```
   # Grid_home/bin/crsctl delete res sun.scal-dg1-rs
   ```
4. **Verify that `sun.`*resource* is deleted.**

   ```
   # Grid_home/bin/crsctl stat res sun.scal-dg1-rs -p
   CRS-210: Could not find resource 'sun.scal-dg1-rs'.
   ```

Removing Support for Oracle RAC

You can remove Support for Oracle RAC from the following entities:

- **A cluster.** See "How to Remove Support for Oracle RAC From a Cluster" on page 127.
- **Selected nodes in a cluster.** See "How to Remove Support for Oracle RAC From Selected Nodes" on page 131.

▼ How to Remove Support for Oracle RAC From a Cluster

Perform this task to remove Support for Oracle RAC from all nodes in a cluster.

On a cluster where multiple Oracle RAC databases are running, perform this task to remove an Oracle RAC database from the cluster. The remaining Oracle RAC databases continue to run in the cluster.

This task involves removing the following resource groups from the cluster in the following order:

- The resource group for the Oracle RAC database
- Resource groups for logical hostname resources
- Resource groups for scalable file-system mount-point resources
- Resource groups for scalable device group resources
- The Oracle RAC framework resource group
- The multiple-owner volume-manager framework resource group, if used

Caution – You might perform this task to remove an Oracle RAC database from a cluster where multiple Oracle RAC databases are running. In this situation, do *not* remove any resource group on whose resources the remaining Oracle RAC databases depend.

For example, you might have configured multiple database file systems to depend on a single device group. In this situation, do *not* remove the resource group that contains the resource for the scalable device group.

Similarly, if multiple databases depend on the Oracle RAC framework resource group, do *not* remove this resource group.

Before You Begin Ensure that the cluster node from which you perform this task is booted in cluster mode.

1 On one node of the cluster, become superuser.

2 Remove the resource group for each Oracle RAC database that you are removing.

For each Oracle RAC database that you are removing, type the following command:

```
# clresourcegroup delete -F rac-db-rg
```

rac-db-rg
Specifies the resource group that you are removing.

3 Use Oracle utilities to remove from the cluster each Oracle RAC database that you no longer require.

4 If you are removing Support for Oracle RAC entirely, use Oracle utilities to remove the following items from all nodes in a cluster:

- The Oracle RAC software
- The Oracle Grid Infrastructure software

5 Disable the Oracle Grid Infrastructure framework resource.

```
# clresource disable crs-framework-rs
```

crs-framework-rs
Specifies the name of the resource that you are disabling. This resource is the instance of the `SUNW.crs_framework` resource type that is configured in the cluster.

6 Remove any resource groups that contain scalable file-system mount-point resources.

If no resource groups that contain scalable file-system mount-point resources are configured, omit this step.

For each resource group that you are removing, type the following command:

```
# clresourcegroup delete -F scal-mp-rg
```

scal-mp-rg
Specifies the resource group that you are removing.

7 Remove any resource groups that contain scalable device group resources.

If no resource groups that contain scalable device group resources are configured, omit this step.

For each resource group that you are removing, type the following command:

```
# clresourcegroup delete -F scal-dg-rg
```

scal-dg-rg
Specifies the resource group that you are removing.

8 Destroy any scalable device groups that were affected by the removal of resource groups in Step 7.

How to perform this step depends on the type of the scalable device group.

- **For each Solaris Volume Manager for Sun Cluster multi-owner disk set, destroy the disk set as follows:**

  a. Remove all metadevices such as volumes, soft partitions, and mirrors from the disk set.

  Use the `metaclear(1M)` command for this purpose.

  ```
  # metaclear -s scal-dg-ms -a
  ```

  `-s` *scal-dg-ms*
  Specifies the name of the disk set from which you are removing metadevices.

  b. Remove all global devices from the disk set.

  ```
  # metaset -s scal-dg-ms -d -f alldevices
  ```

  `-s` *scal-dg-ms*
  Specifies the name of the disk set from which you are removing global devices.

  alldevices
  Specifies a space-separated list that contains *all* global devices that were added to the disk set when the disk set was created. The format of each device ID path name is

`/dev/did/dsk/d`*N*, where *N* is the device number.

c. Remove all nodes from the disk set that you are destroying.

The removal of all nodes from a disk set destroys the disk set.

```
# metaset -s scal-dg-ms -d -h allnodes
```

`-s` *scal-dg-ms*
Specifies the name of the disk set that you are destroying.

`-h` *allnodes*
Specifies a space-separated list that contains *all* nodes that were added to the disk set when the disk set was created.

Note – If you are removing an Oracle RAC database from a cluster where multiple Oracle RAC databases are running, omit the remaining steps in this procedure.

---

9 Remove the Oracle RAC framework resource group.

```
# clresourcegroup delete -F rac-fmwk-rg
```

rac-fmwk-rg
Specifies the resource group that you are removing.

10 Remove the multiple-owner volume-manager framework resource group, if used.

```
# clresourcegroup delete -F vucmm-fmwk-rg
```

vucmm-fmwk-rg
Specifies the resource group that you are removing.

11 Unregister the resource type of each resource that you removed in this procedure.

```
# clresourcetype unregister resource-type-list
```

resource-type-list
Specifies a comma-separated list of the names of the resource types that you are unregistering. For a list of the resource types that are associated with Support for Oracle RAC, see "Automatically Generated Names for Oracle Solaris Cluster Objects" on page 102.

12 (Optional) From each node in the cluster, uninstall the Support for Oracle RAC software packages.

```
# pkg uninstall ha-cluster/data-service/oracle-database ha-cluster/library/ucmm
```

13 Reboot each node in the cluster.

How to Remove Support for Oracle RAC From Selected Nodes

Perform this task to remove Support for Oracle RAC from selected nodes.

On a cluster where multiple Oracle RAC databases are running, perform this task to remove an Oracle RAC database from selected nodes. The Oracle RAC database that you remove continues to run on the other cluster nodes. The remaining Oracle RAC databases continue to run on the selected nodes.

This task involves removing the selected nodes from the following resource groups in the following order:

- The resource group for the Oracle RAC database
- Resource groups for logical hostname resources
- Resource groups for scalable file-system mount-point resources
- Resource groups for scalable device group resources
- The Oracle RAC framework resource group

Caution – You might perform this task to remove an Oracle RAC database from selected nodes of a cluster where multiple Oracle RAC databases are running. In this situation, do *not* remove the nodes from any resource group on whose resources the remaining Oracle RAC databases depend. For example, you might have configured multiple database file systems to depend on a single device group. In this situation, do *not* remove the nodes from the resource group that contains the resource for the scalable device group. Similarly, if multiple databases depend on the Oracle RAC framework resource group, do *not* remove the nodes from this resource group.

1 **Become superuser.**

2 **Remove the nodes from the resource group for each Oracle RAC database that you are removing.**

For each Oracle RAC database that you are removing, perform the following steps:

a. **Take offline the resource group for the Oracle RAC database on the nodes from which you are removing Support for Oracle RAC.**

```
# clresourcegroup offline -n nodelist rac-db-rg
```

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes from which you are taking offline the resource group.

rac-db-rg
Specifies the name of the resource group that you are taking offline.

b. Remove the nodes from the node list of the resource group for the Oracle RAC database.

```
# clresourcegroup remove-node -n nodelist rac-db-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

rac-db-rg
Specifies the name of the resource group from which you are removing nodes.

3 Remove each node that you are removing from the list of nodes where the Oracle Grid Infrastructure resource for the Oracle database runs.

```
# Grid_home/bin/crsctl delete -n node-name
```

Grid_home
Specifies the Oracle Grid Infrastructure home directory. This directory contains the Oracle Grid Infrastructure binary files and Oracle Grid Infrastructure configuration files.

node-name
Specifies the host name of the node where the Oracle Grid Infrastructure resource runs.

4 From each node that you are removing, remove each Oracle Grid Infrastructure resource that represents an Oracle Solaris Cluster resource from whose resource group you are removing nodes.

An Oracle Grid Infrastructure resource is configured for each Oracle Solaris Cluster resource for scalable device groups and scalable file-system mount points on which Oracle components depend.

For each Oracle Grid Infrastructure resource that you are removing, perform the following steps on each node from which you are removing the resource:

a. Stop the Oracle Grid Infrastructure resource that you are removing.

```
# Grid_home/bin/crsctl stop
```

b. Delete the Oracle Grid Infrastructure resource that you want to remove.

```
# Grid_home/bin/crsctl delete res
```

5 Use Oracle utilities to remove the following items from each node from which you are removing Support for Oracle RAC:

- The Oracle RAC database
- Oracle Grid Infrastructure

6 Remove the nodes from any resource groups that contain scalable file-system mount-point resources.

If no resource groups that contain scalable file-system mount-point resources are configured, omit this step.

For each resource group from which you are removing nodes, perform the following steps:

a. Take offline the resource group on the nodes from which you are removing Support for Oracle RAC.

```
# clresourcegroup offline -n nodelist scal-mp-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes from which you are taking offline the resource group.

scal-mp-rg
Specifies the name of the resource group that you are taking offline.

b. Remove the nodes from the node list of the resource group.

```
# clresourcegroup remove-node -n nodelist scal-mp-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

scal-mp-rg
Specifies the name of the resource group from which you are removing nodes.

7 Remove the nodes from any resource groups that contain scalable device group resources.

If no resource groups that contain scalable device group resources are configured, omit this step.

For each resource group from which you are removing nodes, perform the following steps:

a. Take offline the resource group on the nodes from which you are removing Support for Oracle RAC.

```
# clresourcegroup offline -n nodelist scal-dg-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes from which you are taking offline the resource group.

scal-dg-rg
Specifies the name of the resource group that you are taking offline.

b. Remove the nodes from the node list of the resource group.

```
# clresourcegroup remove-node -n nodelist scal-dg-rg
```

-n *nodelist*
Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

scal-dg-rg
Specifies the name of the resource group from which you are removing nodes.

8 Remove the nodes from any scalable device groups that were affected by the removal of nodes from resource groups in Step 7.

How to perform this step depends on the type of the scalable device group.

- **For each Solaris Volume Manager for Sun Cluster multi-owner disk set, type the following command:**

  `# metaset -s` *scal-dg-ms* `-d -h` *nodelist*

  `-s` *scal-dg-ms*
  Specifies the name of the disk set from which you are removing nodes.

  `-h` *nodelist*
  Specifies a space-separated list of the nodes that you are removing from the disk set.

Note – If you are removing an Oracle RAC database from selected nodes of a cluster where multiple Oracle RAC databases are running, omit the remaining steps in this procedure.

---

9 Remove the nodes from the Oracle RAC framework resource group.

a. Take offline the resource group on the nodes from which you are removing Support for Oracle RAC.

`# clresourcegroup offline -n` *nodelist rac-fmwk-rg*

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes from which you are taking offline the resource group.

rac-fmwk-rg
Specifies the name of the resource group that you are taking offline.

b. Remove the nodes from the node list of the resource group.

`# clresourcegroup remove-node -n` *nodelist rac-fmwk-rg*

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

rac-fmwk-rg
Specifies the name of the resource group from which you are removing nodes.

10 Remove the nodes from the multiple-owner volume-manager framework resource group, if used.

a. Take offline the resource group on the nodes from which you are removing Support for Oracle RAC.

`# clresourcegroup offline -n` *nodelist vucmm-fmwk-rg*

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes from which you are taking offline the resource group.

vucmm-fmwk-rg
Specifies the name of the resource group that you are taking offline.

b. Remove the nodes from the node list of the resource group.

```
# clresourcegroup remove-node -n nodelist vucmm-fmwk-rg
```

`-n` *nodelist*
Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

vucmm-fmwk-rg
Specifies the name of the resource group from which you are removing nodes.

