

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

ORACLE®

Part No: E69468
February 2019

Part No: E69468

Copyright © 2000, 2019, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support 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.

Référence: E69468

Copyright © 2000, 2019, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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

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

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

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

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

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Using This Documentation	19
1 Installing Support for Oracle RAC	21
Overview of the Installation Process for Support for Oracle RAC	21
Preinstallation Considerations	21
General Requirements	22
Hardware and Software Requirements	23
Storage Management Requirements	24
Using Oracle Data Guard With Support for Oracle RAC	29
Preparing the Oracle Solaris Cluster Nodes	29
Before You Begin	30
▼ How to Bypass the NIS Name Service	30
▼ How to Create the DBA Group and the DBA User Accounts	31
▼ How to Configure Shared Memory for Support for Oracle RAC Software in the Global Cluster	35
▼ How to Configure Shared Memory for Support for Oracle RAC Software in a Zone Cluster	36
▼ How to Set the Necessary Privileges for Support for Oracle RAC Software in a Zone Cluster	38
▼ How to Configure the Logical Hostname Resources or Virtual IP Addresses for Support for Oracle RAC Software in a Zone Cluster	38
Installing the Support for Oracle RAC Package	39
▼ How to Install the Support for Oracle RAC Package	39
2 Configuring Storage for Oracle Database Files	41
Summary of Configuration Tasks for Storage for Oracle Database Files	41
Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files	41

Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Database Files	43
Tasks for Configuring Hardware RAID Support for Oracle Database Files	45
Tasks for Configuring Oracle ASM for Oracle Database Files	45
Tasks for Configuring Qualified NAS Devices for Oracle Database Files	46
Tasks for Configuring a Cluster File System for Oracle Database Files	47
Installing Storage Management Software With Support for Oracle RAC	47
Using Solaris Volume Manager for Sun Cluster	48
Using Hardware RAID Support	49
Using the StorageTek QFS Shared File System	51
Using Oracle ASM	56
Using a Cluster File System	59
3 Registering and Configuring the Resource Groups	63
Registering and Configuring the Support for Oracle RAC Framework Resource Group	63
Tools for Registering and Configuring the Support for Oracle RAC Framework Resource Group	64
▼ How to Register and Configure the Support for Oracle RAC Framework Resource Group (clsetup)	64
Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group	67
Tools for Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group	68
▼ How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group (clsetup)	68
Creating a Global Device Group for the Support for Oracle RAC Database	71
▼ How to Create a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster for the Oracle RAC Database	71
Registering and Configuring Storage Resources for Oracle Database Files	78
Tools for Registering and Configuring Storage Resources for Oracle Database Files	79
▼ How to Register and Configure Storage Resources for Oracle Database Files (clsetup)	79
Registering and Configuring the Oracle ASM Resource Group	83
Tools for Registering and Configuring the Oracle ASM Resource Group	83
▼ How to Register and Configure the Oracle ASM Resource Group (clsetup)	84

4 Enabling Support for Oracle RAC to Run in a Cluster	89
Overview of Tasks for Enabling Support for Oracle RAC to Run in a Cluster	89
Installing Oracle RAC Software	90
Installing Binary Files and Configuration Files on a Shared File System	90
Overriding Networking Defaults for Oracle Grid Infrastructure	91
Next Steps	91
Verifying the Installation of Oracle RAC	91
▼ How to Verify the Installation of Oracle RAC	91
Creating an Oracle ASM Instance and Disk Groups	92
▼ How to Create an Oracle ASM Instance and Disk Groups	92
Creating an Oracle Grid Infrastructure Framework Resource	92
▼ How to Create an Oracle Grid Infrastructure Framework Resource	93
Creating an Oracle Database	96
▼ How to Specify the Location of Data Files on a Shared File System	96
Configuring Resources for Support for Oracle RAC Database Instances	97
Tools for Registering and Configuring Resources for Support for Oracle RAC Database Instances	97
▼ How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate	98
Verifying the Installation and Configuration of Support for Oracle RAC	103
▼ How to Verify the Configuration of the Support for Oracle RAC Framework Resource Group	104
▼ How to Verify the Configuration of the Multiple-Owner Volume-Manager Framework Resource Group	104
▼ How to Verify the Configuration of Storage Resources for Oracle Database Files	105
▼ How to Verify the Configuration of Resources for Support for Oracle RAC Database Instances	107
▼ How to Verify the Correct Behavior for Shutdown and Booting of the Cluster	109
5 Administering Support for Oracle RAC	111
Overview of Administration Tasks for Support for Oracle RAC	111
Automatically Generated Names for Oracle Solaris Cluster Objects	112
Administering Support for Oracle RAC Databases From Oracle Solaris Cluster Software	112
Effects of State Changes to Oracle Solaris Cluster Resources for Support for Oracle RAC Database Instances	113

Tuning Support for Oracle RAC	115
Guidelines for Setting Timeouts	116
Reservation Step Timeout	116
Oracle Clusterware Framework Resource Method Timeouts	117
Tuning the Support for Oracle RAC Fault Monitors	117
Operation of the Fault Monitor for a Scalable Device Group	118
Operation of the Fault Monitor for Scalable File-System Mount Points	119
Obtaining Core Files for Troubleshooting DBMS Timeouts	120
6 Troubleshooting Support for Oracle RAC	121
Verifying the Status of Support for Oracle RAC	121
▼ How to Verify the Status of Support for Oracle RAC	121
Examples of the Status of Support for Oracle RAC	122
Sources of Diagnostic Information	132
Common Problems and Their Solutions	132
Failure of a Support for Oracle RAC Framework Resource Group	133
Failure of a Multiple-Owner Volume-Manager Framework Resource Group	136
SUNW.qfs Registration Fails Because the Registration File Is Not Found	138
Node Panic Caused by a Timeout	138
Failure of Services that Depend on a SUNW.ScalDeviceGroup Resource	139
Failure of a SUNW.rac_framework or SUNW.vucmm_framework Resource to Start	139
SUNW.rac_framework Failure-to-Start Status Messages	140
SUNW.vucmm_framework Failure-to-Start Status Messages	140
▼ How to Recover From the Timing Out of the START Method	141
Failure of a Resource to Stop	142
7 Modifying an Existing Configuration of Support for Oracle RAC	143
Overview of Tasks for Modifying an Existing Configuration of Support for Oracle RAC	143
Modifying Online the Resource for a Scalable Device Group	144
▼ How to Modify Online the Resource for a Scalable Device Group	144
Extending an Existing Configuration of Support for Oracle RAC	144
▼ How to Add Support for Oracle RAC to Selected Nodes	145
▼ How to Add a Volume Manager Resource to the SUNW.vucmm_framework Resource Group	151
Removing an Oracle Grid Infrastructure Resource	153

▼ How to Remove a Dependency	153
▼ How to Delete the <i>sun.resource</i> Resource	154
Removing Support for Oracle RAC	155
▼ How to Remove Support for Oracle RAC From a Cluster	155
▼ How to Remove Support for Oracle RAC From Selected Nodes	162
8 Upgrading Support for Oracle RAC	173
Upgrading Resources in Support for Oracle RAC	173
Re-creating the Oracle Grid Infrastructure <i>sun.storage_proxy.type</i> Resource Type and Resources	174
Information for Registering the New Versions of Support for Oracle RAC Resource Types	175
Adding Storage Resources for Oracle Database Files	176
Adding Resources for Interoperation With Oracle Grid Infrastructure	177
A Sample Configurations of This Data Service	179
Sample Support for Oracle RAC Configurations in the Global Cluster	180
Sample Support for Oracle RAC Configurations in a Zone Cluster	186
B Preset Actions for DBMS Errors and Logged Alerts	193
C Support for Oracle RAC Extension Properties	201
SUNW.crs_framework Extension Properties	201
SUNW.rac_framework Extension Properties	202
SUNW.scalable_acfs_proxy Extension Properties	202
SUNW.scalable_asm_diskgroup_proxy Extension Properties	203
SUNW.scalable_asm_instance_proxy Extension Properties	205
SUNW.scalable_rac_server_proxy Extension Properties	206
SUNW.ScalDeviceGroup Extension Properties	209
SUNW.ScalMountPoint Extension Properties	211
SUNW.vucmm_framework Extension Properties	214
SUNW.vucmm_svm Extension Properties	214
SUNW.wait_zc_boot Extension Properties	217
D Command-Line Alternatives	219

Setting Support for Oracle RAC Extension Properties	219
Registering and Configuring the Framework Resource Groups by Using Oracle Solaris Cluster Maintenance Commands	220
Overview of the Framework Resource Groups	220
▼ How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands	221
▼ How to Register and Configure the Support for Oracle RAC Framework Resource Group in a Zone Cluster by Using Oracle Solaris Cluster Maintenance Commands	224
Registering and Configuring Oracle ASM Resource Groups (CLI)	226
▼ How to Register and Configure Oracle ASM Resource Groups in the Global Cluster (CLI)	226
▼ How to Register and Configure Oracle ASM Resource Groups in a Zone Cluster (CLI)	228
Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands	229
Resources for Scalable Device Groups and Scalable File-System Mount Points	230
Resources for the StorageTek QFS Metadata Server	231
▼ How to Create a Resource for a Scalable Device Group in the Global Cluster	231
▼ How to Create a Resource for a Scalable Device Group in a Zone Cluster	232
▼ How to Register and Configure Resources for the StorageTek QFS Metadata Server in the Global Cluster	233
▼ How to Register and Configure Resources for the StorageTek QFS Metadata Server for a Zone Cluster	235
▼ How to Create a Resource for a File-System Mount Point in the Global Cluster	236
▼ How to Create a Resource for a File-System Mount Point in Zone Cluster	238
Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands	240
▼ How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster	243
▼ How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure	244
▼ How to Create Oracle Solaris Cluster Resources in a Zone Cluster for Interoperation With Oracle Grid Infrastructure	247

Index 251

Figures

FIGURE 1	Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster	180
FIGURE 2	Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System	181
FIGURE 3	Configuration of Support for Oracle RAC With StorageTek QFS Shared File System and Hardware RAID	182
FIGURE 4	Configuration of Support for Oracle RAC With a NAS Device	183
FIGURE 5	Configuration of Support for Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster	184
FIGURE 6	Configuration of Support for Oracle RAC With Oracle ASM and Hardware RAID	185
FIGURE 7	Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster in a Zone Cluster	186
FIGURE 8	Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System in a Zone Cluster	187
FIGURE 9	Configuration of Support for Oracle RAC With StorageTek QFS Shared File System and Hardware RAID in a Zone Cluster	188
FIGURE 10	Configuration of Support for Oracle RAC With a NAS Device in a Zone Cluster	189
FIGURE 11	Configuration of Support for Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster in a Zone Cluster	190
FIGURE 12	Configuration of Support for Oracle RAC With Oracle ASM and Hardware RAID in a Zone Cluster	191
FIGURE 13	Proxy Resources for Configurations With a Volume Manager	241
FIGURE 14	Proxy Resources for Configurations With a Shared File System	242

Tables

TABLE 1	Tasks for Installing Support for Oracle RAC	21
TABLE 2	Storage Management Schemes for Oracle DBMS Files	25
TABLE 3	Storage Management Schemes for Oracle Grid Infrastructure Files	26
TABLE 4	Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files in the Global Cluster	42
TABLE 5	Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files in a Zone Cluster	42
TABLE 6	Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Files in the Global Cluster	43
TABLE 7	Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Database Files in a Zone Cluster	44
TABLE 8	Tasks for Configuring Hardware RAID Support for Oracle Database Files	45
TABLE 9	Tasks for Configuring Oracle ASM for Oracle Database Files	45
TABLE 10	Tasks for Configuring Qualified NAS Devices for Oracle Database Files	46
TABLE 11	Tasks for Configuring a PxFs-Based Cluster File System for Oracle Database Files	47
TABLE 12	Tasks for Configuring an Oracle ACFS File System for Oracle Database Files	47
TABLE 13	Tasks for Enabling Support for Oracle RAC to Run in a Cluster	89
TABLE 14	Administration Tasks for Support for Oracle RAC	111
TABLE 15	Propagation of State Changes Between Oracle Solaris Cluster Resources and Oracle Grid Infrastructure Resources	114
TABLE 16	Comparisons of States for Oracle Solaris Cluster Resources and Oracle Grid Infrastructure Resources	115
TABLE 17	Resource Types for Support for Oracle RAC Fault Monitors	118
TABLE 18	Tasks for Modifying an Existing Configuration of Support for Oracle RAC	143
TABLE 19	Support for Oracle RAC Resource Types	175