11 (Optional) From each node that you removed, uninstall the Support for Oracle RAC software packages.

```
# pkg uninstall ha-cluster/data-service/oracle-database ha-cluster/library/ucmm
```

12 Reboot each node from which you from which you removed Support for Oracle RAC.

◆ ◆ ◆ APPENDIX A

Sample Configurations of This Data Service

The sections that follow illustrate the configuration of resources and resource groups for typical combinations of storage management schemes on the SPARC platform for both the global cluster and a zone cluster.

Sample Oracle RAC Configurations in the Global Cluster

FIGURE A–1 Configuration of Oracle RAC With Solaris Volume Manager for Sun Cluster

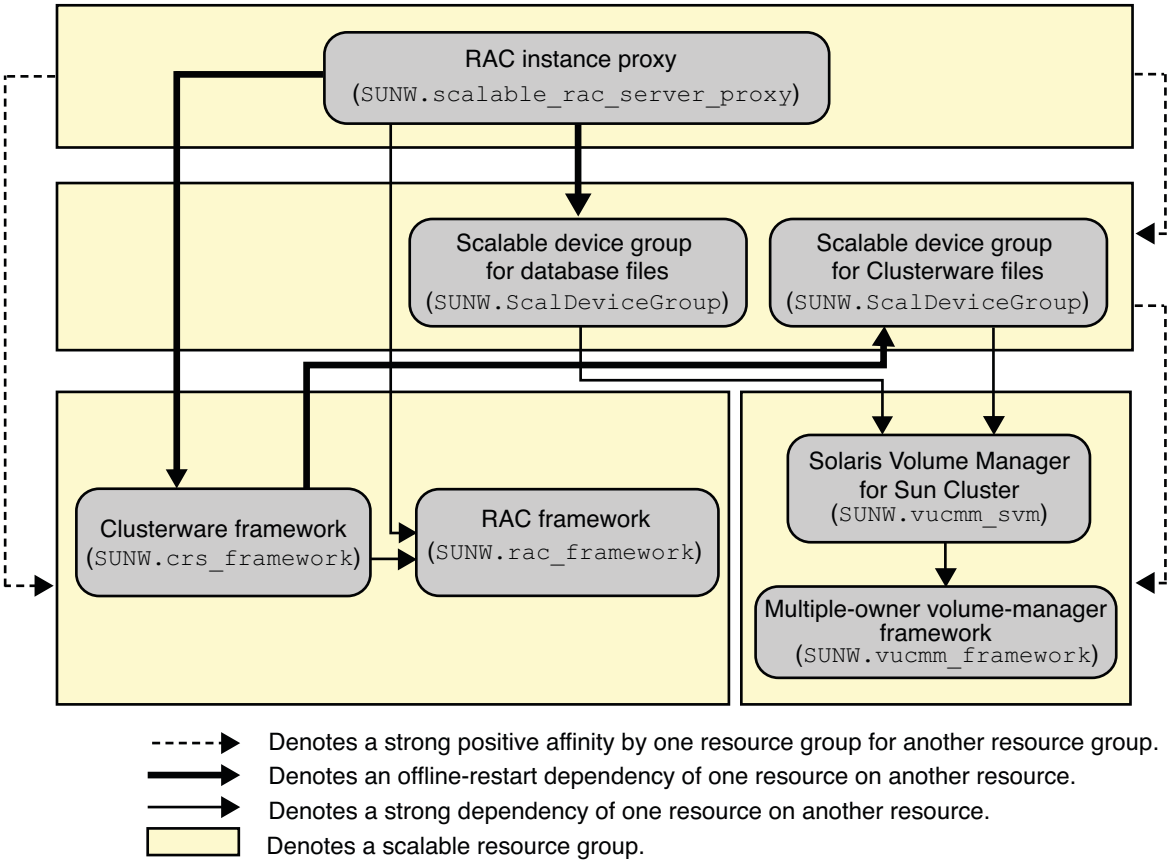

FIGURE A–2 Configuration of Oracle RAC With a NAS Device

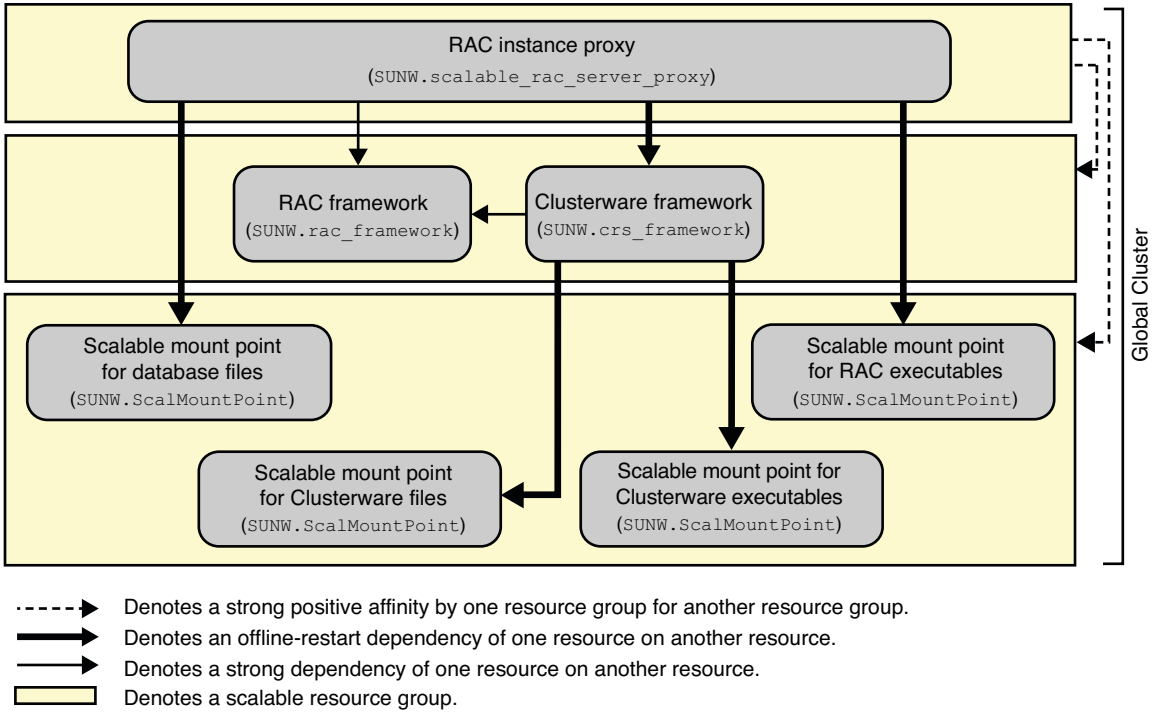

FIGURE A–3 Configuration of Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster

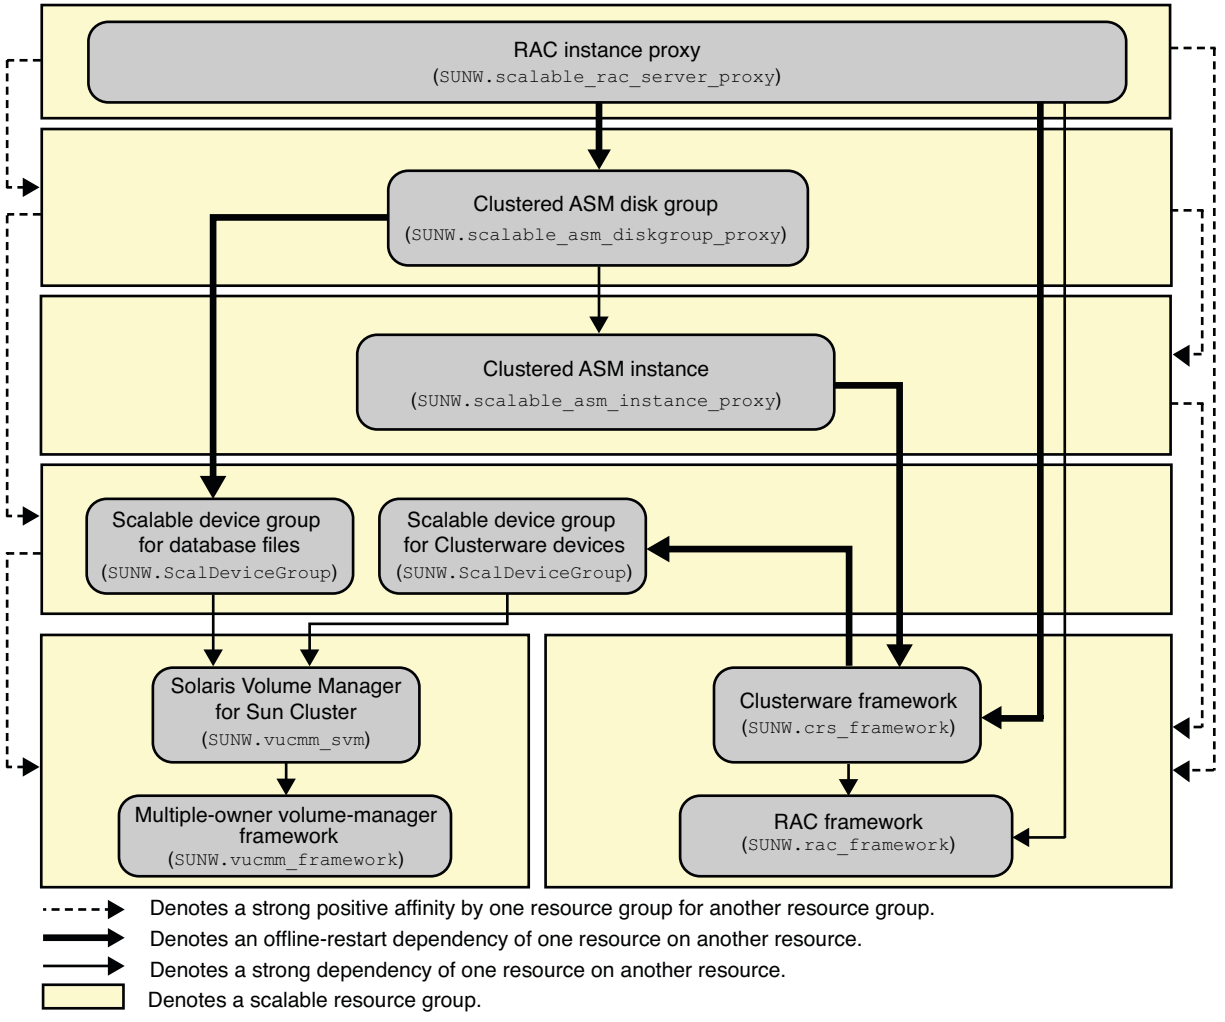

FIGURE A–4 Configuration of Oracle RAC With Oracle ASM and Hardware RAID

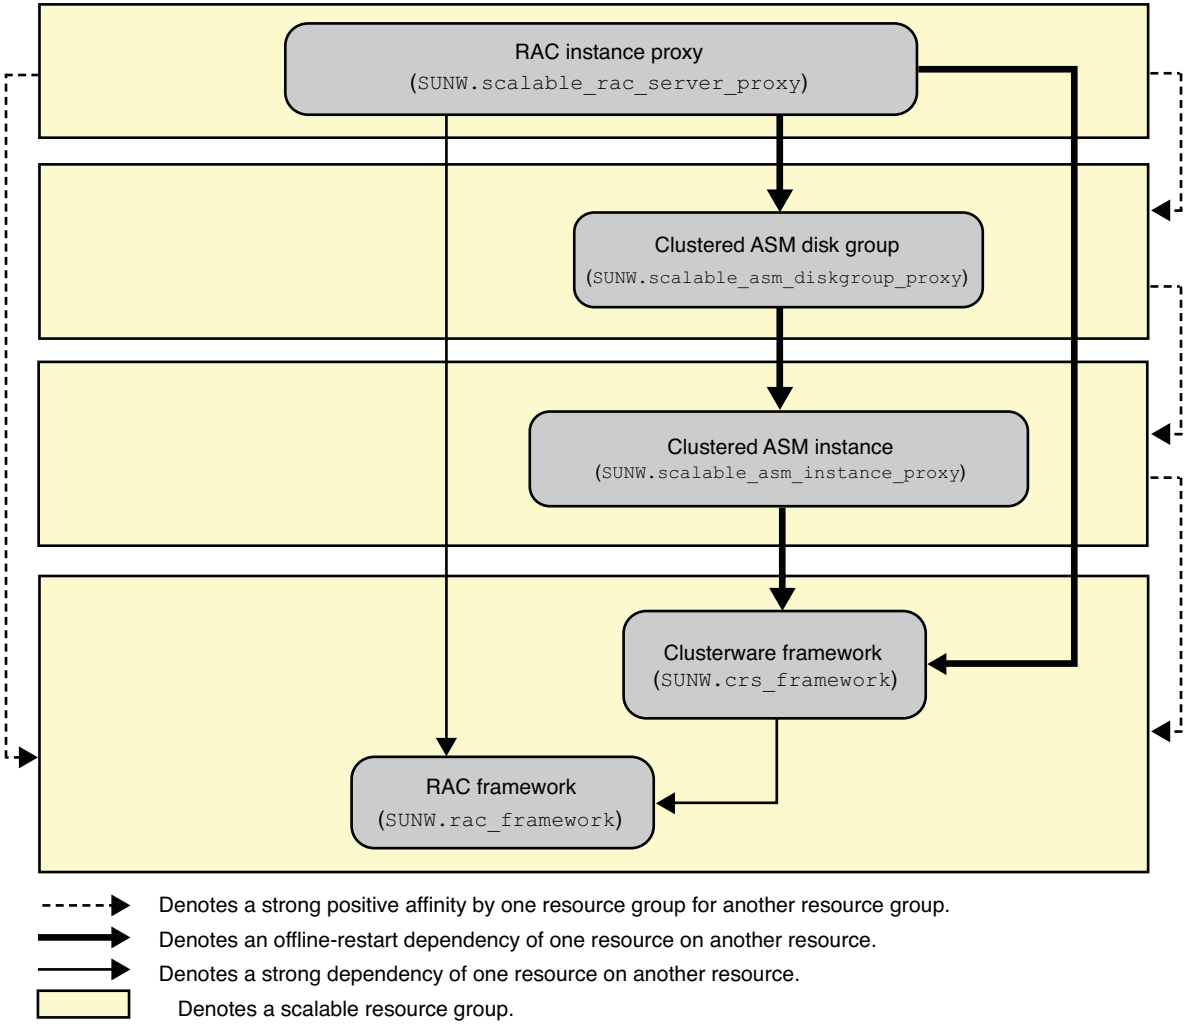

Sample Oracle RAC Configurations in a Zone Cluster

FIGURE A–5 Configuration of Oracle RAC With Solaris Volume Manager for Sun Cluster in a Zone Cluster

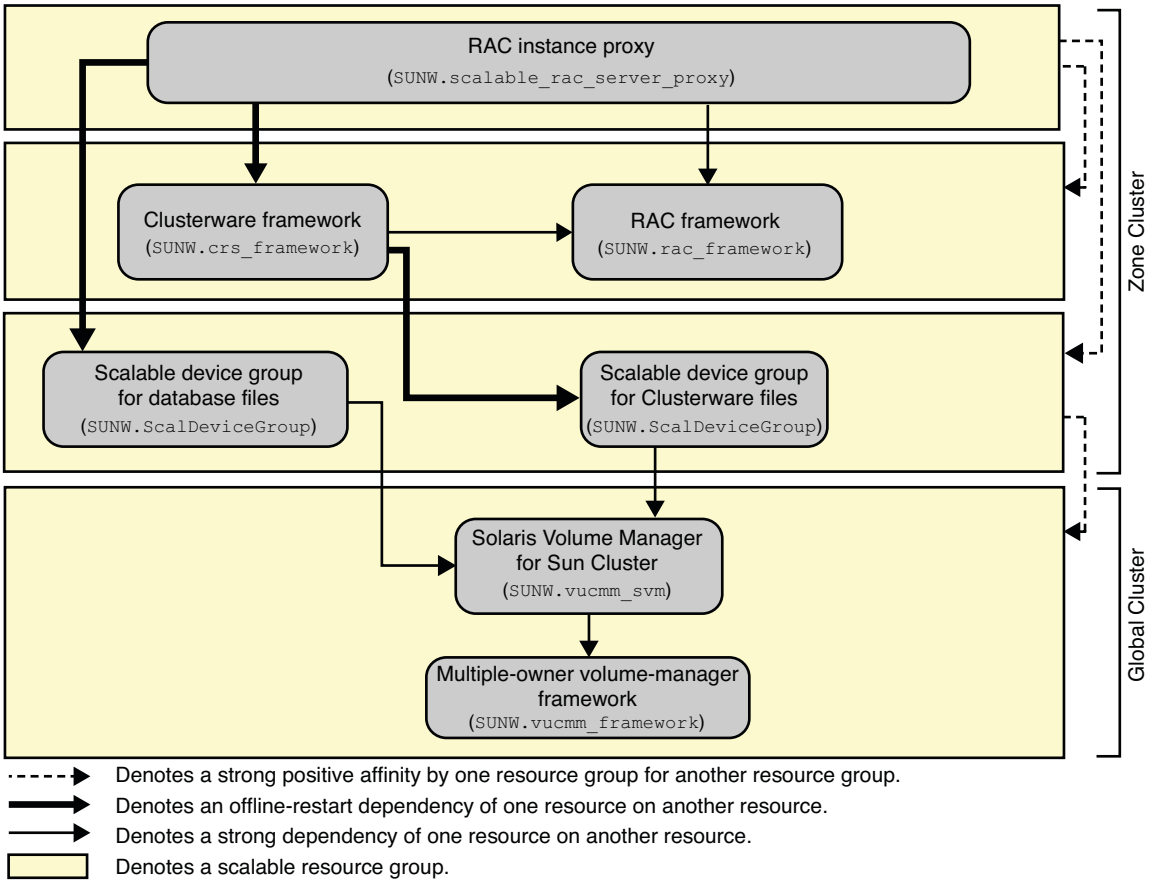

FIGURE A–6 Configuration of Oracle RAC With a NAS Device in a Zone Cluster

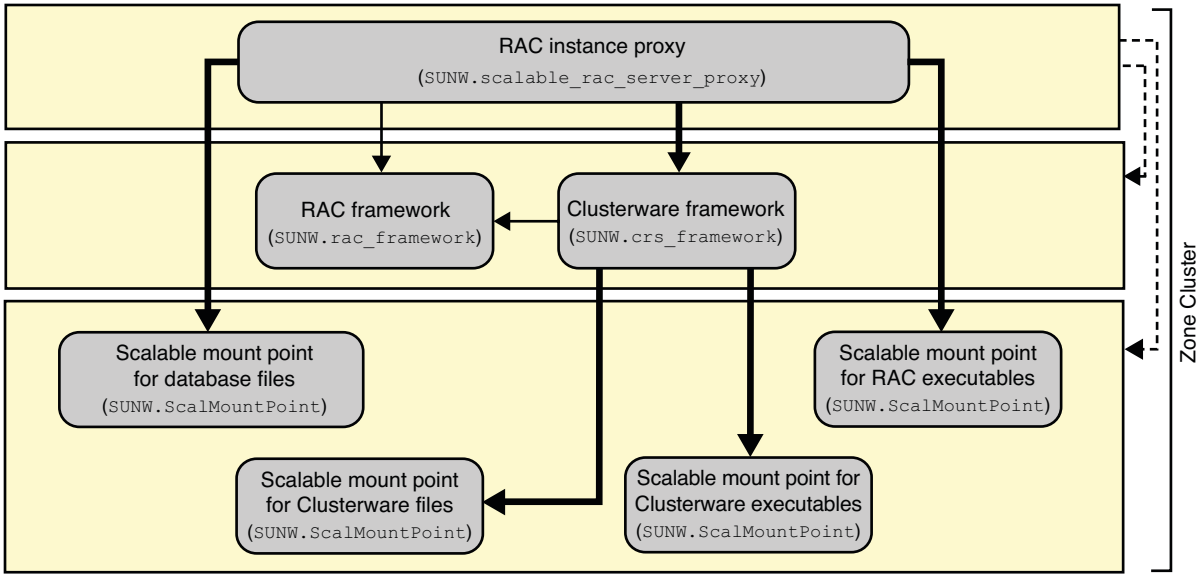

FIGURE A–7 Configuration of Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster in a Zone Cluster

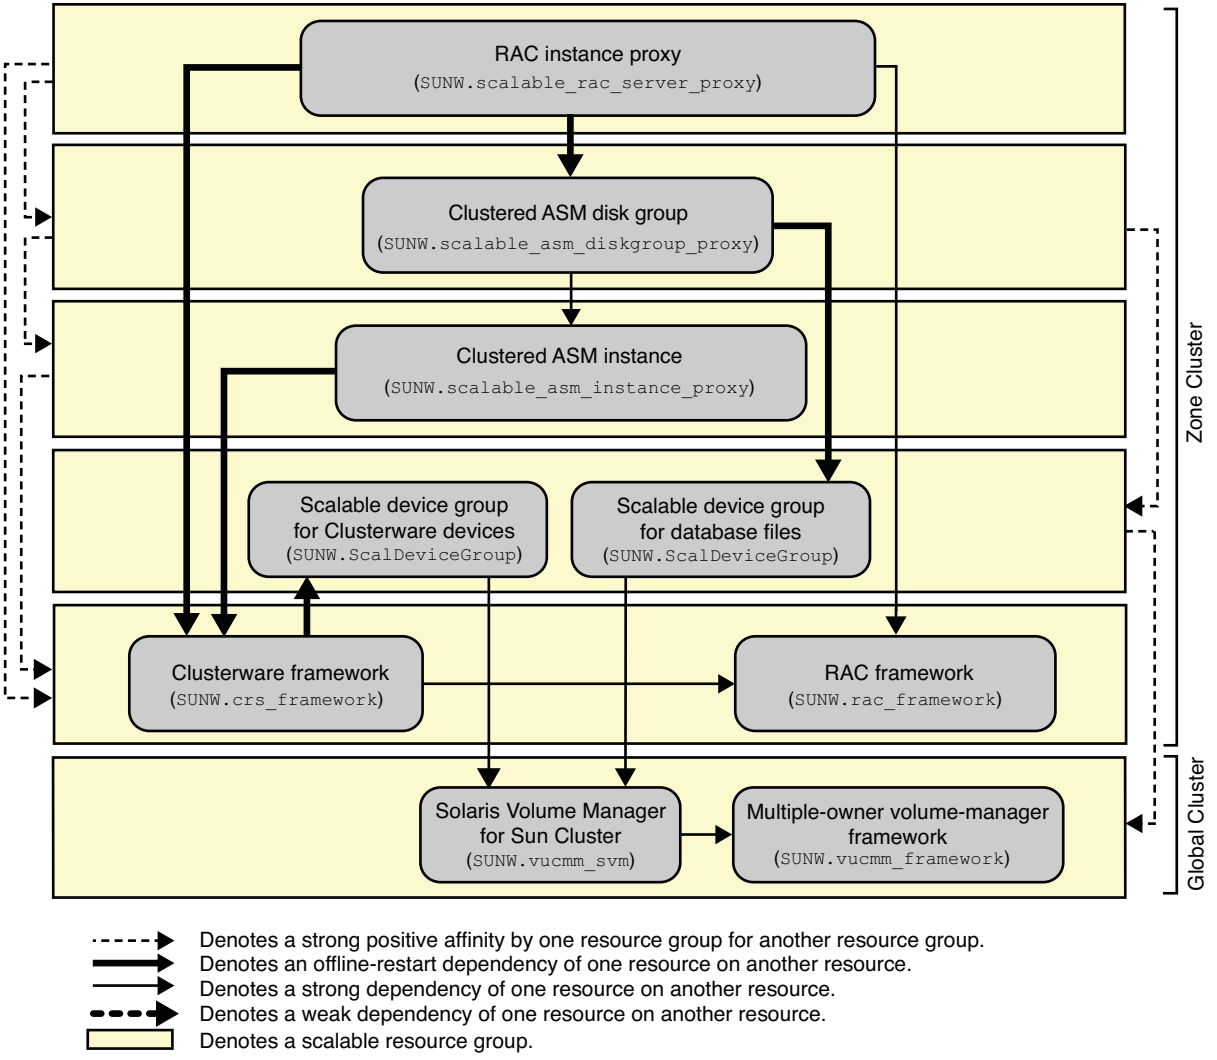

FIGURE A–8 Configuration of Oracle RAC With Oracle ASM and Hardware RAID in a Zone Cluster