TABLE 20	Preset Actions for DBMS Errors	193
TABLE 21	Preset Actions for Logged Alerts	199

Examples

EXAMPLE 1	Setting a Name Service Lookup Entry	31
EXAMPLE 2	Creating the DBA Group and the DBA User Accounts	35
EXAMPLE 3	Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster	76
EXAMPLE 4	Setting the Reservation Step Timeout	116
EXAMPLE 5	Setting the Oracle Clusterware Framework Resource Start Timeout	117
EXAMPLE 6	Status of a Faulty Support for Oracle RAC Framework Resource Group	123
EXAMPLE 7	Status of a Faulty Support for Oracle RAC Database Resource Group	126
EXAMPLE 8	Status of an Operational Support for Oracle RAC Configuration	129
EXAMPLE 9	Adding Support for Oracle RAC to Selected Nodes	149
EXAMPLE 10	Removing Support for Oracle RAC From a Cluster	159
EXAMPLE 11	Removing Support for Oracle RAC From Selected Nodes	168

Using This Documentation

- **Overview** – Describes the Oracle Solaris Cluster Support for Oracle Real Application Clusters (Support for Oracle RAC) data service.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Documentation and resources for this product and related products are available at http://www.oracle.com/pls/topic/lookup?ctx=E69294_01.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

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” on page 21](#)
- [“Preinstallation Considerations” on page 21](#)
- [“Preparing the Oracle Solaris Cluster Nodes” on page 29](#)
- [“Installing the Support for Oracle RAC Package” on page 39](#)

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 Tasks for Installing Support for Oracle RAC

Task	Instructions
Plan your installation	“Preinstallation Considerations” on page 21
Prepare the Oracle Solaris Cluster nodes	“Preparing the Oracle Solaris Cluster Nodes” on page 29
Install data service packages	“Installing the Support for Oracle RAC Package” on page 39

Preinstallation Considerations

This section contains the following preinstallation information:

- [“General Requirements” on page 22](#)
- [“Hardware and Software Requirements” on page 23](#)
- [“Storage Management Requirements” on page 24](#)
- [“Using Oracle Data Guard With Support for Oracle RAC” on page 29](#)

General Requirements

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

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

- **The Support for Oracle RAC framework.** These resources enable Support for 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 Support for Oracle RAC framework. For more information, see [“Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63](#).
- **Storage for Oracle Database files.** These resources provide fault monitoring and automatic fault recovery for volume managers and file systems that store Oracle Database files. Configuring storage resources for Oracle Database files is optional. For more information, see [“Registering and Configuring Storage Resources for Oracle Database Files” on page 78](#).
- **Support for 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 Support for Oracle RAC.

Configuring resources, which enable Oracle Solaris Cluster software to administer Support for Oracle RAC database instances, is optional. For more information, see [“Configuring Resources for Support for Oracle RAC Database Instances”](#) on page 97.

Note - When you use Support for 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 a Support for 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 Support 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 Support for 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”](#) on page 23
- [“Oracle Grid Infrastructure Software Requirements”](#) on page 24
- [“Software License Requirements”](#) on page 24
- [“Supported Topology Requirements”](#) on page 24
- [“Software Update Installation Requirements”](#) on page 24

For information about supported versions, see <http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf>.

Oracle Solaris Cluster Framework Requirements

Support for Oracle RAC requires a functioning cluster with the initial cluster framework already installed. See [Installing and Configuring an Oracle Solaris Cluster 4.4 Environment](#) 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 the information about configuring operating systems in the *Oracle Grid Infrastructure Installation Guide* for your version of Oracle Grid Infrastructure software.

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 Database Files” on page 25](#)
- [“Storage Management Requirements for Oracle Grid Infrastructure” on page 26](#)
- [“Storage Management Requirements for the Support for Oracle RAC Database” on page 27](#)
- [“Storage Management Requirements for Oracle Database Binary Files and Configuration Files” on page 27](#)
- [“Storage Management Schemes Supported by Zone Clusters” on page 28](#)

Storage Management Requirements for Oracle Database Files

Support for Oracle RAC enables you to use the storage management schemes for Oracle Database files that are listed in the following tables. The tables summarize the types of Oracle Database 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 Database files.

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

- Yes Indicates that the storage management scheme can store the type of Oracle Database file.
- * Indicates that the storage management scheme can store the type of Oracle Database file starting with Oracle Database version 12c release 1, but *not* for version 11g release 2.
- No Indicates that the storage management scheme *cannot* store the type of Oracle Database file.

TABLE 2 Storage Management Schemes for Oracle DBMS Files

Oracle DBMS File Type	Solaris Volume Manager for Sun Cluster Scheme	Hardware RAID Scheme	StorageTek QFS Shared File System Scheme	Qualified NAS Devices Scheme	Oracle ASM Scheme	Cluster File System Scheme	Oracle ACFS File System Scheme	Local Disks Scheme
Installation binary files	No	No	Yes	Yes	No	Yes	Yes	Yes
Configuration files	No	No	No	Yes	No	Yes	Yes	Yes
System parameter file (SPFILE)	No	No	Yes	Yes	Yes	Yes	No	No
Alert files	No	No	Yes	Yes	No	Yes	Yes	Yes
Trace files	No	No	Yes	Yes	No	Yes	Yes	Yes
Data files	Yes	Yes	Yes	Yes	Yes	No	*	No
Control files	Yes	Yes	Yes	Yes	Yes	No	*	No
Online redo log files	Yes	Yes	Yes	Yes	Yes	No	*	No
Archived redo log files	No	No	Yes	Yes	Yes	Yes	*	No
Flashback log files	No	No	Yes	Yes	Yes	Yes	*	No

Oracle DBMS File Type	Solaris Volume Manager for Sun Cluster Scheme	Hardware RAID Scheme	StorageTek QFS Shared File System Scheme	Qualified NAS Devices Scheme	Oracle ASM Scheme	Cluster File System Scheme	Oracle ACFS File System Scheme	Local Disks Scheme
Recovery files [†]	No	No	Yes	Yes	Yes	No	*	No

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

TABLE 3 Storage Management Schemes for Oracle Grid Infrastructure Files

Oracle Grid Infrastructure File Type	Solaris Volume Manager for Sun Cluster Scheme	Hardware RAID Scheme	StorageTek QFS Shared File System Scheme	Qualified NAS Devices Scheme	Oracle ASM Scheme	Cluster File System Scheme	Oracle ACFS File System Scheme	Local Disks Scheme
Installation binary files	No	No	No	Yes	No	No	No	Yes
OCR files	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Voting disk	Yes	Yes	Yes	Yes	Yes	Yes	No	No

For more information, see the documentation for database storage options in the Oracle Database installation guide for your version of Oracle Database.

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
- StorageTek QFS shared file systems, with either hardware RAID support or Solaris Volume Manager for Sun Cluster
- Qualified network-attached storage (NAS) devices
- Oracle ASM
- Cluster file systems
- (Starting with Oracle Database 12c) Oracle ACFS file systems

Storage Management Requirements for the Support for Oracle RAC Database

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

- Solaris Volume Manager for Sun Cluster
- Hardware redundant array of independent disks (RAID) support
- StorageTek QFS shared file systems with hardware RAID support or with Solaris Volume Manager for Sun Cluster
- Qualified network-attached storage (NAS) devices
- Oracle ASM
- (Starting with Oracle Database 12c) Oracle ACFS file systems



Caution - To avoid experiencing possible performance degradation or failures of the SUNW.ScalDeviceGroup probe, do not configure Oracle ASM with Solaris Volume Manager mirrored logical volumes.

Storage Management Requirements for Oracle Database Binary Files and Configuration Files

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

- The local disks of each cluster node. See [“Using Local Disks for Oracle Database Binary Files and Configuration Files” on page 28](#) for additional information.
- A shared file system from the following list:

Note - Oracle Grid Infrastructure binaries cannot reside on a cluster file system or on a StorageTek QFS shared file system.

- StorageTek QFS shared file systems
- File systems on a qualified NAS device
- PxFs-based cluster file systems
- ACFS file systems

See [“Using a Shared File System for Oracle Database Binary Files and Configuration Files” on page 28](#) for additional information.

Using Local Disks for Oracle Database Binary Files and Configuration Files

Placing the Oracle Database binary files and configuration files on the individual cluster nodes enables you to upgrade the Oracle Database 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 Database application without shutting down the data service, see your Oracle Database documentation.

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

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

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

Note - Oracle Grid Infrastructure binaries cannot reside on a cluster file system or a StorageTek QFS shared file system.

The following shared file systems are supported:

- StorageTek QFS shared file systems
- File systems on a qualified NAS device
- PxFS-based cluster file systems using Solaris Volume Manager
- ACFS file systems

If you put the Oracle Database binary files and 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 Database application. If a short period of downtime for upgrades is acceptable, place a single copy of the Oracle Database binary files and configuration files on a shared file system.

Storage Management Schemes Supported by Zone Clusters

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

- Solaris Volume Manager for Sun Cluster
- A StorageTek QFS shared file system, using either Solaris Volume Manager or hardware RAID support
- A file system on a qualified NAS device with fencing
- Oracle ASM
- (Starting with Oracle Database 12c) Oracle ACFS file systems

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 Data Guard documentation.

Preparing the Oracle Solaris Cluster Nodes

Preparing the Oracle Solaris Cluster nodes modifies the configuration of the operating system to enable Support for 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 Support for 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 Database installation is incomplete. An incomplete Oracle installation causes Support for Oracle RAC to fail during startup.

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

- Configuring shared memory for Support for Oracle RAC software in a zone cluster
- Setting necessary privileges for Support for Oracle RAC software in a zone cluster

- Configuring logical hostname resources for Support for Oracle RAC software in a zone cluster

This section contains the following information:

- [“Before You Begin” on page 30](#)
- [“How to Bypass the NIS Name Service” on page 30](#)
- [“How to Create the DBA Group and the DBA User Accounts” on page 31](#)
- [“How to Configure Shared Memory for Support for Oracle RAC Software in the Global Cluster” on page 35](#)
- [“How to Configure Shared Memory for Support for Oracle RAC Software in a Zone Cluster” on page 36](#)
- [“How to Set the Necessary Privileges for Support for Oracle RAC Software in a Zone Cluster” on page 38](#)
- [“How to Configure the Logical Hostname Resources or Virtual IP Addresses for Support for Oracle RAC Software in a Zone Cluster” on page 38](#)

Before You Begin

Before you prepare the Oracle Solaris Cluster nodes, ensure that all preinstallation tasks for Support for Oracle RAC are completed. For more information, see your Support for 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. **Assume the root role 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\(8\)](#) and [nsswitch.conf\(5\)](#) man pages.

Example 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:/system/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 31.

▼ 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 RAC 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 RAC documentation.

Perform this task on each cluster node.

1. **On the cluster node where you are performing this task, assume the root role.**

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 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:

```
# passwd 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. Type the password.

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

c. Retype the password.

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 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\(8\)](#)
- [group\(5\)](#)
- [passwd\(5\)](#)
- [shadow\(5\)](#)

Next Steps Go to [“How to Configure Shared Memory for Support for Oracle RAC Software in the Global Cluster”](#) on page 35.

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

To enable the Support for 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. **Assume the `root` role on a cluster node.**

2. Update the shared memory configuration information.

See information about configuring kernel parameters on Oracle Solaris in the Oracle Database installation guide for your version of Oracle Database.

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 Support for 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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

Next Steps If you are using zone clusters, go to [“How to Configure Shared Memory for Support for Oracle RAC Software in a Zone Cluster”](#) on page 36.

If not, go to [“Installing the Support for Oracle RAC Package”](#) on page 39.

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

To configure shared memory for the Support for 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 Support for Oracle RAC Software in the Global Cluster”](#) on page 35.

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 Support 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. Assume the root role 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 Support for 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\(8\)](#) 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 \(clsetup\)](#)” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

Next Steps Go to “[How to Set the Necessary Privileges for Support for Oracle RAC Software in a Zone Cluster](#)” on page 38.

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

You should set the necessary privileges to enable Support for 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. **Assume the root role 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"
clzonecluster:zcname>commit
```

3. **Prevent Oracle Clusterware time synchronization from running in active mode.**
 - a. **Log in to the zone-cluster node as root.**
 - b. **Create an empty `/etc/inet/ntp.conf` file.**

```
# touch /etc/inet/ntp.conf
```

Next Steps Go to [“How to Configure the Logical Hostname Resources or Virtual IP Addresses for Support for Oracle RAC Software in a Zone Cluster”](#) on page 38.