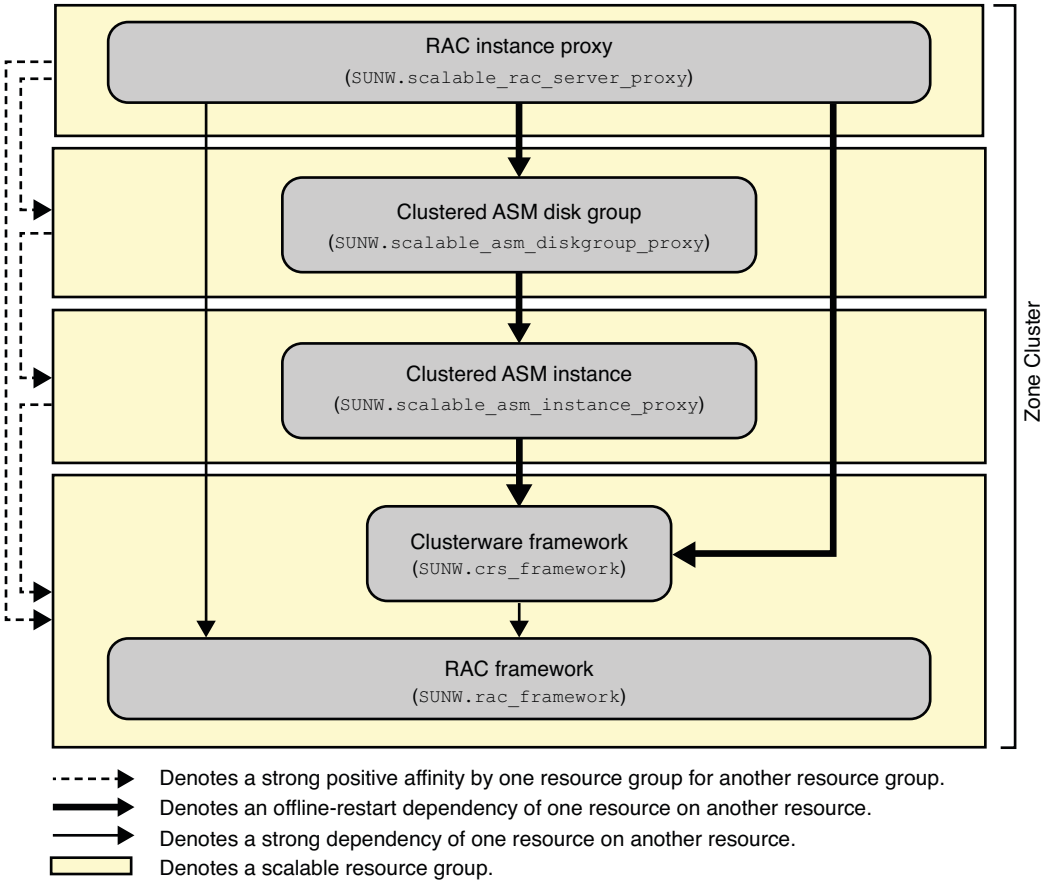

◆ ◆ ◆ APPENDIX B

Preset Actions for DBMS Errors and Logged Alerts

Preset actions for database management system (DBMS) errors and logged alerts are listed as follows:

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

TABLE B–1 Preset Actions for DBMS Errors

Error Number	Action	Connection State	New State	Message
18	NONE	co	di	Max. number of DBMS sessions exceeded
20	NONE	co	di	Max. number of DBMS processes exceeded
28	NONE	on	di	Session killed by DBA, will reconnect
50	RESTART	*	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	RESTART	*	di	Operating system archival error occurred. Check alert log.
447	RESTART	*	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	RESTART	*	di	Oracle background process died

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

Error Number	Action	Connection State	New State	Message
471	RESTART	*	di	Oracle background process died
472	RESTART	*	di	Oracle background process died
473	RESTART	*	di	Oracle background process died
474	RESTART	*	di	SMON died, warm start required
475	RESTART	*	di	Oracle background process died
476	RESTART	*	di	Oracle background process died
477	RESTART	*	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	RESTART	*	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

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

Error Number	Action	Connection State	New State	Message
1034	NONE	on	di	Oracle is not available
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	RESTART	*	*	user storage address cannot be read or written
1054	RESTART	*	*	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	RESTART	*	*	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

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

Error Number	Action	Connection State	New State	Message
2811	RESTART	*	*	Unable to attach shared memory segment
2839	RESTART	*	*	Sync of blocks to disk failed.
2840	RESTART	*	*	
2846	RESTART	*	*	
2847	RESTART	*	*	
2849	RESTART	*	*	
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	RESTART	*	di	
7206	RESTART	*	di	
7208	RESTART	*	di	
7210	RESTART	*	di	
7211	RESTART	*	di	
7212	RESTART	*	di	
7213	RESTART	*	di	
7214	RESTART	*	di	
7215	RESTART	*	di	
7216	RESTART	*	di	
7218	RESTART	*	di	

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

Error Number	Action	Connection State	New State	Message
7219	RESTART	*	*	`slspool: unable to allocate spooler argument buffer.`
7223	RESTART	*	*	`slspool: fork error, unable to spawn spool process. - Resource limit reached`
7224	RESTART	*	*	
7229	RESTART	*	*	
7232	RESTART	*	*	
7234	RESTART	*	*	
7238	RESTART	*	*	`slemcl: close error.`
7250	RESTART	*	*	
7251	RESTART	*	*	
7252	RESTART	*	*	
7253	RESTART	*	*	
7258	RESTART	*	*	
7259	RESTART	*	*	
7263	RESTART	*	*	
7269	RESTART	*	*	
7279	RESTART	*	*	
7280	RESTART	*	*	
7296	RESTART	*	*	
7297	RESTART	*	*	
7306	RESTART	*	*	
7310	RESTART	*	*	
7315	RESTART	*	*	
7321	RESTART	*	*	
7322	RESTART	*	*	
7324	RESTART	*	*	
7325	RESTART	*	*	
7351	RESTART	*	*	

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

Error Number	Action	Connection State	New State	Message
7361	RESTART	*	*	
7404	RESTART	*	*	
7414	RESTART	*	*	
7415	RESTART	*	*	
7417	RESTART	*	*	
7418	RESTART	*	*	
7419	RESTART	*	*	
7430	RESTART	*	*	
7455	RESTART	*	*	
7456	RESTART	*	*	
7466	RESTART	*	*	
7470	RESTART	*	*	
7475	RESTART	*	*	
7476	RESTART	*	*	
7477	RESTART	*	*	
7478	RESTART	*	*	
7479	RESTART	*	*	
7481	RESTART	*	*	
9706	RESTART	*	*	
9716	RESTART	*	*	
9718	RESTART	*	*	
9740	RESTART	*	*	
9748	RESTART	*	*	
9747	RESTART	*	*	
9749	RESTART	*	*	
9751	RESTART	*	*	
9755	RESTART	*	*	

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

Error Number	Action	Connection State	New State	Message
9757	RESTART	*	*	
9756	RESTART	*	*	
9758	RESTART	*	*	
9761	RESTART	*	*	
9765	RESTART	*	*	
9779	RESTART	*	*	
9829	RESTART	*	*	
9831	RESTART	*	*	
9834	RESTART	*	*	
9836	RESTART	*	*	
9838	RESTART	*	*	
9837	RESTART	*	*	
9844	RESTART	*	*	
9845	RESTART	*	*	
9846	RESTART	*	*	
9847	RESTART	*	*	
9853	RESTART	*	*	
9854	RESTART	*	*	
9856	RESTART	*	*	
9874	RESTART	*	*	
9876	RESTART	*	*	
9877	RESTART	*	*	
9878	RESTART	*	*	
9879	RESTART	*	*	
9885	RESTART	*	*	
9888	RESTART	*	*	
9894	RESTART	*	*	

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

Error Number	Action	Connection State	New State	Message
9909	RESTART	*	*	
9912	RESTART	*	*	
9913	RESTART	*	*	
9919	RESTART	*	*	
9943	RESTART	*	*	
9947	RESTART	*	*	
9948	RESTART	*	*	
9949	RESTART	*	*	
9950	RESTART	*	*	
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
99999	RESTART	*	di	Monitor detected death of Oracle background processes.

TABLE B–2 Preset Actions for Logged Alerts

Alert String	Action	Connection State	New State	Message
ORA-07265	RESTART	*	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)

◆ ◆ ◆ APPENDIX C

Support for Oracle RAC Extension Properties

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

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.

For information about all system-defined properties, see the r_properties(5) man page and the rg_properties(5) man page.

SUNW.crs_framework Extension Properties

The SUNW.crs_framework resource type has no extension properties.

SUNW.rac_framework Extension Properties

`reservation_timeout`
This property specifies the timeout (in seconds) for the reservation step of a reconfiguration of Support for Oracle RAC.

Data Type: Integer

Default: 325

Range: 100 – 99999 seconds

Tunable: Any time

SUNW.scalable_asm_diskgroup_proxy Extension Properties

`asm_diskgroups`
This property specifies the Oracle ASM disk group. If required, more than one Oracle ASM disk group can be specified as a comma separated list.

Data Type: String array

Default: Not applicable

Range: Not applicable

Tunable: When disabled

`debug_level` (integer)

Note – All SQL*Plus and `srvmgr` messages that the Oracle ASM disk group resource issues are written to the log file `/var/opt/SUNWscor/oracle_asm/message_log.${RESOURCE}`.

This property indicates the level to which debug messages for the Oracle ASM disk group resources are logged. When the debug level is increased, more debug messages are written to the system log `/var/adm/messages` as follows:

0	No debug messages
1	Function Begin and End messages
2	All debug messages and function Begin/End messages

You can specify a different value of the `debug_level` extension property for each node that can master the resource.

Range: 0–2

Default: 0

Tunable: Any time

`Proxy_probe_interval`(integer)
Specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the clustered Oracle ASM disk group resource for which this resource is acting as a proxy.

Range: 5–120

Default: 30

Tunable: Any time

`proxy_probe_timeout`(integer)
This property specifies the timeout value, in seconds, for the probe command.

Range: 5–120

Default: 60

Tunable: Any time

SUNW.scalable_asm_instance_proxy Extension Properties

`asm_diskgroups`
This property specifies the Oracle ASM disk group. If required, more than one Oracle ASM disk group can be specified as a comma separated list.

Data Type: String array

Default: Not applicable

Range: Not applicable

Tunable: When disabled

`crs_home`
This property specifies the full path to the Oracle Grid Infrastructure home directory. The Oracle Grid Infrastructure home directory contains the binary files, log files and parameter files for the Oracle Grid Infrastructure software.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: When disabled

`debug_level`

Note – All SQL*Plus and `srvmgr` messages that the clustered Oracle ASM instance proxy resource issues are written to the log file `/var/opt/SUNWscor/oracle_asm/message_log.${RESOURCE}`.

This property indicates the level to which debug messages from the monitor for the clustered Oracle ASM instance proxy are logged. When the debug level is increased, more debug messages are written to the system log `/var/adm/messages` as follows:

0
No debug messages

1
Function Begin and End messages

2
All debug messages and function Begin/End messages

You can specify a different value of the `debug_level` extension property for each node that can master the resource.

Data Type: Integer

Range: 0–2

Default: 0

Tunable: Any time

`oracle_home`
This property specifies the full path to the Oracle home directory. The Oracle home directory contains the binary files, log files, and parameter files for the Oracle software.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: When disabled

`oracle_sid`
This property specifies the Oracle System Identifier (SID). The Oracle SID uniquely identifies the clustered Oracle ASM database instance on the node where the instance is running.

You *must* specify a different value of the `oracle_sid` extension property for each node that can master the resource. The value for each node must correctly identify the instance that is running on the node.

Data Type: String

Range: Not applicable

Default: `NULL`

Tunable: When disabled

`proxy_probe_timeout`
This property specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the Oracle Grid Infrastructure resource for which this resource is acting as a proxy.

Data Type: Integer

Range: 5–120

Default: 60

Tunable: Any time

`proxy_probe_interval`
This property specifies the interval , in seconds, between probes of the Oracle Grid Infrastructure resource for which this resource is acting as a proxy.

Data Type: Integer

Range: 5–120

Default: 60

Tunable: Any time

SUNW.scalable_rac_server_proxy Extension Properties

`client_retries`
This property specifies the maximum number of attempts by the resource's remote procedure call (RPC) client to connect to the proxy daemon.

Data Type: Integer

Range: 1–25

Default: 3

Tunable: When disabled

`client_retry_interval`
This property specifies the interval, in seconds, between attempts by the resource's RPC client to connect to the proxy daemon.

Data Type: Integer

Range: 1–3600

Default: 5

Tunable: When disabled

`crs_home`
This property specifies the directory in which the Oracle Grid Infrastructure software is located.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: When disabled

`db_name`
This property specifies the name that uniquely identifies the specific Oracle RAC database that is associated with this resource. This identifier distinguishes the database from other databases that might run simultaneously on your system. The name of the Oracle RAC database is specified during the installation of Oracle RAC.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: When disabled

`debug_level`
This property indicates the level to which debug messages from the component for the Oracle RAC proxy server 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/scalable_rac_server_proxy/message_log.`*rs*, where *rs* is the name of the resource that represents the Oracle RAC proxy server component.

You can specify a different value of the `debug_level` extension property for each node that can master the resource.

Data Type: Integer

Range: 0–100

Default: 1, which logs `syslog` messages

Tunable: Any time

`monitor_probe_interval`
This property specifies the interval, in seconds, between probes of the Oracle Grid Infrastructure resource for which this resource is acting as a proxy.

Data Type: Integer

Range: 10–3600

Default: 300

Tunable: Any time

`oracle_home`
This property specifies the full path to the Oracle home directory. The Oracle home directory contains the binary files, log files, and parameter files for the Oracle software.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: When disabled

`oracle_sid`
This property specifies the Oracle System Identifier (SID). The Oracle SID uniquely identifies the Oracle RAC database instance on the node where the instance is running.

You *must* specify a different value of the `oracle_sid` extension property for each node that can master the resource. The value for each node must correctly identify the instance that is running on the node.

Data Type: String

Range: Not applicable

Default: `NULL`

Tunable: When disabled

`proxy_probe_timeout`
This property specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the Oracle Grid Infrastructure resource for which this resource is acting as a proxy.

Data Type: Integer

Range: 5–3600

Default: 120

Tunable: Any time

`startup_wait_count`
This property specifies the maximum number of attempts by this resource to confirm that the Oracle Grid Infrastructure software is started completely. The interval between attempts is twice the value of the `proxy_probe_timeout` extension property.

The resource requires confirmation that Oracle Grid Infrastructure software is started before attempting to start an Oracle RAC database instance. If the maximum number of attempts is exceeded, the resource does not attempt to start the database instance.

Data Type: Integer

Range: 10–600

Default: 20

Tunable: When disabled

`user_env`
This property specifies the name of the file that contains the environment variables that are to be set before the database starts up or shuts down. You must define all environment variables whose values differ from Oracle defaults in this file.

For example, a user's `listener.ora` file might not be located under the `/var/opt/oracle` directory or the *oracle-home*`/network/admin` directory. In this situation, the `TNS_ADMIN` environment variable must be defined.

The definition of each environment variable that is defined must follow the format *variable-name*=*value*. Each definition must start on a new line in the environment file.

You can specify a different value of the `user_env` extension property for each node that can master the resource.

Data Type: String

Range: Not applicable

Default: No default defined

Tunable: Any time

SUNW.ScalDeviceGroup Extension Properties

`debug_level`
This property specifies the level to which debug messages from the resource of this type are logged. When the debug level is increased, more debug messages are written to the log files.

Data type: Integer

Default: 0

Range: 0–10

Tunable: Any time

`diskgroupname`
This property specifies the name of the device group that the resource represents. You must set this property to the following item:

- The name of an existing Solaris Volume Manager for Sun Cluster multi-owner disk set. This name was specified in the `metaset(1M)` command with which the disk set was created.

The requirements for the device group that you specify are as follows:

- The device group must be a valid, existing multi-owner disk set or shared-disk group.
- The device group must be hosted on all nodes that can master the resource.
- The device group must be accessible from all nodes that can master the scalable device group resource.
- The device group must contain at minimum one volume.

Data type: String

Default: No default defined

Range: Not applicable

Tunable: When disabled

`logicaldevicelist`
This property specifies a comma-separated list of logical volumes that the fault monitor of the resource is to monitor. This property is optional. If you do not specify a value for this property, all logical volumes in the device group are to be monitored.

The status of the device group is derived from the statuses of the individual logical volumes that are monitored. If all monitored logical volumes are healthy, the device group is healthy. If any monitored logical volume is faulty, the device group is faulty.

The status of an individual logical volume is obtained by querying the volume's volume manager. If the status of a Solaris Volume Manager for Sun Cluster volume cannot be determined from a query, the fault monitor performs file input/output (I/O) operations to determine the status.

If a device group is discovered to be faulty, monitoring of the resource that represents the group is stopped and the resource is put into the disabled state.

Note – For mirrored disks, if one submirror is faulty, the device group is still considered to be healthy.

The requirements for each logical volume that you specify are as follows:

- The logical volume must exist.
- The logical volume must be contained in the device group that the `diskgroupname` property specifies.
- The logical volume must be accessible from all nodes that can master the scalable device group resource.

Data type: String array

Default: ""

Range: Not applicable

Tunable: Any time

`monitor_retry_count`
This property specifies the maximum number of restarts by the process monitor facility (PMF) that are allowed for the fault monitor.

Data type: Integer

Default: 4

Range: No range defined

Tunable: Any time

`monitor_retry_interval`
This property specifies the period of time in minutes during which the PMF counts restarts of the fault monitor.

Data type: Integer

Default: 2

Range: No range defined

Tunable: Any time

SUNW.ScalMountPoint Extension Properties

`debug_level`
This property specifies the level to which debug messages from the resource for a file-system mount point are logged. When the debug level is increased, more debug messages are written to the log files.

Data type: Integer

Default: 0

Range: 0–10

Tunable: Any time

`filesystemtype`
This property specifies the type of file system whose mount point the resource represents. You must specify this property. Set this property to one of the following values:

`nas`
Specifies that the file system is a file system on a qualified NAS device.

Data type: String

Default: No default defined

Range: Not applicable

Tunable: When disabled

`iotimeout`
This property specifies the timeout value in seconds that the fault monitor uses for file input/output (I/O) probes. To determine if the mounted file system is available, the fault monitor performs I/O operations such as opening, reading, and writing to a test file on the file system. If an I/O operation is not completed within the timeout period, the fault monitor reports an error.

Data type: Integer

Default: 300

Range: 5–300

Tunable: Any time

`monitor_retry_count`
This property specifies the maximum number of restarts by the process monitor facility (PMF) that are allowed for the fault monitor.

Data type: Integer

Default: 4

Range: No range defined

Tunable: Any time

`monitor_retry_interval`
This property specifies the period of time in minutes during which the PMF counts restarts of the fault monitor.

Data type: Integer

Default: 2

Range: No range defined

Tunable: Any time

`mountoptions`
This property specifies a comma-separated list of mount options that are to be used when the file system that the resource represents is mounted. This property is optional. If you do not specify a value for this property, mount options are obtained from the file system's table of defaults.

- For a file system on a qualified NAS device, these options are obtained from the `/etc/vfstab` file.

Mount options that you specify through this property override the mount options in the file system's table of defaults.

Data type: String

Default: ""

Range: Not applicable

Tunable: When disabled

`mountpointdir`
This property specifies the mount point of the file system that the resource represents. The mount point is the full path to the directory where the file system is attached to the file system hierarchy when the file system is mounted. You must specify this property.

The directory that you specify must already exist.

Data type: String

Default: No default defined

Range: Not applicable

Tunable: When disabled

`targetfilesystem`
This property specifies the file system that is to be mounted at the mount point that the `mountpointdir` extension property specifies. You must specify this property. The type of the file system must match the type that the `filesystemtype` property specifies. The format of this property depends on the type of the file system as follows:

- For a file system on a qualified NAS device, set this property to *nas-device*:*path*. The replaceable items in this format are as follows:

  nas-device
  Specifies the name of the qualified NAS device that is exporting the file system. You can optionally qualify this name with a domain.

  path
  Specifies the full path to the file system that the qualified NAS device is exporting.

  The qualified NAS device and the file system must already be configured for use with Sun Cluster. For more information, see *Oracle Solaris Cluster 4.0 With Network-Attached Storage Device Manual*.

Data type: String

Default: No default defined

Range: Not applicable

Tunable: When disabled

SUNW.vucmm_framework Extension Properties

`reservation_timeout`
This property specifies the timeout (in seconds) for the reservation step of a reconfiguration of the framework.

Data Type: Integer

Default: 325

Range: 100 – 99999 seconds

Tunable: Any time

SUNW.vucmm_svm Extension Properties

`debug_level`
This property specifies the level to which debug messages from the Solaris Volume Manager for Sun Cluster component are logged. When the debug level is increased, more messages are written to the log files during reconfiguration.

Data Type: Integer

Default: 1, which logs `syslog` messages

Range: 0 – 10

Tunable: Any time

`svm_abort_step_timeout`
This property specifies the timeout (in seconds) for the abort step of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_return_step_timeout`
This property specifies the timeout (in seconds) for the return step of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_start_step_timeout`
This property specifies the timeout (in seconds) for the start step of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_step1_timeout`
This property specifies the timeout (in seconds) for step 1 of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_step2_timeout`
This property specifies the timeout (in seconds) for step 2 of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_step3_timeout`
This property specifies the timeout (in seconds) for step 3 of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 30 – 99999 seconds

Tunable: Any time

`svm_step4_timeout`
This property specifies the timeout (in seconds) for step 4 of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 120

Range: 100 – 99999 seconds

Tunable: Any time

`svm_stop_step_timeout`
This property specifies the timeout (in seconds) for the stop step of a reconfiguration of the Solaris Volume Manager for Sun Cluster module of the volume manager reconfiguration framework.

Data Type: Integer

Default: 40

Range: 30 – 99999 seconds

Tunable: Any time

SUNW.wait_zc_boot Extension Properties

`zcname`
This property specifies the name of the zone cluster that needs to be booted before the dependent resource.

Data type: String

Default: No default defined

Range: Not applicable

Tunable: When disabled

◆ ◆ ◆ APPENDIX D

Command-Line Alternatives

Oracle Solaris Cluster maintenance commands enable you to automate the creation, modification, and removal of the framework resource groups by using scripts. Automating this process reduces the time for propagating the same configuration information to many nodes in a cluster.