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

To support Oracle Grid Infrastructure virtual IP resources in Support for 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 Support for Oracle RAC.

1. **Assume the root role 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
```

3. If the SCAN hostname resolves to multiple IP addresses, configure a separate global network resource for each IP address that the SCAN hostname resolves to.

```
# clzonecluster configure zcname
clzonecluster:zcname>add net
clzonecluster:zcname:net>set address=SCAN-address1
clzonecluster:zcname:net>end
clzonecluster:zcname>add net
clzonecluster:zcname:net>set address=SCAN-address2
clzonecluster:zcname:net>end
clzonecluster:zcname>add net
clzonecluster:zcname:net>set address=SCAN-address3
clzonecluster:zcname:net>end
clzonecluster:zcname>commit
```

Next Steps Go to [“Installing the Support for Oracle RAC Package”](#) on page 39.

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, assume the root role.**
2. **Ensure that the data service package is available from the configured publisher and that the solaris and ha-cluster publishers are valid.**

```
# pkg list -a ha-cluster/data-service/oracle-database ha-cluster/library/ucmm
# pkg publisher
PUBLISHER                TYPE    STATUS  P  LOCATION
solaris                   origin  online  F  solaris-repository
ha-cluster                 origin  online  F  ha-cluster-repository
```

For information about setting the solaris publisher, see [“Adding, Modifying, or Removing Package Publishers” in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

Tip - Use the -nv options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the -nv options, run the command again without the -n option to actually perform the installation or update. If you do get error messages, run the command again with more -v options (for example, -nvv) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,” in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

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 about updating your software, see [Chapter 10, “Updating Software Packages” in *Updating Your Oracle Solaris Cluster 4.4 Environment*](#).

◆◆◆ 2 CHAPTER 2

Configuring Storage for Oracle Database Files

This chapter explains how to configure storage for Oracle Database files.

- [“Summary of Configuration Tasks for Storage for Oracle Database Files” on page 41](#)
- [“Installing Storage Management Software With Support for Oracle RAC” on page 47](#)

Summary of Configuration Tasks for Storage for Oracle Database Files

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

- [“Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files” on page 41](#)
- [“Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Database Files” on page 43](#)
- [“Tasks for Configuring Hardware RAID Support for Oracle Database Files” on page 45](#)
- [“Tasks for Configuring Oracle ASM for Oracle Database Files” on page 45](#)
- [“Tasks for Configuring Qualified NAS Devices for Oracle Database Files” on page 46](#)
- [“Tasks for Configuring a Cluster File System for Oracle Database Files” on page 47](#)

Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files

The following tables summarize the tasks for configuring the StorageTek QFS shared file system and provides cross-references to detailed instructions for performing the tasks. The first

table provides information on Support for Oracle RAC running in the global cluster and the second table provide information on Support for Oracle RAC running in a zone cluster.

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

TABLE 4 Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files in the Global Cluster

Task	Instructions
Install and configure the StorageTek QFS shared file system.	“Using the StorageTek QFS Shared File System” on page 51
Install and configure the other storage management scheme that you are using with the StorageTek QFS shared file system.	Follow the procedure for your storage management scheme: <ul style="list-style-type: none"> ■ “Using Solaris Volume Manager for Sun Cluster” on page 48. ■ “Using Hardware RAID Support” on page 49.
Register and configure the Support for Oracle RAC framework resource group.	Follow the procedure for your choice of configuration tool: <ul style="list-style-type: none"> ■ “Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63 (clsetup) ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221.
If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Support for 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 71
Register and configure storage resources for Oracle Database files.	Follow the procedure for your choice of configuration tool: <ul style="list-style-type: none"> ■ “Registering and Configuring Storage Resources for Oracle Database Files” on page 78 (clsetup) ■ “Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229

TABLE 5 Tasks for Configuring the StorageTek QFS Shared File System for Oracle Database Files in a Zone Cluster

Task	Instructions
Install and configure the StorageTek QFS shared file system in the global cluster.	“Using the StorageTek QFS Shared File System” on page 51
Install and configure the other storage management scheme that you are using with the StorageTek QFS shared file system in the global cluster.	Follow the procedure for your storage management scheme:

Task	Instructions
	<ul style="list-style-type: none"> ■ “Using Solaris Volume Manager for Sun Cluster” on page 48 ■ “Using Hardware RAID Support” on page 49
Register and configure the Support for Oracle RAC framework resource group in the global cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63 (clsetup) ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221
If you are using Solaris Volume Manager for Sun Cluster, create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Support for 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 71
Configure StorageTek QFS shared file system for the zone cluster.	See “How to Add a Oracle HSM Shared File System to a Zone Cluster (CLI)” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>
Register and configure the storage resources for Oracle Database files in the zone cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring Storage Resources for Oracle Database Files” on page 78 (clsetup) ■ “Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229.

Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Database 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 6 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 48

Task	Instructions
Register and configure the multiple-owner volume manager resource group.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group (clsetup)” on page 68 ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221
Create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Support for 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 71
Register and configure storage resources for Oracle Database files	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring Storage Resources for Oracle Database Files” on page 78 (clsetup) ■ “Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229

TABLE 7 Tasks for Configuring Solaris Volume Manager for Sun Cluster for Oracle Database 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 48
Register and configure the multiple-owner volume manager framework resource group in the global cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group (clsetup)” on page 68 ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221
Create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for the Support for 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 71
Configure Solaris Volume Manager devices in a zone cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “How to Add a Global Storage Device to a Zone Cluster (clsetup)” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>

Task	Instructions
	<ul style="list-style-type: none"> ■ “How to Add a Disk Set to a Zone Cluster (CLI)” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>
Register and configure storage resources for Oracle Database files in the zone cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring Storage Resources for Oracle Database Files” on page 78 (clsetup) ■ “Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229

Tasks for Configuring Hardware RAID Support for Oracle Database Files

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

TABLE 8 Tasks for Configuring Hardware RAID Support for Oracle Database Files

Task	Instructions
Configure hardware RAID support.	“Using Hardware RAID Support” on page 49

Note - For information configuring hardware RAID for a zone cluster, see [“Adding Storage Devices to a Zone Cluster” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

Tasks for Configuring Oracle ASM for Oracle Database Files

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

TABLE 9 Tasks for Configuring Oracle ASM for Oracle Database Files

Task	Instructions
Configure devices for Oracle ASM.	“Using Oracle ASM” on page 56

Note - For information about configuring Oracle ASM for a zone cluster, see [“Adding Storage Devices to a Zone Cluster”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

Tasks for Configuring Qualified NAS Devices for Oracle Database 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 10 Tasks for Configuring Qualified NAS Devices for Oracle Database Files

Task	Instructions
Install and configure the qualified NAS device.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ Managing Network-Attached Storage Devices in an Oracle Solaris Cluster 4.4 Environment (Oracle Solaris Cluster maintenance commands) ■ To use Oracle Solaris Cluster Manager for this task, in the Storage panel, go to NAS Devices and click the New NAS Device button.
Register and configure the Support for Oracle RAC framework resource group in a global cluster or zone cluster.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63 (c1setup) ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221.
Register and configure storage resources for Oracle files, including Support for Oracle RAC to support NAS NFS.	<p>Follow the procedure for your choice of configuration tool:</p> <ul style="list-style-type: none"> ■ “Registering and Configuring Storage Resources for Oracle Database Files” on page 78 (c1setup) ■ “Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229

Tasks for Configuring a Cluster File System for Oracle Database Files

Cluster file system types are the PxFs-based cluster file system and the Oracle ACFS file system. The following tables summarize the tasks for configuring a PxFs-based cluster file system or an Oracle ACFS 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 11 Tasks for Configuring a PxFs-Based Cluster File System for Oracle Database Files

Task	Instructions
Install and configure the cluster file system.	“Using a Cluster File System” on page 59
Register and configure the Support for Oracle RAC framework resource group.	Follow the procedure for your choice of configuration tool: <ul style="list-style-type: none"> ■ “Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63 (clsetup) ■ “How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221

TABLE 12 Tasks for Configuring an Oracle ACFS File System for Oracle Database Files

Task	Instructions
Install and configure the Oracle ACFS file system, and register and configure the Support for Oracle RAC framework resource group.	“Creating Oracle ACFS File Systems” in <i>Installing and Configuring an Oracle Solaris Cluster 4.4 Environment</i>

Installing Storage Management Software With Support for Oracle RAC

Install the software for the storage management schemes that you are using for Oracle Database files. For more information, see [“Storage Management Requirements” on page 24](#).

Note - For information about how to install and configure qualified NAS devices with Support for Oracle RAC, see [Managing Network-Attached Storage Devices in an Oracle Solaris Cluster 4.4 Environment](#).

This section contains the following information:

- [“Using Solaris Volume Manager for Sun Cluster” on page 48](#)
- [“Using Hardware RAID Support” on page 49](#)
- [“Using the StorageTek QFS Shared File System” on page 51](#)
- [“Using Oracle ASM” on page 56](#)
- [“Using a Cluster File System” on page 59](#)

Using Solaris Volume Manager for Sun Cluster

Always install Solaris Volume Manager software, which includes the Solaris Volume Manager for Oracle Solaris 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 node into the zone cluster. All administration tasks for Solaris Volume Manager for Oracle Solaris Cluster are performed in the global-cluster node, even when the Solaris Volume Manager for Oracle Solaris Cluster volume is used in a zone cluster.

When a Support for 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 Oracle Solaris Cluster volume in the global cluster. In this case, the scalable device group resource belongs to this zone cluster.

When a Support for Oracle RAC installation inside a zone cluster runs directly on the Solaris Volume Manager for Oracle Solaris 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 Database files that you can store by using Solaris Volume Manager for Sun Cluster, see [“Storage Management Requirements” on page 24](#).

▼ 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 *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

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 a Solaris Volume Manager for Sun Cluster volume into a zone cluster, see [“How to Add a Global Storage Device to a Zone Cluster \(clsetup\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* or [“How to Add a Disk Set to a Zone Cluster \(CLI\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

Next Steps Ensure that all other storage management schemes that you are using for Oracle Database files are installed. After all storage management schemes that you are using for Oracle Database files are installed, go to [Chapter 3, “Registering and Configuring the Resource Groups”](#).

Using Hardware RAID Support

For information about the types of Oracle Database files that you can store by using hardware RAID support, see [“Storage Management Requirements”](#) on page 24.

▼ 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.