This appendix contains the following sections:

Setting Support for Oracle RAC Extension Properties

The procedures in the sections that follow contain instructions for registering and configuring resources. These instructions explain how to set *only* extension properties that Support for Oracle RAC requires you to set. Optionally, you can set additional extension properties to override their default values. For more information, see the following sections:

- "Tuning Support for Oracle RAC" on page 105
- Appendix C, "Support for Oracle RAC Extension Properties"

Registering and Configuring the Framework Resource Groups by Using Oracle Solaris Cluster Maintenance Commands

The tasks in this section are an alternative for the resource-configuration steps in "How to Register and Configure the Oracle RAC Framework Resource Group by Using `clsetup`" on page 52. Instructions include configuring the multiple-owner volume-manager framework, which currently you cannot configure by using `clsetup`. The following information is in this section:

- "Overview of the Framework Resource Groups" on page 172
- "How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands" on page 173

Overview of the Framework Resource Groups

This section describes the following framework resource groups:

- "Oracle RAC Framework Resource Group" on page 172
- "Multiple-Owner Volume-Manager Framework Resource Group" on page 172

Oracle RAC Framework Resource Group

The Oracle RAC framework resource group enables Oracle RAC to run with Oracle Solaris Cluster. This resource group contains an instance of the following single-instance resource types:

- `SUNW.rac_framework`, which represents the framework that enables Oracle RAC to be managed by using Oracle Solaris Cluster commands

---

Note – The resource types that are defined for the Oracle RAC framework resource group do *not* enable the Resource Group Manager (RGM) to manage instances of Oracle RAC.

Multiple-Owner Volume-Manager Framework Resource Group

The multiple-owner volume-manager framework resource group enables a multiple-owner shared-storage feature to be used by Oracle RAC.

The multiple-owner volume-manager framework resource group is based on the `SUNW.vucmm_framework` resource type. This resource group contains the volume manager resources for the multiple-owner volume-manager framework, `SUNW.vucmm_svm`.

▼ How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands

Perform this procedure on only one node of the global cluster.

1 Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.

2 Create a scalable Oracle RAC resource group.

Note – If you are performing the steps in this procedure to register and configure the Oracle RAC resource framework in a zone cluster and Oracle RAC support is not also required in the global cluster, you do not need to also create an Oracle RAC framework resource group in the global cluster. In that case, skip this step and proceed to Step 5.

Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

```
# clresourcegroup create -n nodelist \
-p maximum_primaries=num-in-list \
-p desired_primaries=num-in-list \
[-p rg_description="description"] \
-p rg_mode=Scalable \
rac-fmwk-rg
```

`-n nodelist=`*nodelist*
: Specifies a comma-separated list of cluster nodes on which Support for Oracle RAC is to be enabled. The Support for Oracle RAC software packages must be installed on each node in this list.

`-p maximum_primaries=`*num-in-list*
: Specifies the number of nodes on which Support for Oracle RAC is to be enabled. This number must equal the number of nodes in *nodelist*.

`-p desired_primaries=`*num-in-list*
: Specifies the number of nodes on which Support for Oracle RAC is to be enabled. This number must equal the number of nodes in *nodelist*.

`-p rg_description="`*description*"
: Specifies an optional brief description of the resource group. This description is displayed when you use Oracle Solaris Cluster maintenance commands to obtain information about the resource group.

`-p rg_mode=Scalable`
Specifies that the resource group is scalable.

rac-fmwk-rg
Specifies the name that you are assigning to the Oracle RAC resource group.

3 Register the `SUNW.rac_framework` resource type.

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

4 Add an instance of the `SUNW.rac_framework` resource type to the resource group that you created in Step 2.

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

`-g` *rac-fmwk-rg*
Specifies the resource group to which you are adding the resource. This resource group must be the resource group that you created in Step 2.

rac-fmwk-rs
Specifies the name that you are assigning to the `SUNW.rac_framework` resource.

5 Create a scalable multiple-owner volume-manager framework resource group, if one does not already exist.

If you do not intend to create a multiple-owner volume-manager framework resource group, skip to Step 8.

```
# clresourcegroup create -n nodelist -S vucmm-fmwk-rg
```

`-n nodelist=`*nodelist*
Specifies the same node list that you configured for the scalable Oracle RAC resource group.

vucmm-fmwk-rg
Specifies the name that you are assigning to the multiple-owner volume-manager framework resource group.

6 Register the `SUNW.vucmm_framework` resource type.

If you did not create a multiple-owner volume-manager framework resource group, skip to Step 8.

```
# clresourcetype register SUNW.vucmm_framework
```

7 Add an instance of the `SUNW.vucmm_framework` resource type to the resource group that you created in Step 5.

If you did not create a multiple-owner volume-manager framework resource group, skip to Step 8.

```
# clresource create -g vucmm-fmwk-rg -t SUNW.vucmm_framework vucmm-fmwk-rs
```

`-g` *vucmm-fmwk-rg*
Specifies the resource group to which you are adding the resource. This resource group must be the resource group that you created in Step 5.

vucmm-fmwk-rs
Specifies the name that you are assigning to the `SUNW.vucmm_framework` resource.

8 Register and add an instance of the resource type that represents the volume manager that you are using for Oracle files, if any.

If you are not using a volume manager, omit this step.

- **If you are using Solaris Volume Manager for Sun Cluster, register and add the instance as follows:**

  a. Register the resource type.

  ```
  # clresourcetype register svm-rt
  ```

  svm-rt
  Specifies the `SUNW.vucmm_svm` resource type you register.

  b. Add an instance of the resource type to the resource group to contain the volume manager resource.

  Ensure that this instance depends on the framework resource that you created.

  ```
  # clresource create -g fmwk-rg \
  -t svm-rt \
  -p resource_dependencies=fmwk-rs \
  svm-rs
  ```

  `-g` *fmwk-rg*
  Specifies the resource group that you created in Step 5.

  `-p resource_dependencies=`*fmwk-rs*
  Specifies that this instance depends on the framework resource that you created in Step 7.

  svm-rs
  Specifies the name that you are assigning to the `SUNW.vucmm_svm` resource.

9 Bring online and in a managed state the Oracle RAC framework resource group, the multiple-owner volume-manager framework resource group if used, and their resources.

```
# clresourcegroup online -emM rac-fmwk-rg [vucmm-fmwk-rg]
```

rac-fmwk-rg
Specifies that the Oracle RAC resource group that you created in Step 2 is to be moved to the `MANAGED` state and brought online.

vucmm-fmwk-rg
Specifies that the multiple-owner volume-manager framework resource group that you created in Step 5 is to be moved to the `MANAGED` state and brought online.

Registering and Configuring Oracle ASM Resource Groups (CLI)

The following information is in this section:

▼ How to Register and Configure Oracle ASM Resource Groups in the Global Cluster (CLI)

Before You Begin

- Ensure that the Oracle Grid Infrastructure framework resource is created and that dependencies are configured between the Oracle RAC framework resource and the Oracle Grid Infrastructure framework resource.
- Ensure that the Oracle RAC framework resource group, the multiple-owner volume-manager framework resource group if used, and their resources are online.

1 Register the Oracle ASM resource types for the data service.

a. Register the `SUNW.scalable_asm_instance_proxy` resource type.

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

b. Register the `SUNW.scalable_asm_diskgroup_proxy` resource type.

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

2 Create the Oracle ASM instance and disk-group resource groups.

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

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

asm-dg-rg
Specifies the name of the Oracle ASM disk-group resource group.

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

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

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

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

5 **Create a `SUNW.scalable_asm_instance_proxy` resource and set the resource dependencies.**

- **For a local file system, perform the following commands:**

```
# clresource create -g asm-inst-rg \
-t SUNW.scalable_asm_instance_proxy \
-p ORACLE_HOME=Grid_home \
-p CRS_HOME=Grid_home \
-p "ORACLE_SID{node1}"=asm-instance1 \
-p "ORACLE_SID{node2}"=asm-instance2 \
-p resource_dependencies_offline_restart=crs-fmwk-rs \
-d asm-inst-rs
```

6 **Add a `SUNW.scalable_asm_diskgroup_proxy` resource type to the *asm-dg-rg* resource group.**

```
# 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
```

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

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

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

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

9 **Verify the Oracle ASM configuration.**

```
# clresource status +
```

▼ How to Register and Configure Oracle ASM Resource Groups in a Zone Cluster (CLI)

Before You Begin

- Ensure that the Oracle Grid Infrastructure framework resource is created and that dependencies are configured between the Oracle RAC framework resource and the Oracle Grid Infrastructure framework resource.
- Ensure that the Oracle RAC framework resource group, the multiple-owner volume-manager framework resource group if used, and their resources are online.

Perform all steps from the global zone.

1 **Register the Oracle ASM resource types for the data service.**

a. **Register the `SUNW.scalable_asm_instance_proxy` resource type.**

```
# clresourcetype register -Z zcname SUNW.scalable_asm_instance_proxy
```

b. **Register the `SUNW.scalable_asm_diskgroup_proxy` resource type.**

```
# clresourcetype register -Z zcname SUNW.scalable_asm_diskgroup_proxy
```

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

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

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

asm-dg-rg
Specifies the name of the Oracle ASM disk-group resource group.

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

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

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

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

5 **Create a `SUNW.scalable_asm_instance_proxy` resource and set the resource dependencies.**

- **For a local file system, perform the following command:**

```
# clresource create -Z zcname -g asm-inst-rg \
-t SUNW.scalable_asm_instance_proxy \
-p ORACLE_HOME=Grid_home \
-p CRS_HOME=Grid_home \
-p "ORACLE_SID{node1}"=asm-instance1 \
-p "ORACLE_SID{node2}"=asm-instance2 \
-p resource_dependencies_offline_restart=crs-fmwk-rs \
-d asm-inst-rs
```

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

`-t SUNW.scalable_asm_instance_proxy`
Specifies the type of the resource to add.

`-p ORACLE_HOME=`*Grid_home*
Sets the path to the Oracle Grid Infrastructure home directory.

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

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

6 **Add a `SUNW.scalable_asm_diskgroup_proxy` resource type to the** *asm-dg-rg* **resource group.**

```
# clresource create -Z zcname -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
```

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

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

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

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

9 **Verify the Oracle ASM configuration.**

```
# clresource status -Z zcname +
```

Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands

The tasks in this section are alternatives for the resource-configuration steps in "How to Register and Configure Storage Resources for Oracle Files by Using `clsetup`" on page 66. The following information is in this section:

The following resources to represent storage for Oracle files are required:

- Resources for scalable device groups and scalable file-system mount points

Resources for Scalable Device Groups and Scalable File-System Mount Points

If you are using Solaris Volume Manager for Sun Cluster, configure storage resources as follows:

- Create one scalable resource group to contain all resources for scalable device groups.
- Create one resource for each Solaris Volume Manager for Sun Cluster multi-owner disk set that you are using for Oracle files.

If you are using qualified NAS devices, configure storage resources as follows:

- Create one scalable resource group to contain all resources for scalable file-system mount points.
- Create one resource for each NFS file system on a qualified NAS device that you are using for Oracle files.

▼ How to Create a Resource for a Scalable Device Group in the Global Cluster

Perform this procedure on only one node of the cluster.

1 **Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2 **Create a scalable resource group to contain the scalable device group resource.**

Set a strong positive affinity by the resource group for the Oracle RAC framework resource group.

Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

```
# clresourcegroup create -p nodelist=nodelist \
-p desired_primaries=num-in-list \
-p maximum_primaries=num-in-list \
-p rg_affinities=++rac-fmwk-rg \
[-p rg_description="description"] \
-p rg_mode=Scalable \
scal-dg-rg
```

3 **Register the `SUNW.ScalDeviceGroup` resource type.**

```
# clresourcetype register SUNW.ScalDeviceGroup
```

4 **For each scalable device group that you are using for Oracle files, add an instance of the `SUNW.ScalDeviceGroup` resource type to the resource group that you created in Step 2.**

Set a strong dependency for the instance of `SUNW.ScalDeviceGroup` on the resource in the Oracle RAC framework resource group that represents the volume manager for the device group. Limit the scope of this dependency to only the node where the `SUNW.ScalDeviceGroup` resource is running.

```
# clresource create -t SUNW.ScalDeviceGroup -g scal-dg-rg \
-p resource_dependencies=fm-vol-mgr-rs{local_node} \
-p diskgroupname=disk-group \
scal-dg-rs
```

5 **Bring online and in a managed state the resource group that you created in Step 2.**

```
# clresourcegroup online -emM scal-dg-rg
```

How to Create a Resource for a Scalable Device Group in a Zone Cluster

Perform this procedure from the global cluster.

1. **Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2. **Create a scalable resource group to contain the scalable device group resource.**

   Set a strong positive affinity by the resource group for the Oracle RAC framework resource group.

   Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

   ```
   # clresourcegroup create -Z zcname -p nodelist=nodelist \
   -p desired_primaries=num-in-list \
   -p maximum_primaries=num-in-list \
   -p rg_affinities=++rac-fmwk-rg \
   [-p rg_description="description"] \
   -p rg_mode=Scalable \
   scal-dg-rg
   ```

3. **Register the `SUNW.ScalDeviceGroup` resource type.**

   ```
   # clresourcetype register -Z zcname SUNW.ScalDeviceGroup
   ```

4. **For each scalable device group that you are using for Oracle files, add an instance of the `SUNW.ScalDeviceGroup` resource type to the resource group that you created in Step 2.**

   Set a strong dependency for the instance of `SUNW.ScalDeviceGroup` on the resource in the Oracle RAC framework resource group that represents the volume manager for the device group. Limit the scope of this dependency to only the node where the `SUNW.ScalDeviceGroup` resource is running.

   ```
   # clresource create -Z zcname -t SUNW.ScalDeviceGroup -g scal-dg-rg \
   -p resource_dependencies=fm-vol-mgr-rs{local_node} \
   -p diskgroupname=disk-group \
   scal-dg-rs
   ```

5. **Bring online and in a managed state the resource group that you created in Step 2.**

   ```
   # clresourcegroup online -Z zcname-emM scal-dg-rg
   ```

▼ How to Create a Resource for a File-System Mount Point in the Global Cluster

Perform this procedure on only one node of the cluster.

1 **Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2 **Create a scalable resource group to contain the resource for a scalable file-system mount point.**

If you are also using a volume manager, set a strong positive affinity by the resource group for the resource group that contains the volume manager's scalable device-group resource. This resource group is created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.

Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

```
# clresourcegroup create -n nodelist \
-p desired_primaries=num-in-list \
-p maximum_primaries=num-in-list \
[-p rg_affinities=++scal-dg-rg] \
[-p rg_description="description"] \
-p rg_mode=Scalable \
scal-mp-rg
```

3 **Register the `SUNW.ScalMountPoint` resource type.**

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

4 **For each shared file system that requires a scalable file-system mount point resource, add an instance of the `SUNW.ScalMountPoint` resource type to the resource group that you created in Step 2.**

- **For each file system on a qualified NAS device, type the following command:**

  If you are also using a volume manager, set an offline-restart dependency by the instance of `SUNW.ScalMountPoint` on the resource for the scalable device group that is to store the file system. This resource is created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.

  ```
  # clresource create -t SUNW.ScalMountPoint -g scal-mp-rg \
  [-p resource_dependencies_offline_restart=scal-dg-rs] \
  -p mountpointdir=mp-path \
  -p filesystemtype=nas \
  -p targetfilesystem=nas-device:fs-name \
  nas-mp-rs
  ```

5 **Bring online and in a managed state the resource group that you created in Step 2.**

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

Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands

The tasks in this section are an alternative for the resource-configuration steps in "How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate" on page 87. The following information is in this section:

- "How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster" on page 185
- "How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure" on page 187
- "How to Create Oracle Solaris Cluster Resources in a Zone Cluster for Interoperation With Oracle Grid Infrastructure" on page 190

Resources for interoperation with Oracle Grid Infrastructure enable you to administer Oracle RAC database instances by using Oracle Solaris Cluster interfaces. These resources also ensure that dependencies by Oracle Grid Infrastructure resources on Oracle Solaris Cluster resources are met. These resources enable the high-availability frameworks that are provided by Oracle Solaris Cluster software and Oracle Grid Infrastructure to interoperate.

The following resources for interoperation are required:

- An Oracle Solaris Cluster resource to act as a proxy for the Oracle RAC database
- An Oracle Solaris Cluster resource to represent the Oracle Grid Infrastructure framework
- Oracle Grid Infrastructure resources to represent scalable device groups
- Oracle Grid Infrastructure resources to represent scalable file-system mount points

You must assign to an Oracle Grid Infrastructure resource that represents an Oracle Solaris Cluster resource a name in the following form:

`sun.`*`node`*`.`*`sc-rs`*

node
: Specifies the name of the node where the Oracle Grid Infrastructure resource is to run.

sc-rs
: Specifies the name of the Oracle Solaris Cluster resource that the Oracle Grid Infrastructure resource represents.

For example, the name of the Oracle Grid Infrastructure resource for node `pclus1` that represents the Oracle Solaris Cluster resource `scal-dg-rs` must be as follows:

```
sun.pclus1.scal-dg-rs
```

FIGURE D–1 Proxy Resources for Configurations With a Volume Manager

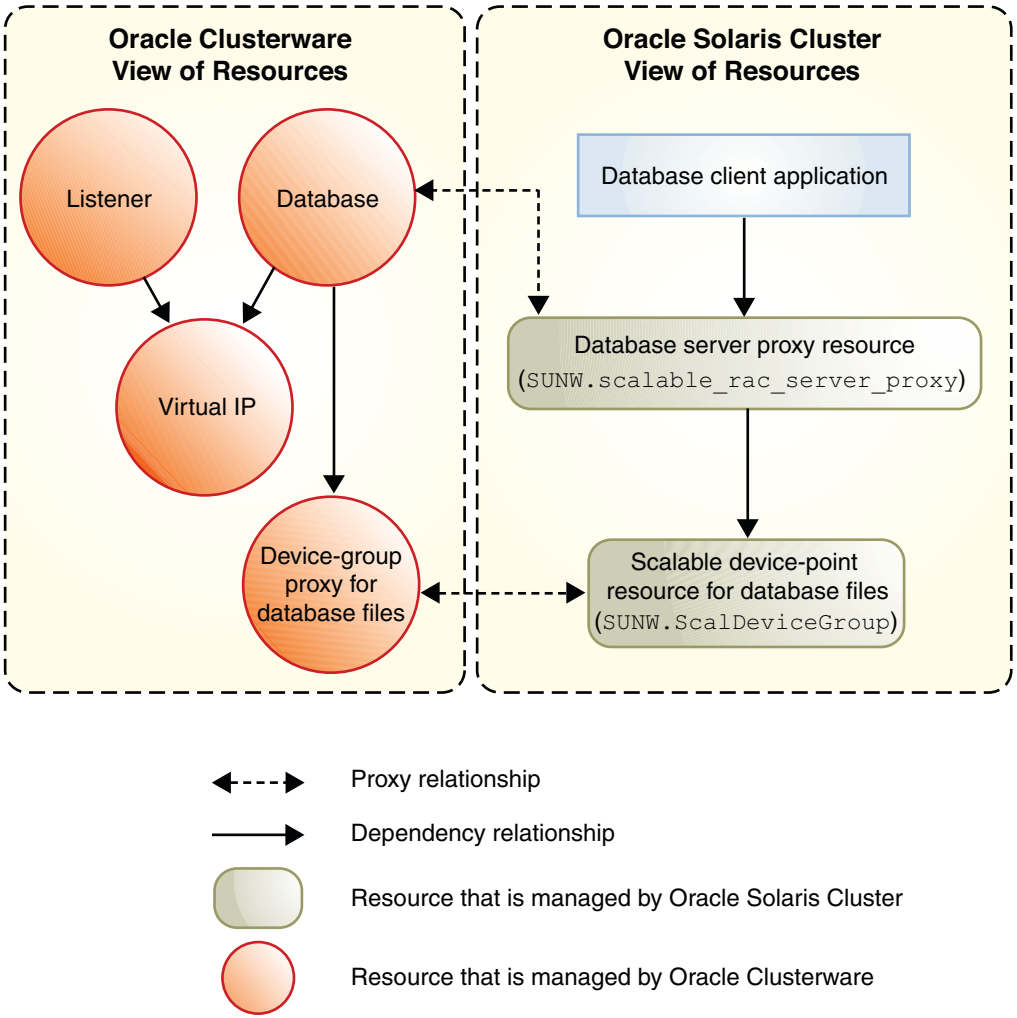

FIGURE D–2 Proxy Resources for Configurations With a Shared File System

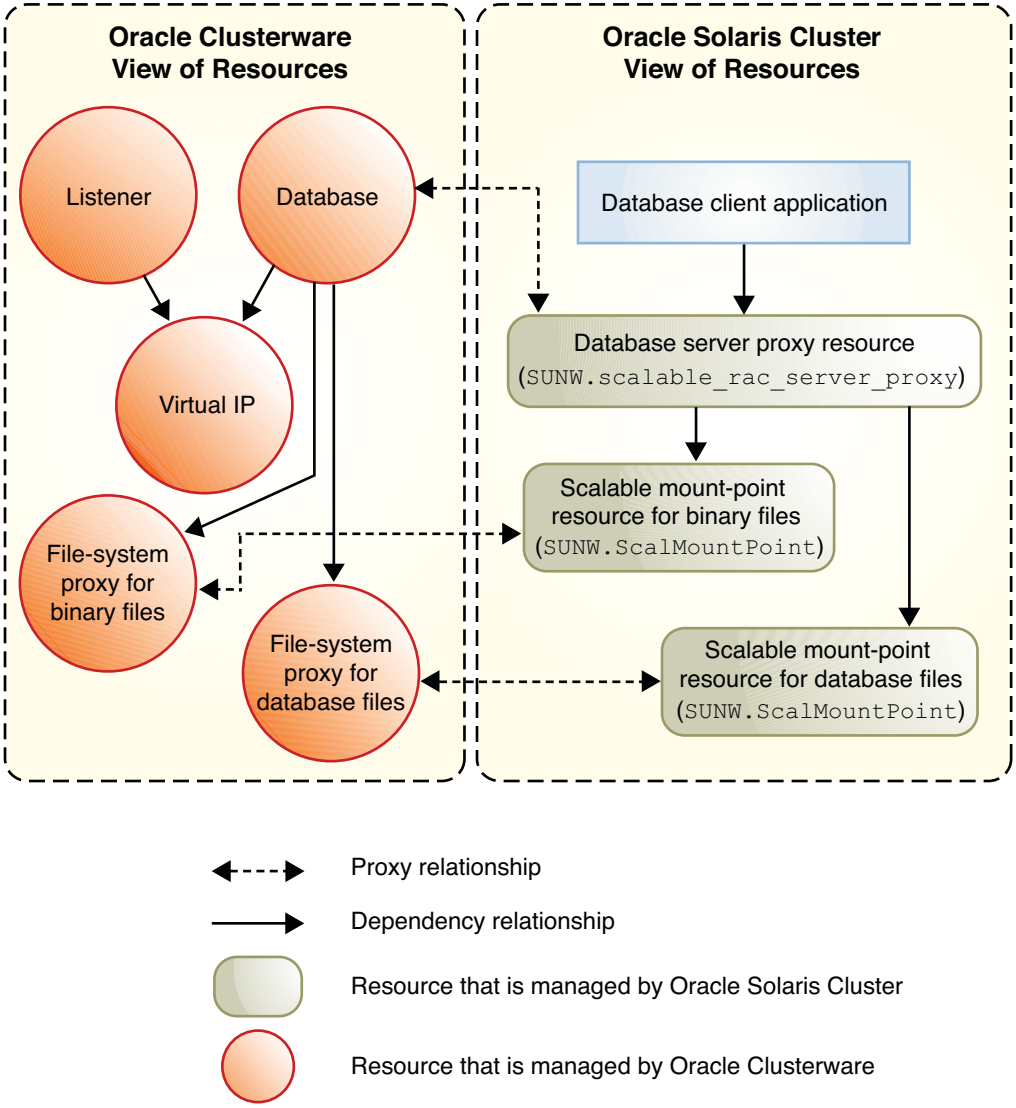

How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster

For Oracle RAC without Oracle ASM, perform this procedure to manually create an Oracle Grid Infrastructure resource that proxies the Oracle Solaris Cluster `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource.