```
# format
```

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(8CL)` 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/rdisk/c0t2d0
d2              phys-schost-1:/dev/rdisk/c0t3d0
d3              phys-schost-2:/dev/rdisk/c4t4d0
d3              phys-schost-1:/dev/rdisk/c1t5d0
d4              phys-schost-2:/dev/rdisk/c3t5d0
d4              phys-schost-1:/dev/rdisk/c2t5d0
d5              phys-schost-2:/dev/rdisk/c4t4d1
d5              phys-schost-1:/dev/rdisk/c1t5d1
d6              phys-schost-2:/dev/rdisk/c3t5d1
d6              phys-schost-1:/dev/rdisk/c2t5d1
d7              phys-schost-2:/dev/rdisk/c0t2d0
d8              phys-schost-2:/dev/rdisk/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/rdisk/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 Global Storage Device to a Zone Cluster \(clsetup\)](#)” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* or “[How to Add a DID Device to a Zone Cluster \(CLI\)](#)” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

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 cluster 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 `sN` 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/rdisk/d4`. If you choose to use slice `s0` on these devices, specify the raw device `/dev/did/rdisk/d4s0`.

Next Steps Ensure that all other storage management schemes that you are using for Oracle Database files are installed. After all storage management schemes that you are using for Oracle Database files are installed, go to [Chapter 3, “Registering and Configuring the Resource Groups”](#).

Using the StorageTek QFS Shared File System

The StorageTek QFS shared file system is always installed in the global-cluster node, even when a file system is used by a zone cluster. You configure specific StorageTek QFS shared file system into a specific zone cluster using the `clzc` command. The scalable mount-point resource belongs to this zone cluster. The metadata server resource, `SUNW.qfs`, belongs to the global cluster.

You must use the StorageTek QFS shared file system with one storage management scheme from the following list:

- Hardware RAID support

- Solaris Volume Manager for Sun Cluster

Distributing Oracle Database Files Among StorageTek QFS Shared File Systems

You can store all the files that are associated with Support for Oracle RAC on the StorageTek QFS shared file system.

Distribute these files among several file systems as explained in the subsections that follow.

- [“StorageTek QFS File Systems for RDBMS Binary Files and Related Files” on page 52](#)
- [“StorageTek QFS File Systems for Database Files and Related Files” on page 52](#)

StorageTek QFS File Systems for RDBMS Binary Files and Related Files

For RDBMS binary files and related files, create one file system in the cluster to store the files. The RDBMS binary files and related files are as follows:

- Oracle Database relational database management system (RDBMS) binary files
- Oracle Database 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`)
- Oracle Clusterware binary files

Note - Beginning with Oracle Database version 11g release 2, Oracle Clusterware binaries cannot reside on StorageTek QFS shared file systems.

StorageTek QFS File Systems for Database Files and Related Files

For database files and related files, determine whether you require one file system for each database or multiple file systems for each database.

- For simplicity of configuration and maintenance, create one file system to store these files for all Support for Oracle RAC instances of the database.

- To facilitate future expansion, create multiple file systems to store these files for all Support for Oracle RAC instances of the database.

Note - If you are adding storage for an existing database, you must create additional file systems for the storage that you are adding. In this situation, distribute the database files and related files among the file systems that you will use for the database.

Each file system that you create for database files and related files must have its own metadata server. For information about the resources that are required for the metadata servers, see [“Resources for the StorageTek QFS Metadata Server” on page 231](#).

The database files and related files are as follows:

- Data files
- Control files
- Online redo log files
- Archived redo log files
- Flashback log files
- Recovery files
- Oracle Database cluster registry (OCR) files
- Oracle Clusterware voting disk

Optimizing the Performance of the StorageTek QFS Shared File System

For optimum performance with Solaris Volume Manager for Sun Cluster, configure the volume manager and the file system as follows:

- Use Solaris Volume Manager for Sun Cluster to mirror the logical unit numbers (LUNs) of your disk arrays.
- If you require striping, configure the striping by using the file system's stripe option.

Mirroring the LUNs of your disk arrays involves the following operations:

- Creating RAID-0 metadevices
- Using the RAID-0 metadevices or Solaris Volume Manager soft partitions of such metadevices as StorageTek QFS devices

The input/output (I/O) load on your system might be heavy. In this situation, ensure that the LUN for Solaris Volume Manager metadata or hardware RAID metadata maps to a different physical disk than the LUN for data. Mapping these LUNs to different physical disks ensures that contention is minimized.

▼ How to Install and Configure the StorageTek QFS Shared File System

Before You Begin You might use Solaris Volume Manager metadevices as devices for the shared file systems. In this situation, ensure that the metaset and its metadevices are created and available on all cluster nodes before configuring the shared file systems.

1. **Ensure that the StorageTek QFS software is installed on all nodes of the global cluster where Support for Oracle RAC is to run.**
2. **Ensure that each StorageTek QFS shared file system is correctly created for use with Support for Oracle RAC.**

For each StorageTek QFS shared file system, set the correct mount options for the types of Oracle Database files that the file system is to store.

- For the file system that contains binary files, configuration files, alert files, and trace files, use the default mount options.

Note - Beginning with Oracle version 11g release 2, Oracle Clusterware binaries cannot reside on StorageTek QFS shared file systems.

- For the file systems that contain database data files, control files, online redo log files, and archived redo log files, or Oracle Clusterware OCR and Voting files, set the mount options as follows:
 - In the `/etc/vfstab` file set the `shared` option.
 - In the `/etc/opt/SUNWsamfs/samfs.cmd` file or the `/etc/vfstab` file, set the following options:

```
fs=fs-name
stripe=width
mh_write
qwrite
forcedirectio
rdlease=300    Set this value for optimum performance.
wrlease=300    Set this value for optimum performance.
aplease=300    Set this value for optimum performance.
```

fs-name Specifies the name that uniquely identifies the file system.

width Specifies the required stripe width for devices in the file system. The required stripe width is a multiple of the file system's disk

allocation unit (DAU). *width* must be an integer that is greater than or equal to 1.

Note - Ensure that settings in the `/etc/vfstab` file do not conflict with settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file. Settings in the `/etc/vfstab` file override settings in the `/etc/opt/SUNWsamfs/samfs.cmd` file.

3. Mount each StorageTek QFS shared file system that you are using for Oracle files.

```
# mount mount-point
```

mount-point Specifies the mount point of the file system that you are mounting.

4. If you are using a zone cluster, configure the StorageTek QFS shared file system into the zone cluster. Otherwise, go to [Step 5](#).

For information about configuring StorageTek QFS shared file system into a zone cluster, see [“How to Add a Cluster File System to a Zone Cluster \(clsetup\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* or [“How to Add a Oracle HSM Shared File System to a Zone Cluster \(CLI\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

5. Change the ownership of each file system that you are using for Oracle Database files.

Note - If you have configured StorageTek QFS shared file system for a zone cluster, perform this step in that zone cluster.

Change the file-system ownership as follows:

- Owner: the database administrator (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 31.

```
# chown user-name:group-name mount-point
```

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

<i>group-name</i>	Specifies the name of the DBA group. This group is normally named dba.
<i>mount-point</i>	Specifies the mount point of the file system whose ownership you are changing.

6. Grant to the owner of each file system whose ownership you changed in [Step 5](#) read access and write access to the file system.

Note - When StorageTek QFS shared file system is configured for a zone cluster, you need to perform this step in that zone cluster.

```
# chmod u+rw mount-point
```

<i>mount-point</i>	Specifies the mount point of the file system to whose owner you are granting read access and write access.
--------------------	--

Next Steps Ensure that all other storage management schemes that you are using for Oracle Database files are installed. After all storage management schemes that you are using for Oracle Database 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 57](#).
- **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 71](#).

Note - If you use Oracle ASM in a global cluster as well as in a zone cluster configured in that global cluster, you must ensure that, in each particular cluster, Oracle ASM can see only those devices that are intended for its use, whether in the global zone or in a zone cluster. If Oracle ASM can see devices that are used by Oracle ASM in a different cluster, this can cause start problems for Oracle Clusterware or Oracle Grid Infrastructure, because Oracle ASM sees the devices as already mounted elsewhere.

For information about the types of Oracle Database files that you can store by using Oracle ASM, see [“Storage Management Requirements” on page 24](#).

Note - When a Support for Oracle RAC installation in a zone cluster uses Oracle ASM, you must configure all the devices needed by that Support for 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, assume the root role.**
2. **Determine the identities of device identity (DID) devices that correspond to shared disks that are available in the cluster.**

Use the `cldevice(8CL)` 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/rdisk/c3t216000C0FF084E77d0
d5                  phys-schost-1:/dev/rdisk/c5t216000C0FF084E77d0
d5                  phys-schost-2:/dev/rdisk/c4t216000C0FF084E77d0
d5                  phys-schost-4:/dev/rdisk/c2t216000C0FF084E77d0
d6                  phys-schost-3:/dev/rdisk/c4t216000C0FF284E44d0
d6                  phys-schost-1:/dev/rdisk/c6t216000C0FF284E44d0
d6                  phys-schost-2:/dev/rdisk/c5t216000C0FF284E44d0
d6                  phys-schost-4:/dev/rdisk/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/rdisk/d5
Full Device Path:        phys-schost-1:/dev/rdisk/c5t216000C0FF084E77d0
Replication:              none
```

```

default_fencing:                global

DID Device Name:                /dev/did/rdisk/d6
Full Device Path:              phys-schost-1:/dev/rdisk/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 Global Storage Device to a Zone Cluster \(clsetup\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* or [“How to Add a DID Device to a Zone Cluster \(CLI\)”](#) in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

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 `sX` to the DID device name that you obtained in [Step 3](#), where `X` is the slice number.

```

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

For more information about changing the ownership and permissions of raw devices for use by Oracle ASM, see your Oracle ASM 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/rdisk/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/rdisk/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/rdisk/*'
```

For more information, see your Oracle ASM documentation.

Next Steps Ensure that all other storage management schemes that you are using for Oracle Database files are installed. After all storage management schemes that you are using for Oracle Database files are installed, go to [Chapter 3, “Registering and Configuring the Resource Groups”](#).

Using a Cluster File System

Support for Oracle RAC is supported on two types of cluster file systems:

- Cluster file systems that 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 Global Devices, Device Groups, and Cluster File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*
 - “Creating Cluster File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*
- Oracle Automatic Storage Management Cluster File System (Oracle ACFS)
 - For general information about how to create and mount an Oracle ACFS file system, see “Creating Oracle ACFS File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*.

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 Database Files That You Can Store on a PxFS-Based Cluster File System” on page 60

- [“Types of Oracle Database Files That You Can Store on an Oracle ACFS File System” on page 60](#)
- [“Optimizing Performance and Availability When Using a PxFs-Based Cluster File System” on page 61](#)
- [“How to Use a PxFs-Based Cluster File System” on page 61](#)
- [“How to Use an Oracle ACFS File System” on page 62](#)

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

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

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

- Oracle Database RDBMS binary files
- Oracle Database 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 Grid Infrastructure cluster registry (OCR) files
- Oracle Grid Infrastructure voting disk

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

Types of Oracle Database Files That You Can Store on an Oracle ACFS File System

You can store only these files that are associated with Support for Oracle RAC on the Oracle ACFS file system:

- Oracle Database RDBMS binary files
- Oracle Database configuration files (for example, `init.ora`, `tnsnames.ora`, `listener.ora`, and `sqlnet.ora`)

- Alert files (for example, `alert_sid.log`)
- Trace files (*.trc)
- (Starting with Oracle Database 12c) Data files
- (Starting with Oracle Database 12c) Control files
- (Starting with Oracle Database 12c) Online redo log files
- (Starting with Oracle Database 12c) Archived redo log files
- (Starting with Oracle Database 12c) Flashback log files
- (Starting with Oracle Database 12c) Recovery files

Note - You *must not* store Oracle Grid Infrastructure binary files, cluster registry (OCR) files, or voting disks on an Oracle ACFS file system. In addition, if you are using Oracle Database version 11g release 2, you also *must not* store data files, control files, online redo log files, archived redo log files, flashback log files, or Oracle Database recovery files on an Oracle ACFS 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 [“Device Group Ownership” in *Concepts for Oracle Solaris Cluster 4.4*](#).

▼ How to Use a PxFs-Based Cluster File System

1. Create and mount the cluster file system.

See [“Creating Cluster File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#) 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 Database 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.

Oracle Database File Type	Options
Oracle Database RDBMS binary files	global, logging
Oracle Database 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
Oracle Grid Infrastructure 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 Database files are installed. After all storage management schemes that you are using for Oracle Database files are installed, go to [Chapter 3, “Registering and Configuring the Resource Groups”](#).

▼ **How to Use an Oracle ACFS File System**

● **Create and mount the Oracle ACFS file system.**

See “Creating Oracle ACFS File Systems” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment* for information about how to create and mount the Oracle ACFS file system.

Next Steps Go to [Chapter 3, “Registering and Configuring the Resource Groups”](#).

Registering and Configuring the Resource Groups

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

- “Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63
- “Registering and Configuring the Multiple-Owner Volume-Manager Framework Resource Group” on page 67
- “Creating a Global Device Group for the Support for Oracle RAC Database” on page 71
- “Registering and Configuring Storage Resources for Oracle Database Files” on page 78
- “Registering and Configuring the Oracle ASM Resource Group” on page 83

Registering and Configuring the Support for Oracle RAC Framework Resource Group

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

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

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

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

- [“Tools for Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 64](#)
- [“How to Register and Configure the Support for Oracle RAC Framework Resource Group \(c1setup\)” on page 64](#)

Tools for Registering and Configuring the Support for Oracle RAC Framework Resource Group

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

- **The c1setup utility.** For more information, see [“How to Register and Configure the Support for Oracle RAC Framework Resource Group \(c1setup\)” on page 64](#).
- **Oracle Solaris Cluster maintenance commands.** For more information, see [Appendix D, “Command-Line Alternatives”](#).

The c1setup utility provides a wizard for configuring resources for the Support for 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 c1setup utility runs only in a node of the global cluster.

▼ How to Register and Configure the Support for Oracle RAC Framework Resource Group (c1setup)

Note - You can also use the Oracle Solaris Cluster Manager browser interface to perform this task, as well as optionally configuring resources for Oracle Clusterware and Oracle ASM. For Oracle Solaris Cluster Manager log-in instructions, see [“How to Access Oracle Solaris Cluster Manager” in *Administering an Oracle Solaris Cluster 4.4 Configuration*](#). After you log in, click Tasks and then click Oracle Real Application Clusters to start the wizard.

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

Note - The c1setup utility currently allows ongoing administration of an Oracle RAC framework running only in the global cluster. For ongoing administration of a Support for Oracle RAC framework configured in a zone cluster, you need to instead use the Oracle Solaris Cluster maintenance commands.

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 Support 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. **Assume the root role on any cluster node.**
2. **Start the c1setup utility, verify prerequisites, and navigate to the Support for Oracle RAC configuration section.**
3. **Choose the Support for Oracle RAC location.**
Follow the prompts to select either the global cluster or a zone cluster.
4. **Choose to configure the RAC Framework Resource Group.**
5. **Follow the prompts in the c1setup utility to provide the following information.**

Data to Provide	Guidelines	Your Values
Cluster nodes on which to run Support for Oracle RAC	<p>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.</p> <p>To choose nodes, type a comma-separated or space-separated list of the option numbers for the nodes you choose.</p>	
(Zone cluster only) Zone cluster name		

6. **Confirm the object names.**

- a. **If you require a different name for any Oracle Solaris Cluster objects, choose the name to change and type the new name in response to the prompt.**
 - b. **To confirm your selection of Oracle Solaris Cluster object names, type d.**
7. **To create the configuration, type c.**

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.

When the configuration process completes, you can press Return repeatedly to exit from the clsetup utility. If you prefer, you can leave the clsetup utility running while you perform other required tasks before using the utility again.
8. **Ensure that the Support for Oracle RAC framework resource group and its resources are online.**

- **For the global cluster, type the following commands:**

Display the resource group status .

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

If needed, bring online the resource group and resources.

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

- **For a zone cluster, from the global zone type the following commands:**

Display the resource group status.

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

If needed, bring online the resource group and resources.

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

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 Resource name: rac-framework-rs Resource group: rac-framework-rg	None.	Support for Oracle RAC framework resource.

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 67
None	“Registering and Configuring Storage Resources for Oracle Database Files” on page 78

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

Registering and configuring the multiple-owner volume-manager framework resource group enables Support for 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 node supports any volume manager used by Support for 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” on page 68](#)
- [“How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group \(clsetup\)” on page 68](#)

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 \(clsetup\)”](#) on page 68.
- **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 node of the global cluster.

▼ How to Register and Configure the Multiple-Owner Volume-Manager Framework Resource Group (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.

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

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

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.

Before You Begin

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

1. **Assume the root role on any cluster node.**
2. **Start the clsetup utility , verify prerequisites, and navigate to the Support for Oracle RAC configuration section.**
3. **Choose the Global Cluster as the Support for Oracle RAC cluster location.**
You can configure a multiple-owner volume-manager framework resource group only in the global cluster.
4. **Choose to configure the Multiple-Owner Volume-Manager Framework Resource Group.**
5. **Follow the prompts in the clsetup utility to choose the multiple-owner volume managers to use.**
Only installed volume managers are listed.
6. **Confirm the object names.**
 - a. **If you require a different name for any Oracle Solaris Cluster objects, choose the name to change and type the new name in response to the prompt.**
 - b. **To confirm your selection of Oracle Solaris Cluster object names, type d.**
7. **To create the configuration, type c.**
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.

When the configuration process completes, you can press Return repeatedly to exit from the clsetup utility. If you prefer, you can leave the clsetup utility running while you perform other required tasks before using the utility again.

8. Ensure that the multiple-owner volume-manager framework resource group and its resources are online.

By default, vucmm_framework_rg is the name that the clsetup utility assigns to the multiple-owner volume-manager framework resource group.

Display the resource group status .

```
# clresourcegroup status vucmm_framework_rg
```

If needed, bring online the resource group and resources.

```
# clresourcegroup online -eM vucmm_framework_rg
```

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 Oracle Solaris 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 71
None	“Registering and Configuring Storage Resources for Oracle Database Files” on page 78

Creating a Global Device Group for the Support for Oracle RAC Database

If you are using a volume manager for Oracle Database files, the volume manager requires a global device group for the Support for 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 71.

▼ 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 Support for Oracle RAC database, the StorageTek QFS shared file system, 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 39.
- 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 67.



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.

- Unless you are using the StorageTek QFS shared file system, do not create any file systems in the multi-owner disk set. In configurations without the StorageTek QFS shared file system, only the raw data file uses this 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/dN`, where *N* is the device number.

3. For the disk set that you created in [Step 1](#), create the volumes that the Support for Oracle RAC database or StorageTek QFS shared file system will use.

Tip - If you are creating many volumes for Oracle data files, you can simplify this step by using soft partitions. However, if you are using the StorageTek QFS shared file system and the I/O load on your system is heavy, use separate partitions for data and metadata. Otherwise, the performance of your system might be impaired. 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 `dV`, 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. 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`.

5. 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`.

6. 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.

7. (Configurations *without* the StorageTek QFS shared file system only) On each node that can own the disk set, change the ownership of each volume that you created in [Step 3](#).

If you are using the StorageTek QFS shared file system, omit this step.

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 31.

Ensure that you change ownership only of volumes that the Support for 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/dV`
- For raw devices: `/dev/md/setname/rdisk/dV`

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](#).

8. **(Configurations *without* the StorageTek QFS shared file system only) Grant read access and write access to the owner of each volume whose ownership you changed in [Step 7](#).**

If you are using the StorageTek QFS shared file system, omit this step.

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 Support for 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 7](#).

9. **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/rdisk/d` path for the Oracle ASM disk group, add the value `/dev/md/*/rdisk/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/*/rdisk/d*'
```

For more information, see your Oracle Database documentation.

Example 3 Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster

This example shows the sequence of operations that is required to create a multi-owner disk set in Solaris Volume Manager for Sun Cluster for a four-node cluster.

The disk set is to be used with the StorageTek QFS shared file system. This example does not show the creation of the StorageTek QFS shared file system on the devices that are added to the disk set.

1. To create the multi-owner disk set, the following command is run:

```
# metaset -s oradg -M -a -h pclus1 pclus2 pclus3 pclus4
```

The multi-owner disk set is named `oradg`. The nodes `pclus1`, `pclus2`, `pclus3`, and `pclus4` are added to this disk set.

2. To add global devices to the disk set, the following command is run:

```
# metaset -s oradg -a /dev/did/dsk/d8 /dev/did/dsk/d9 /dev/did/dsk/d15 \  
/dev/did/dsk/d16
```

The preceding command adds the following global devices to the disk set:

- `/dev/did/dsk/d8`
- `/dev/did/dsk/d9`
- `/dev/did/dsk/d15`
- `/dev/did/dsk/d16`

3. To create volumes for the disk set, the following commands are run:

```
# metainit -s oradg d10 1 1 /dev/did/dsk/d9s0  
# metainit -s oradg d11 1 1 /dev/did/dsk/d16s0  
# metainit -s oradg d20 1 1 /dev/did/dsk/d8s0  
# metainit -s oradg d21 1 1 /dev/did/dsk/d15s0
```

Each volume is created by a one-on-one concatenation of a slice as shown in the following table. The slices are *not* striped.

Volume	Slice
d10	/dev/did/dsk/d9s0
d11	/dev/did/dsk/d16s0
d20	/dev/did/dsk/d8s0
d21	/dev/did/dsk/d15s0

4. To verify that each node is correctly added to the multi-owner disk set, the following command is run:

```
# metaset -s oradg
```

```
Multi-owner Set name = oradg, Set number = 1, Master = pclus2
```

Host	Owner	Member
pclus1	multi-owner	Yes
pclus2	multi-owner	Yes
pclus3	multi-owner	Yes
pclus4	multi-owner	Yes

```
Drive Dbase
```

```
d8 Yes
```

```
d9 Yes
```

```
d15 Yes
```

```
d16 Yes
```

5. To verify that the multi-owner disk set is correctly configured, the following command is run:

```
# cldevicegroup show oradg
```

```
=== Device Groups ===
```

```
Device Group Name:          oradg
Type:                      Multi-owner_SVM
failback:                  false
Node List:                  pclus1, pclus2, pclus3, pclus4
preferenced:                false
numsecondaries:             0
diskset name:               oradg
```

6. To verify the online status of the multi-owner disk set, the following command is run:

```
# cldevicegroup status oradg

=== Cluster Device Groups ===

--- Device Group Status ---

Device Group Name      Primary      Secondary      Status
-----

--- Multi-owner Device Group Status ---

Device Group Name      Node Name      Status
-----
oradg                  pclus1         Online
                       pclus2         Online
                       pclus3         Online
                       pclus4         Online
```

Next Steps Go to [“Registering and Configuring Storage Resources for Oracle Database Files” on page 78.](#)

Registering and Configuring Storage Resources for Oracle Database 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 StorageTek QFS shared file system, with either Solaris Volume Manager for Sun Cluster multi-owner disk sets or hardware RAID
 - A file system on a qualified NAS device with Support for Oracle RAC:

Note - You can configure a NAS NFS file system in a zone cluster with the existing tools. See [“Tasks for Configuring Qualified NAS Devices for Oracle Database Files” on page 46.](#)

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

- [“Tools for Registering and Configuring Storage Resources for Oracle Database Files” on page 79](#)
- [“How to Register and Configure Storage Resources for Oracle Database Files \(c1setup\)” on page 79](#)

Tools for Registering and Configuring Storage Resources for Oracle Database 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 c1setup(8CL) utility.** For more information, see [“How to Register and Configure Storage Resources for Oracle Database Files \(c1setup\)” on page 79.](#)
- **Oracle Solaris Cluster maintenance commands.** For more information, see [“Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229.](#)

The c1setup 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.

Note - The c1setup utility runs only in a node of the global cluster.

▼ How to Register and Configure Storage Resources for Oracle Database Files (c1setup)

Perform this procedure from only one node of the cluster.

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.

Before You Begin Ensure that the following prerequisites are met:

- The Support for Oracle RAC framework resource group is created and is online. For more information, see [“Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63](#).
- 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 67](#).



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 47](#)
 - [“Creating a Global Device Group for the Support for Oracle RAC Database” on page 71](#)
- 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 Database files, if any
- The mount point of each shared file system that you are using for Oracle Database files, if any

1. **On one node of the cluster, assume the root role.**
2. **Start the `c1setup` utility, verify prerequisites, and navigate to the Support for Oracle RAC configuration section**
3. **Choose the Support for Oracle RAC location.**
Follow the prompts to select either the global cluster or a zone cluster.
4. **Choose Storage Resources for Oracle Files to configure.**
5. **Follow the prompts in the `c1setup` utility to provide the following information.**

Data to Provide	Guidelines	Your Values
If prompted, a storage management scheme for your Oracle Database files	<p>If no suitable resources for device groups exist, the clsetup utility prompts you to create a resource. Select the scalable device group to use, and whether to use logical devices or the entire disk group.</p> <p>If you select more than one existing resource, the selected resources must be in the same resource group.</p>	
Shared file-system mount points to use for Oracle files	<p>If no suitable resources for the shared file-system mount points exist, the clsetup utility prompts you to create a resource.</p> <p>To choose resources, type a comma-separated or space-separated list of the option numbers for the resources to use. If you select more than one existing resource, the selected resources must be in the same resource group.</p>	

6. Confirm the object names.

a. If you require a different name for any Oracle Solaris Cluster objects, choose the name to change and type the new name in response to the prompt.

b. To confirm your selection of Oracle Solaris Cluster object names, type d.