This Oracle Grid Infrastructure resource ensures that the corresponding Oracle database is not started until the Oracle Grid Infrastructure resource is online. The resource comes online only if the corresponding `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource is online. And the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource only comes online if the actual volume-manager disk set or disk group or the mount point is online.

1 Become superuser on one node of the cluster.

2 Create the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type.

```
# Grid_home/bin/crsctl \
add type sun.storage_proxy.type \
-basetype cluster_resource \
-attr \
"ATTRIBUTE=ACTION_SCRIPT,TYPE=string", \
"ATTRIBUTE=HOSTING_MEMBERS,TYPE=string", \
"ATTRIBUTE=CARDINALITY,TYPE=string", \
"ATTRIBUTE=PLACEMENT,TYPE=string", \
"ATTRIBUTE=SCRIPT_TIMEOUT,TYPE=int", \
"ATTRIBUTE=RESTART_ATTEMPTS,TYPE=int", \
"ATTRIBUTE=ACL,TYPE=string", \
"ATTRIBUTE=VERSION,TYPE=string"
```

3 Create the Oracle Solaris Cluster `sun.`*resource* resource of type `sun.storage_proxy.type`.

The Oracle Grid Infrastructure resource name uses the form sun.*sc-resource*, where *sc-resource* is the name of the `SUNW.ScalDeviceGroup` or `SUNW.ScalMountPoint` resource.

```
# Grid_home/bin/crsctl add resource sun.sc-resource \
-type sun.storage_proxy.type \
-attr "ACTION_SCRIPT='/opt/SUNWscor/dsconfig/bin/scproxy_crs_action' \
CARDINALITY='number-nodes' \
SCRIPT_TIMEOUT='20' \
PLACEMENT='restricted' \
RESTART_ATTEMPTS='60' \
HOSTING_MEMBERS='nodelist' \
VERSION='1' "
```

`CARDINALITY`
: The number of nodes in the cluster membership

`HOSTING_MEMBERS`
: The space-separated list of nodes in the cluster membership

4 Determine the DBA group of the Oracle Grid Infrastructure installation.

```
# Grid_home/bin/osdbagrp
griddba-group
```

5 Set the primary group of the Oracle Grid Infrastructure storage proxy resource to the group determined in Step 4.

```
# Grid_home/bin/crsctl setperm resource sun.sc-resource -g "griddba-group"
```

6 **Determine the DBA group of the Oracle Database Software installation.**

```
# oracle_home/bin/osdbagrp
dba-group
```

7 **Set the group permissions of the Oracle Grid Infrastructure storage proxy resource to the group determined in Step 6.**

Omit this step if the Oracle Grid Infrastructure installation DBA group `griddba-group`, determined in Step 4, and the Oracle Database Software installation DBA group `dba-group`, determined in Step 6, are the same DBA group.

```
# Grid_home/bin/crsctl setperm resource sun.sc-resource -u "group:dba-group:r-x"
```

8 **Bring online the Oracle Grid Infrastructure storage proxy resource.**

```
# Grid_home/bin/crsctl start resource sun.sc-resource
```

See Also If you need to remove an Oracle Grid Infrastructure resource, perform procedures in "Removing an Oracle Grid Infrastructure Resource" on page 126.

▼ How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure

Note – If you are configuring Oracle RAC to run in a zone cluster, instead perform "How to Create Oracle Solaris Cluster Resources in a Zone Cluster for Interoperation With Oracle Grid Infrastructure" on page 190.

Perform this procedure on only one node of the cluster.

1 **Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2 **Register the `SUNW.crs_framework` resource type.**

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

3 **Add an instance of the `SUNW.crs_framework` resource type to the Oracle RAC framework resource group.**

For information about this resource group, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.

Set a strong dependency by the instance of `SUNW.crs_framework` on the instance of `SUNW.rac_framework` in the Oracle RAC framework resource group.

You might have configured a storage resource for the storage that you are using for Oracle Grid Infrastructure files. In this situation, set an offline-restart dependency by the instance of `SUNW.crs_framework` on the storage resource. Limit the scope of this dependency to only the node where the storage resource is running.

- If you are using a volume manager for database files, set the dependency on the resource that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.
- If you are using a file system for database files, set the dependency on the resource that you created in "How to Create a Resource for a File-System Mount Point in the Global Cluster" on page 182.

You might have configured a storage resource for the file system that you are using for Oracle Grid Infrastructure executables. In this situation, set an offline-restart dependency by the instance of `SUNW.crs_framework` on the storage resource. Limit the scope of this dependency to only the node where the storage resource is running. Set the dependency on the resource that you created in "How to Create a Resource for a File-System Mount Point in the Global Cluster" on page 182.

```
# clresource create -t SUNW.crs_framework \
-g rac-fmwk-rg \
-p resource_dependencies=rac-fmwk-rs \
[-p resource_dependencies_offline_restart=db-storage-rs{local_node} \
[,bin-storage-rs{local_node}]] \
crs-fmwk-rs
```

4 Create a scalable resource group to contain the proxy resource for the Oracle RAC database server.

Set a strong positive affinity by the scalable resource group for the Oracle RAC framework resource group.

You might have configured a storage resource for the storage that you are using for database files. In this situation, set a strong positive affinity by the scalable resource group for the resource group that contains the storage resource for database files.

- If you are using a volume manager for database files, set a strong positive for the resource group that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.
- If you are using a file system for database files, set a strong positive affinity for the resource group that you created in "How to Create a Resource for a File-System Mount Point in the Global Cluster" on page 182.

Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

```
# clresourcegroup create -n nodelist \
-p maximum_primaries=num-in-list \
-p desired_primaries=num-in-list \
-p rg_affinities=++rac-fmwk-rg[,++db-storage-rg] \
[-p rg_description="description"] \
-p rg_mode=Scalable \
rac-db-rg
```

5 **Register the `SUNW.scalable_rac_server_proxy` resource type.**

```
# clresourcetype register SUNW.scalable_rac_server_proxy
```

6 **Add an instance of the `SUNW.scalable_rac_server_proxy` resource type to the resource group that you created in Step 4.**

Set a strong dependency by the instance of `SUNW.scalable_rac_server_proxy` on the instance of `SUNW.rac_framework` in the Oracle RAC framework resource group.

Set an offline-restart dependency by the instance of `SUNW.scalable_rac_server_proxy` on the instance of `SUNW.crs_framework` that you created in Step 3.

You might have configured a storage resource for the storage that you are using for database files. In this situation, set an offline-restart dependency by the instance of `SUNW.scalable_rac_server_proxy` on the storage resource. Limit the scope of this dependency to only the node where the storage resource is running.

- If you are using a volume manager for database files, set the dependency on the resource that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.
- If you are using a file system for database files, set the dependency on the resource that you created in "How to Create a Resource for a File-System Mount Point in the Global Cluster" on page 182.

Set a different value of the `oracle_sid` extension property for each node that can master the resource.

```
# clresource create -g rac-db-rg \
-t SUNW.scalable_rac_server_proxy \
-p resource_dependencies=rac-fmwk-rs \
-p resource_dependencies_offline_restart=crs-fmk-rs[, db-storage-rs] \
-p oracle_home=ora-home \
-p crs_home=Grid_home \
-p db_name=db-name \
-p oracle_sid{node1-id}=sid-node1 \
[ -p oracle_sid{node2-id}=sid-node2...] \
rac-srvr-proxy-rs
```

7 **Bring online the resource group that you created in Step 4.**

```
# clresourcegroup online -emM rac-db-rg
```

How to Create Oracle Solaris Cluster Resources in a Zone Cluster for Interoperation With Oracle Grid Infrastructure

Note – If you are configuring Oracle RAC to run in the global cluster, instead perform "How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure" on page 187.

Perform this procedure on only one node of the cluster.

Note – When a step in the procedure requires running the Oracle Solaris Cluster commands in a zone cluster, you should run the command from the global cluster and use the `-Z` option to specify the zone cluster.

1 **Become superuser or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2 **Register the `SUNW.crs_framework` resource type.**

```
# clresourcetype register -Z zcname SUNW.crs_framework
```

3 **Add an instance of the `SUNW.crs_framework` resource type to the Oracle RAC framework resource group.**

For information about this resource group, see "Registering and Configuring the Oracle RAC Framework Resource Group" on page 51.

Set a strong dependency by the instance of `SUNW.crs_framework` on the instance of `SUNW.rac_framework` in the Oracle RAC framework resource group.

You might have configured a storage resource for the storage that you are using for Oracle Grid Infrastructure files. In this situation, set an offline-restart dependency by the instance of `SUNW.crs_framework` on the storage resource. Limit the scope of this dependency to only the node where the storage resource is running.

- If you are using a volume manager for database files, set the dependency on the resource that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.

```
# clresource create -Z zcname -t SUNW.crs_framework \
-g rac-fmwk-rg \
-p resource_dependencies=rac-fmwk-rs \
[-p resource_dependencies_offline_restart=db-storage-rs{local_node} \
[,bin-storage-rs{local_node}]] \
crs-fmwk-rs
```

4 **Create a scalable resource group to contain the proxy resource for the Oracle RAC database server.**

Set a strong positive affinity by the scalable resource group for the Oracle RAC framework resource group.

You might have configured a storage resource for the storage that you are using for database files. In this situation, set a strong positive affinity by the scalable resource group for the resource group that contains the storage resource for database files.

- If you are using a volume manager for database files, set a strong positive affinity for the resource group that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.

Tip – If you require Support for Oracle RAC to run on all cluster nodes, specify the `-S` option in the command that follows and omit the options `-n`, `-p maximum_primaries`, `-p desired_primaries`, and `-p rg_mode`.

```
# clresourcegroup create -Z zcname -n nodelist \
-p maximum_primaries=num-in-list \
-p desired_primaries=num-in-list \
-p rg_affinities=++rac-fmwk-rg[,db-storage-rg] \
[-p rg_description="description"] \
-p rg_mode=Scalable \
rac-db-rg
```

5 **Register the `SUNW.scalable_rac_server_proxy` resource type.**

```
# clresourcetype register -Z zcname SUNW.scalable_rac_server_proxy
```

6 **Add an instance of the `SUNW.scalable_rac_server_proxy` resource type to the resource group that you created in Step 4.**

Set a strong dependency by the instance of `SUNW.scalable_rac_server_proxy` on the instance of `SUNW.rac_framework` in the Oracle RAC framework resource group.

Set an offline-restart dependency by the instance of `SUNW.scalable_rac_server_proxy` on the instance of `SUNW.crs_framework` that you created in Step 3.

You might have configured a storage resource for the storage that you are using for database files. In this situation, set an offline-restart dependency by the instance of `SUNW.scalable_rac_server_proxy` on the storage resource. Limit the scope of this dependency to only the node where the storage resource is running.

- If you are using a volume manager for database files, set the dependency on the resource that you created in "How to Create a Resource for a Scalable Device Group in the Global Cluster" on page 180.

Set a different value of the `oracle_sid` extension property for each node that can master the resource.

```
# clresource create -Z zcname -g rac-db-rg \
-t SUNW.scalable_rac_server_proxy \
-p resource_dependencies=rac-fmwk-rs \
-p resource_dependencies_offline_restart=crs-fmk-rs \
[, db-storage-rs, bin-storage-rs] \
-p oracle_home=ora-home \
-p crs_home=Grid_home \
-p db_name=db-name \
-p oracle_sid{node1-id}=sid-node1 \
[ -p oracle_sid{node2-id}=sid-node2...] \
rac-srvr-proxy-rs
```

7 Bring online the resource group that you created in Step 4.

```
# clresourcegroup online -Z zcname -emM rac-db-rg
```

Index

C

D

E

F

N

O

P

Q

R

S

T

U

V

W

Z