7. To create the configuration, type c.

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.

When the configuration process completes, you can press Return repeatedly to exit from the clsetup utility. If you prefer, you can leave the clsetup utility running while you perform other required tasks before using the utility again.

8. Ensure that the created resource groups are online.

Display the resource group status .

```
# clresourcegroup status
```

If needed, bring online each resource group.

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

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
<p>Resource type: SUNW.ScalDeviceGroup</p> <p>Resource name: <i>scal_{dg-name}-rs</i>, where <i>dg-name</i> is the name of the device group that the resource represents</p> <p>Resource group: <i>scal_{dg}-rg</i></p>	<p>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, the Solaris Volume Manager for Oracle Solaris Cluster resource.</p>	<p>Scalable device-group resource. One resource is created for each scalable device group that you are using for Oracle files.</p>
<p>Resource type: SUNW.qfs</p> <p>Resource name: <i>qfs-mp-dir-rs</i>, where <i>mp-dir</i> is the mount point of the file system, with / replaced by –</p> <p>Resource group: <i>qfs_{mds}-rg</i></p>	<p>Strong dependency on the scalable <i>wait_zc_boot</i> resource and scalable device-group resource, if any.</p> <p>If you are using StorageTek QFS without a volume manager, this resource does not depend on any other resources.</p>	<p>Resource for the StorageTek QFS metadata server. One resource is created for each StorageTek QFS shared file system that you are using for Oracle files.</p>
<p>Resource type: SUNW.ScalMountPoint</p> <p>Resource name: <i>scal-mp-dir-rs</i>, where <i>mp-dir</i> is the mount point of the file system, with / replaced by –</p> <p>Resource group: <i>scal_{mnt}-rg</i></p>	<p>Strong dependency on the resource for the StorageTek QFS metadata server, if any.</p> <p>Offline-restart dependency on the scalable device-group resource, if any.</p> <p>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.</p>	<p>Scalable file system mount point resource. One resource is created for each shared file system that you are using for Oracle files.</p>
<p>Resource type: SUNW.wait_zc_boot</p> <p>Resource name: <i>wait-zc-rs</i>, where <i>zc</i> is the zone cluster name.</p> <p>Resource group: <i>scal_{mnt}-rg</i></p>	<p>None</p>	<p>Resource to ensure that the StorageTek QFS shared file system configured to the zone cluster is mounted only after the zone cluster is booted.</p>

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 83.

Otherwise, go to [Chapter 4, “Enabling Support for 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” on page 83](#)
- [“How to Register and Configure the Oracle ASM Resource Group \(clsetup\)” on page 84](#)

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 \(clsetup\)” on page 84](#).
- **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 node of the global cluster.

▼ How to Register and Configure the Oracle ASM Resource Group (clsetup)

Note - You can also use the Oracle Solaris Cluster Manager browser interface to perform this task, as well as optionally configuring resources for Oracle Clusterware and Oracle ASM. For Oracle Solaris Cluster Manager log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*. After you log in, click Tasks and then click Oracle Real Application Clusters to start the wizard.

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.

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

- Before You Begin**
- Ensure that Oracle ASM disk groups are configured. For more information, see [“Using Oracle ASM”](#) on page 56.
 - Any project you created to run Oracle ASM exists in the `projects` name service database for the user that runs that application. If no custom project is created for the application, the default project is used. For more information, see the `projects(1)` man page.

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.

Note - If your Support for Oracle RAC configuration requires using Oracle ASM with Solaris Volume Manager mirrored logical volumes, you might experience either of the following conditions:

- Performance degradation. See Article 603825.1 at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for additional information.
- Failures of the SUNW.ScalDeviceGroup probe. These failures result in a loss of availability of any service that is dependent on the SUNW.ScalDeviceGroup resource.

You can mitigate these issues by increasing the IOTimeout property setting for the SUNW.ScalDeviceGroup resource type.

```
# clresource set -p IOTimeout=timeout-value-in-seconds resource-name
```

The default for the IOTimeout property is 30 seconds but can be set much higher. Determine the optimal setting for each system individually, starting with a value of at least 240 seconds.

1. **Assume the root role on any cluster node.**
2. **Start the c1setup utility, verify prerequisites, and navigate to the Support for Oracle RAC configuration section.**
3. **Choose the Support for Oracle RAC location.**
Select the global cluster or a zone cluster. The c1setup utility displays a list of Support for Oracle RAC components to configure.
4. **Choose to configure Automatic Storage Management (ASM).**
The c1setup utility displays a list of the cluster nodes to select.
5. **Follow the prompts in the c1setup utility to provide the following information.**

Data to Provide	Guidelines	Your
Cluster nodes on which to run Oracle ASM	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 choose nodes, type a comma-separated or space-separated list of the option numbers for the nodes you choose.	
Oracle ASM instance resource to use	If no Oracle ASM instance resource is available, the c1setup utility prompts you to create a resource.	
Oracle Grid Infrastructure home directory	Provide the full path.	

Data to Provide	Guidelines	Yes
Project where Oracle ASM is to run	If you did not create a project where Oracle ASM will run, select the default project.	
Oracle ASM SIDs	If the list of Oracle ASM system identifiers (SIDs) that are discovered in the cluster is not correct, you can change the SIDs.	
Oracle ASM disk-group resources	Choose whether to create a disk-group resource or to use an existing disk-group resource.	
Oracle ASM storage resources to manage the file-system mount point for the Oracle Grid Infrastructure home directory	If no storage resource exist, the c1setup utility creates one when you complete Oracle ASM configuration.	
Oracle ASM disk-group resources to manage Oracle ASM disk groups	Choose whether to create a disk-group resource or to use an existing disk-group resource.	
Oracle ASM disk groups		

6. Confirm the object names.

a. If you require a different name for any Oracle Solaris Cluster objects, choose the name to change and type the new name in response to the prompt.

b. To confirm your selection of Oracle Solaris Cluster object names, type d.

7. To create the configuration, type c.

The c1setup utility displays a progress message to indicate that the utility is running commands to create the configuration. When configuration is complete, the c1setup utility displays the commands that the utility ran to create the configuration.

When the configuration process completes, you can press Return repeatedly to exit from the c1setup utility. If you prefer, you can leave the c1setup utility running while you perform other required tasks before using the utility again.

Resource Configuration

The following table lists the default resource configuration that the c1setup utility creates when you complete this task.

Resource Name, Resource Type, and Resource Group	Dependencies	Description
Resource type: SUNW.scalable_rac_server_proxy	Strong dependency on the Support for Oracle RAC framework resource.	Support for Oracle RAC instance proxy resource
Resource name: rac_server_proxy-rs	Offline-restart dependency on the clustered	
Resource group: rac_server_proxy-rg	Oracle ASM disk-group resource.	

Resource Name, Resource Type, and Resource Group	Dependencies	Description
	<p>Offline-restart dependency on the Oracle Grid Infrastructure framework resource.</p> <p>When configured with Solaris Volume Manager for Sun Cluster, strong positive affinity by the Support for Oracle RAC instance proxy resource group for the Support for Oracle RAC framework resource group.</p> <p>When configured with hardware RAID, strong positive affinity by the Support for Oracle RAC instance proxy resource group for the clustered Oracle ASM disk-group resource group.</p>	
<p>Resource type: SUNW.oracle_asm_diskgroup</p> <p>Resource name: asm-dg-rs</p> <p>Resource group: asm-dg-rg</p>	<p>Strong positive affinity by the clustered Oracle ASM disk-group resource group for the clustered Oracle ASM instance resource group.</p> <p>When configured with Solaris Volume Manager for Sun Cluster:</p> <ul style="list-style-type: none"> ■ Strong dependency on the clustered Oracle ASM instance resource. ■ Offline-restart dependency on the scalable device-group resource for database files. <p>When configured with hardware RAID, offline-restart dependency on the clustered Oracle ASM disk-group resource.</p>	Clustered Oracle ASM disk-group resource
<p>SPARC: Resource type: SUNW.scalable_oracle_asm_instance_proxy</p> <p>Resource name: asm-inst-rs</p> <p>Resource group: asm-inst-rg</p>	<p>Offline-restart dependency on the Oracle Grid Infrastructure resource.</p> <p>When configured with hardware RAID, strong positive affinity by the clustered Oracle ASM instance resource group for the Oracle Clusterware framework resource group.</p>	Clustered Oracle ASM instance resource
<p>SPARC: Resource type: SUNW.sqfs</p> <p>Resource name: asm-home-sqfs-rs</p> <p>Resource group: asm-home-sqfs-rg</p>	<p>Strong dependency on the scalable wait_zc_boot resource and scalable device-group resource, if any.</p> <p>If you are using StorageTek QFS without a volume manager, this resource does not depend on any other resources.</p>	Resource for the Oracle ASM home on a StorageTek QFS metadata server. One resource is created for each StorageTek QFS shared file system that you are using for Oracle files.
<p>Resource type: SUNW.ScalMountPoint</p> <p>Resource name: asm-mp-rs</p>	<p>Strong dependency on the resource for the StorageTek QFS metadata server, if any.</p>	Resource for the Oracle Grid Infrastructure home on a scalable file system mount point. One resource is created for each shared

How to Register and Configure the Oracle ASM Resource Group (clsetup)

Resource Name, Resource Type, and Resource Group	Dependencies	Description
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.	file system that you are using for Oracle files.
Resource type: SUNW.ScalDeviceGroup Resource name: <i>scaldg-name-rs</i> , where <i>dg-name</i> is the name of the device group that the resource represents Resource group: <i>scaldg-rg</i>	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 Oracle Solaris 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 Support for Oracle RAC to Run in a Cluster”](#).

Enabling Support for Oracle RAC to Run in a Cluster

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

- [“Overview of Tasks for Enabling Support for Oracle RAC to Run in a Cluster” on page 89](#)
- [“Installing Oracle RAC Software” on page 90](#)
- [“Verifying the Installation of Oracle RAC” on page 91](#)
- [“Creating an Oracle ASM Instance and Disk Groups” on page 92](#)
- [“Creating an Oracle Grid Infrastructure Framework Resource” on page 92](#)
- [“Creating an Oracle Database” on page 96](#)
- [“Configuring Resources for Support for Oracle RAC Database Instances” on page 97](#)
- [“Verifying the Installation and Configuration of Support for Oracle RAC” on page 103](#)

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

Table 13, “Tasks for Enabling Support for Oracle RAC to Run in a Cluster,” on page 89 summarizes the tasks for enabling Support for Oracle RAC to run in a cluster.

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

TABLE 13 Tasks for Enabling Support for Oracle RAC to Run in a Cluster

Task	Instructions
Install the Support for Oracle RAC software.	“Installing Oracle RAC Software” on page 90
Verify the installation of the Support for Oracle RAC software.	“Verifying the Installation of Oracle RAC” on page 91

Task	Instructions
Create an Oracle ASM instance.	“Creating an Oracle ASM Instance and Disk Groups” on page 92
Create an Oracle Clusterware framework resource.	“Creating an Oracle Grid Infrastructure Framework Resource” on page 92
Create your Oracle database.	“Creating an Oracle Database” on page 96
Verify the Support for Oracle RAC installation and configuration.	“Verifying the Installation and Configuration of Support for Oracle RAC” on page 103

Installing Oracle RAC Software

This section contains the following information:

- [“Installing Binary Files and Configuration Files on a Shared File System” on page 90](#)
- [“Overriding Networking Defaults for Oracle Grid Infrastructure” on page 91](#)
- [“Next Steps” on page 91](#)

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

Installing Binary Files and Configuration Files on a Shared File System

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

- The StorageTek QFS shared file system
- The PxFs-based cluster file system
- The Oracle ACFS file system
- A file system on a qualified NAS device

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

To install Oracle Database 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 91](#).

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.

▼ How to Verify the Installation of Oracle RAC

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

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

For more information, see your Oracle Database 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 Support for 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. Assume the root role 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 Support for 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

Before You Begin Ensure that the following prerequisites are met:

- The Support for Oracle RAC framework is created and is online. See [“Registering and Configuring the Support for Oracle RAC Framework Resource Group”](#) on page 63.
- Storage resources for Oracle Database files are configured. See [“Registering and Configuring Storage Resources for Oracle Database Files”](#) on page 78.
- The Support for Oracle RAC software is installed. See [“Installing Oracle RAC Software”](#) on page 90.
- The installation of the Support for Oracle RAC software is verified. See [“Verifying the Installation of Oracle RAC”](#) on page 91.

Ensure that you have the following information:

- The full path to the Oracle Grid Infrastructure home directory

1. **On one node of the cluster, assume the root role.**
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, Support for Oracle RAC Create Configuration.

The `clsetup` utility prompts you to select the Support for 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 Support for Oracle RAC cluster.

- **If you select the global cluster option, the `clsetup` utility displays the list of components of Support for 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.

The `clsetup` utility displays the list of components of Support for 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.

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.**
- **If the directory is not listed, type the full path to the Oracle Grid Infrastructure home directory.**

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.**

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

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

a. **Type e.**

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.**

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.

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.**

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

b. **At the New Value prompt, type the new name.**

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

14. **Confirm your selection of Oracle Solaris Cluster object names.**

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

15. **To create the configuration, type c.**

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” in *Oracle Database Administrator's Guide* and *Oracle Database SQL Language Reference* (http://docs.oracle.com/cd/E11882_01/server.112/e41084/toc.htm).

You can create the Oracle database to be either administrator managed or policy managed.

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

Note - Starting with Oracle Database version 12c, an Oracle ACFS file system can be used for data files.

- [“How to Specify the Location of Data Files on a Shared File System” on page 96](#)

▼ 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.**

Next Steps After you have created the Oracle database, go to [“Configuring Resources for Support for Oracle RAC Database Instances” on page 97](#).

Configuring Resources for Support for Oracle RAC Database Instances

Resources for Support 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 Support for Oracle RAC database instances:

- [“Tools for Registering and Configuring Resources for Support for Oracle RAC Database Instances” on page 97](#)
- [“How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate” on page 98](#)

Tools for Registering and Configuring Resources for Support for Oracle RAC Database Instances

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

- **The `clsetup` utility.** For more information, see [“How to Enable Oracle Solaris Cluster and Oracle Grid Infrastructure to Interoperate” on page 98](#).
- **Oracle Solaris Cluster maintenance commands.** For more information, see [“Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands” on page 240](#).

The `clsetup` utility provides a wizard for configuring resources for Support 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.

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

▼ 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 Support for 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 Support for Oracle RAC framework resource group is created and is online. See [“Registering and Configuring the Support for Oracle RAC Framework Resource Group” on page 63](#).
- Storage resources for Oracle Database files are configured. See [“Registering and Configuring Storage Resources for Oracle Database Files” on page 78](#).
- The Support for Oracle RAC software is installed. See [“Installing Oracle RAC Software” on page 90](#).
- The installation of the Support for Oracle RAC software is verified. See [“Verifying the Installation of Oracle RAC” on page 91](#).
- The Oracle database is created. See [“Creating an Oracle Database” on page 96](#).
- The Oracle database is running.

To determine whether the Oracle database is running, use the Oracle Database `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 Database home directory
- The names of all Oracle Solaris Cluster storage resources for Oracle files that you are using

1. On one node of the cluster, assume the root role.

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, Support for Oracle RAC Create Configuration.

The `clsetup` utility prompts you to select the Support for 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 Support for Oracle RAC cluster.

- If you select the global cluster option, the `clsetup` utility displays the list of components of Support for 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.

The `clsetup` utility displays the list of components of Support for 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 Support for Oracle RAC databases that are configured on the cluster.

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

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

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

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

- a. **Type e.**

- The `clsetup` utility prompts you for the name of the Support for Oracle RAC database that you are configuring.

- b. **Type the name of the Support for Oracle RAC database that you are configuring.**

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

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

- **If the directory is listed, type the number that corresponds to the directory that you are selecting.**

The `clsetup` utility displays a list of Oracle Database 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 Support for Oracle RAC is to run.

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

- a. **Type e.**

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

- b. **Type the full path to the Oracle Database home directory.**

- 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 Support for Oracle RAC is to run.

12. Type a comma-separated or space-separated list of the numbers that correspond to the storage resources that you are using.

- **If you choose 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.**

13. 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.

14. 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.

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

b. At the New Value prompt, type the new name.

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

15. Confirm your selection of Oracle Solaris Cluster object names.

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

16. To create the configuration, type `c`.

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.

17. Press Return to continue.

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

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

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	Strong dependency on the Support for Oracle RAC framework resource.	Oracle Grid Infrastructure framework resource.
Resource name: crs_framework-rs	Offline-restart dependencies on all scalable device-group resources for Oracle files, if any.	
Resource group: rac-framework-rg		

Resource Type, Resource Name, and Resource Group	Dependencies	Description
	<p>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.</p> <p>If you are not using a shared file system or a volume manager, this resource has no offline-restart dependencies on any other resources.</p>	
Resource type: SUNW. scalable_rac_server_proxy	Strong dependency on the Support for Oracle RAC framework resource.	Proxy resource for the Support for Oracle RAC database server.
Resource name: rac_server_proxy-rs Resource group: rac_server_proxy-rg	Offline-restart dependency on the Oracle Grid Infrastructure framework resource.	

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 Database Files” on page 78](#) 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 Support for 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 103.

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 Support 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(7)` man page.

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

1. Verifying the configuration of the Support for 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 Support for Oracle RAC Framework Resource Group”](#) on page 104 and [“How to Verify the Configuration of the Multiple-Owner Volume-Manager Framework Resource Group”](#) on page 104.
2. Verifying the configuration of storage resources for Oracle Database files. See [“How to Verify the Configuration of Storage Resources for Oracle Database Files”](#) on page 105.
3. Verifying the configuration of resources for Support for Oracle RAC database instances. See [“How to Verify the Configuration of Resources for Support for Oracle RAC Database Instances”](#) on page 107.
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 109.

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

The Support for Oracle RAC framework resource group is created when the task in [“Registering and Configuring the Support for Oracle RAC Framework Resource Group”](#) on page 63 is performed.

1. **Assume the root role on a cluster node or assume a role that provides `solaris.cluster.read` and `solaris.cluster.admin` RBAC authorizations.**
2. **Verify that the Support for Oracle RAC framework resource group is correctly configured.**

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

```
rac-fmk-rg
```

Specifies the name of the Support for Oracle RAC framework resource group.

3. **If the Support for Oracle RAC framework resource group is not already online, bring online the Support for Oracle RAC framework resource group.**

```
# clresourcegroup online -eM rac-fmwk-rg
```

```
rac-fmk-rg
```

Specifies the name of the Support for Oracle RAC framework resource group.

4. **Verify that the Support for 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 Oracle RAC configuration.

1. **Assume the root role 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 -eM vucmm-fmwk-rg
```

```
vucmm-fmk-rg
```

Specifies the name of the multiple-owner volume-manager framework resource group.

4. **Verify that the Oracle RAC resource group is online.**

```
# clresourcegroup status
```

▼ How to Verify the Configuration of Storage Resources for Oracle Database Files

Perform this task only if your configuration of storage management schemes requires storage resources for Oracle Database files.

1. **Assume the root role on a cluster node.**
2. **Verify that all resource groups that contain storage resources for Oracle Database 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 Database files.

3. **Take offline the resource group that contains the resources on which the Oracle 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 Database files.

- If you are using a shared file system for Oracle Database 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 Database 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 Support 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 Support for Oracle RAC database resource group is offline.
- The status of the Support for 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 -eM 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 Support for Oracle RAC are online.

```
# clresourcegroup status
```

Note - Bringing online the resource group that contains the resources on which the Oracle RAC database resource depends also brings online the Oracle RAC database resource group. If you perform this step before the Oracle 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 Oracle RAC database resource group is online on all nodes.

▼ How to Verify the Configuration of Resources for Support for Oracle RAC Database Instances

Resources for Support 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 98](#)
- [“Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands” on page 240](#)

1. Assume the root role 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 Oracle 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 Oracle 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 Oracle RAC database resource. If you perform this step before the Oracle 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 Oracle 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 Support for 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 *Administering an Oracle Solaris Cluster 4.4 Configuration*.

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 *Administering an Oracle Solaris Cluster 4.4 Configuration*.

4. Confirm that the cluster starts correctly.

5. Verify that all Support for Oracle RAC resource groups and their resources are online.

```
# clresourcegroup status
```

Note - The bringing online of all Support for 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.

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” on page 111](#)
- [“Automatically Generated Names for Oracle Solaris Cluster Objects” on page 112](#)
- [“Administering Support for Oracle RAC Databases From Oracle Solaris Cluster Software” on page 112](#)
- [“Effects of State Changes to Oracle Solaris Cluster Resources for Support for Oracle RAC Database Instances” on page 113](#)
- [“Tuning Support for Oracle RAC” on page 115](#)
- [“Tuning the Support for Oracle RAC Fault Monitors” on page 117](#)

Overview of Administration Tasks for Support for Oracle RAC

[Table 14, “Administration Tasks for Support for Oracle RAC,” on page 111](#) summarizes the administration tasks for Support for Oracle RAC.

Perform these tasks whenever they are required.

TABLE 14 Administration Tasks for Support for Oracle RAC

Task	Instructions
Administer Support for Oracle RAC databases from Oracle Solaris Cluster.	“Administering Support for Oracle RAC Databases From Oracle Solaris Cluster Software” on page 112
Tune Support for Oracle RAC extension properties.	“Tuning Support for Oracle RAC” on page 115
Tune Support for Oracle RAC fault monitors.	“Tuning the Support for Oracle RAC Fault Monitors” on page 117

Task	Instructions
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 <i>dg-name</i> -rs, where <i>dg-name</i> is the name of the device group that the resource represents
SUNW.ScalMountPoint	scal- <i>mp-dir</i> -rs, where <i>mp-dir</i> is the mount point of the file system, with / replaced by - asm-mp-rs
SUNW.qfs	qfs- <i>mp-dir</i> -rs, where <i>mp-dir</i> is the mount point of the file system, with / replaced by - asm-home-sqfs-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	<i>lh-name</i> , where <i>lh-name</i> is the logical hostname that you specified when you created the resource

Administering Support for Oracle RAC Databases From Oracle Solaris Cluster Software

Administering Support for 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 Support for Oracle RAC database instances. For information about how to create these resources, see [“Configuring Resources for Support for Oracle RAC Database Instances” on page 97](#).

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 Database `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 node. You can administer the resource groups and resources belonging to a zone cluster from that zone cluster. In the global-cluster 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 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.

For information about the effects on Oracle Database components of changes to the states of Oracle Solaris Cluster resources, see [“Effects of State Changes to Oracle Solaris Cluster Resources for Support for Oracle RAC Database Instances” on page 113](#).

Effects of State Changes to Oracle Solaris Cluster Resources for Support 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 Support for 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 15 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 16 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 *Concepts for Oracle Solaris Cluster 4.4*](#).

For detailed information about the state of Oracle Grid Infrastructure resources, see your Oracle Grid Infrastructure 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#) to configure the resources later.

This section contains the following information about tuning the Support for Oracle RAC data service:

- [“Guidelines for Setting Timeouts” on page 116](#)
- [“Reservation Step Timeout” on page 116](#)
- [“Oracle Clusterware Framework Resource Method Timeouts” on page 117](#)

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.

If timeouts occur during reconfiguration processes, increase the values of the appropriate timeout property to accommodate your cluster configuration.

Reservation Step Timeout

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 202](#).

EXAMPLE 4 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 Support for Oracle RAC framework

component is represented by an instance of the `SUNW.rac_framework` resource type that is named `rac-framework-rs`.

Oracle Clusterware Framework Resource Method Timeouts

The start and stop timing of the Oracle Clusterware framework resource `SUNW.crs_framework` is dependent on a number of factors, including things such as the following:

- The number of shared physical devices
- The Oracle ASM storage configuration
- The number of Oracle Clusterware resources
- The use of an Administrative Database and its configuration

If you see that the Oracle Clusterware framework resource start or stop method execution time is getting close to the timeout value, or if an actual timeout occurs, increase the timeout for that method.

EXAMPLE 5 Setting the Oracle Clusterware Framework Resource Start Timeout

```
# clresource set -p Start_timeout=600 crs-framework-rs
```

This example sets the timeout for the start method of the `crs-framework-rs` resource to 600 seconds.

If a stop method timeout occurs, you first need to clear the `STOP_FAILED` condition before you start the resource with the new timeout value. See the [clresource\(8CL\)](#) man page for instructions on clearing this condition.

If a start method timeout occurs, first take the resource offline to clear this condition, then start the resource with the new timeout value.

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 17 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

Standard 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*](#). 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 118](#)
- [“Operation of the Fault Monitor for Scalable File-System Mount Points” on page 119](#)
- [“Obtaining Core Files for Troubleshooting DBMS Timeouts” on page 120](#)

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 cluster node.
 - The resource is placed into the disabled state on the current cluster node, causing the file system to be unmounted from that node.
- If the file system is a StorageTek QFS shared file system, the response is as follows:
 - If the cluster node on which the error occurred is hosting the metadata server resource, the metadata server resource is failed over to another node.
 - The file system is unmounted.

If the failover attempt fails, the file system remains unmounted and a warning is given.

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.

```
# coreadm -g /var/cores/%f.%n.%p.core -e global -e process \  
-e global-setid -e proc-setid -e log
```

For more information, see the [coreadm \(1M\)](#) man page.

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” on page 121
- “Sources of Diagnostic Information” on page 132
- “Common Problems and Their Solutions” on page 132

Verifying the Status of Support for Oracle RAC

The status of resource groups and resources for Support for Oracle RAC indicates the status of Support for 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(8CL)` command.
- To obtain status information for resources, use the `clresource(8CL)` command.

▼ How to Verify the Status of Support for Oracle RAC

1. **Assume the root role 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\(8CL\)](#)
- [clresourcegroup\(8CL\)](#)

Examples of the Status of Support for Oracle RAC

The following examples show the status of resource groups and resources for a configuration of Support for Oracle RAC on a four-node cluster. Each node is a machine that uses the SPARC processor.

The configuration in this example uses a StorageTek QFS shared file system on Solaris Volume Manager for Sun Cluster to store Oracle files. The configuration includes a multiple-owner volume-manager framework resource group to contain the volume manager resource.

The resource groups and resources for this configuration are shown in the following table.

Resource Group	Purpose	Resource Type	Resource Instance Type
rac-framework-rg	Support for Oracle RAC framework resource group	SUNW.rac_framework SUNW.crs_framework	rac-framework-rs crs_framework-rs
vucmm-framework-rg	multiple-owner volume-manager framework resource group	SUNW.vucmm_framework SUNW.vucmm_svm	vucmm-framework-rs vucmm-svm-rs
sca1dg-rg	Resource group for scalable device-group resources	SUNW.ScalDeviceGroup	sca1oradg-rs

Resource Group	Purpose	Resource Type	Resource Instance Type
qfsmnds-rg	Resource group for StorageTek QFS metadata server resources	SUNW.qfs	qfs-db_qfs-OraHome-rs qfs-db_qfs-OraData-rs
scalmnt-rg	Resource group for scalable file-system mount-point resources	SUNW.ScalMountPoint	scal-db_qfs-OraHome-rs scal-db_qfs-OraData-rs
rac_server_proxy-rg	Support for Oracle RAC database resource group	SUNW. scalable_rac_server_proxy	rac_server_proxy-rs
rac_server_proxy-rg	Support for Oracle RAC database resource group	SUNW. scalable_rac_server_proxy	rac_server_proxy-rs

EXAMPLE 6 Status of a Faulty Support for Oracle RAC Framework Resource Group

This example provides the following status information for a Support for Oracle RAC framework resource group that is faulty.

- A reconfiguration error has prevented the `rac_framework` resource from starting on cluster node `pclus1`.
- The effects of this reconfiguration error on resource groups and other resources are as follows:
 - The `rac-framework-rg` resource group is offline and in the start failed state on cluster node `pclus1`.
 - The `rac_svm` resource and the `crs_framework` resource are offline on cluster node `pclus1`.
 - All other multi-master resource groups and the resources that the groups contain are offline on cluster node `pclus1`.
 - All failover resource groups have failed over from cluster node `pclus1` to a secondary node.
- All multi-master resource groups and the resources that the groups contain are online on the remaining nodes.

```
# clresourcegroup status +
```

```
=== Cluster Resource Groups ===
```

```
Group Name      Node Name      Suspended      Status
-----
rac-framework-rg  pclus1        No             Online faulted
                  pclus2        No             Online
                  pclus3        No             Online
                  pclus4        No             Online
```

```

vucmm-framework-rg    pclus1    No    Online
                      pclus2    No    Online
                      pclus3    No    Online
                      pclus4    No    Online

scaldg-rg             pclus1    No    Online
                      pclus2    No    Online
                      pclus3    No    Online
                      pclus4    No    Online

qfsmds-rg            pclus1    No    Offline
                      pclus2    No    Online
                      pclus3    No    Offline
                      pclus4    No    Offline

scalmnt-rg           pclus1    No    Online
                      pclus2    No    Online
                      pclus3    No    Online
                      pclus4    No    Online

rac_server_proxy-rg  pclus1    No    Pending online blocked
                      pclus2    No    Online
                      pclus3    No    Online
                      pclus4    No    Online

```

clresource status -g rac-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
rac-framework-rs	pclus1	Start failed	Faulted - Error in previous reconfiguration.
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
crs_framework-rs	pclus1	Offline	Offline
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

clresource status -g vucmm-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----

```

vucmm-framework-rs pclus1      Online      Online
                   pclus2      Online      Online
                   pclus3      Online      Online
                   pclus4      Online      Online

vucmm-svm-rs       pclus1      Offline     Offline
                   pclus2      Online      Online
                   pclus3      Online      Online
                   pclus4      Online      Online

```

clresource status -g scaldg-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
scalordg-rs	pclus1	Online	Online - Diskgroup online
	pclus2	Online	Online - Diskgroup online
	pclus3	Online	Online - Diskgroup online
	pclus4	Online	Online - Diskgroup online

clresource status -g qfsmds-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
qfs-db_qfs-OraHome-rs	pclus1	Offline	Offline
	pclus2	Online	Online - Service is online.
	pclus3	Offline	Offline
	pclus4	Offline	Offline
qfs-db_qfs-OraData-rs	pclus1	Offline	Offline
	pclus2	Online	Online - Service is online.
	pclus3	Offline	Offline
	pclus4	Offline	Offline

clresource status -g scalmnt-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
scal-db_qfs-OraHome-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

```
scal-db_qfs-OraData-rs  pclus1    Online   Online
                        pclus2    Online   Online
                        pclus3    Online   Online
                        pclus4    Online   Online
```

```
# clresource status -g rac_server_proxy-rg +
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	State	Status Message
rac_server_proxy-rs	pclus1	Offline	Offline
	pclus2	Online	Online - Oracle instance UP
	pclus3	Online	Online - Oracle instance UP
	pclus4	Online	Online - Oracle instance UP

EXAMPLE 7 Status of a Faulty Support for Oracle RAC Database Resource Group

This example provides the following status information for a Support for Oracle RAC database resource group that is faulty:

- The Support for Oracle RAC database on pclus1 has failed to start. The effects of this failure are as follows:
 - The rac_server_proxy-rg resource group is online, but faulted on node pclus1.
 - The rac_server_proxy-rs resource is offline on node pclus1.
- All other multi-master resource groups and the resources that the groups contain are online on all nodes.
- All failover resource groups and the resources that the groups contain are online on their primary nodes and offline on the remaining nodes.

```
# clresourcegroup status +
```

```
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
rac-framework-rg	pclus1	No	Online
	pclus2	No	Online
	pclus3	No	Online
	pclus4	No	Online
vucmm-framework-rg	pclus1	No	Online
	pclus2	No	Online
	pclus3	No	Online
	pclus4	No	Online

```

scaldg-rg          pclus1    No      Online
                   pclus2    No      Online
                   pclus3    No      Online
                   pclus4    No      Online

qfsmds-rg         pclus1    No      Online
                   pclus2    No      Offline
                   pclus3    No      Offline
                   pclus4    No      Offline

scalmnt-rg        pclus1    No      Online
                   pclus2    No      Online
                   pclus3    No      Online
                   pclus4    No      Online

rac_server_proxy-rg pclus1    No      Online faulted
                   pclus2    No      Online
                   pclus3    No      Online
                   pclus4    No      Online

```

clresource status -g rac_server_proxy-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
rac_server_proxy-rs	pclus1	Offline	Offline - Oracle instance DOWN
	pclus2	Online	Online - Oracle instance UP
	pclus3	Online	Online - Oracle instance UP
	pclus4	Online	Online - Oracle instance UP

clresource status -g rac-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
rac-framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
crs_framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

clresource status -g vucmm-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----
vucmm-framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
vucmm-svm-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

clresource status -g scaldg-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----
scalordg-rs	pclus1	Online	Online - Diskgroup online
	pclus2	Online	Online - Diskgroup online
	pclus3	Online	Online - Diskgroup online
	pclus4	Online	Online - Diskgroup online

clresource status -g qfsmds-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----
qfs-db_qfs-OraHome-rs	pclus1	Online	Online - Service is online.
	pclus2	Offline	Offline
	pclus3	Offline	Offline
	pclus4	Offline	Offline
qfs-db_qfs-OraData-rs	pclus1	Online	Online - Service is online.
	pclus2	Offline	Offline
	pclus3	Offline	Offline
	pclus4	Offline	Offline

clresource status -g scalmnt-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----


```

scal-db_qfs-OraHome-rs  pclus1    Online  Online
                        pclus2    Online  Online
                        pclus3    Online  Online
                        pclus4    Online  Online

scal-db_qfs-OraData-rs pclus1    Online  Online
                        pclus2    Online  Online
                        pclus3    Online  Online
                        pclus4    Online  Online

```

EXAMPLE 8 Status of an Operational Support for Oracle RAC Configuration

This example shows the status of a Support for Oracle RAC configuration that is operating correctly. The example indicates that the status of resource groups and resources in this configuration is as follows:

- All multi-master resource groups and the resources that the groups contain are online on all nodes.
- All failover resource groups and the resources that the groups contain are online on their primary nodes and offline on the remaining nodes.

```
# clresourcegroup status +
```

```
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
rac-framework-rg	pclus1	No	Online
	pclus2	No	Online
	pclus3	No	Online
	pclus4	No	Online
vucmm-framework-rg	pclus1	No	Online
	pclus2	No	Online
	pclus3	No	Online
	pclus4	No	Online
scal-dg-rg	pclus1	No	Online
	pclus2	No	Online
	pclus3	No	Online
	pclus4	No	Online
qfsmds-rg	pclus1	No	Online
	pclus2	No	Offline
	pclus3	No	Offline
	pclus4	No	Offline

```

scalmnt-rg      pclus1      No      Online
                pclus2      No      Online
                pclus3      No      Online
                pclus4      No      Online

rac_server_proxy-rg  pclus1      No      Online
                    pclus2      No      Online
                    pclus3      No      Online
                    pclus4      No      Online
    
```

clresource status -g rac-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
rac-framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
crs_framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

clresource status -g vucmm-framework-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
vucmm-framework-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
vucmm-svm-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

clresource status -g scaldg-rg +

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
---------------	-----------	-------	----------------

```

scaloradg-rs      pclus1      Online      Online - Diskgroup online
                  pclus2      Online      Online - Diskgroup online
                  pclus3      Online      Online - Diskgroup online
                  pclus4      Online      Online - Diskgroup online

```

```
# clresource status -g qfsmnds-rg +
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	State	Status Message
qfs-db_qfs-OraHome-rs	pclus1	Online	Online - Service is online.
	pclus2	Offline	Offline
	pclus3	Offline	Offline
	pclus4	Offline	Offline
qfs-db_qfs-OraData-rs	pclus1	Online	Online - Service is online.
	pclus2	Offline	Offline
	pclus3	Offline	Offline
	pclus4	Offline	Offline

```
# clresource status -g scalmnt-rg +
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	State	Status Message
scal-db_qfs-OraHome-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online
scal-db_qfs-OraData-rs	pclus1	Online	Online
	pclus2	Online	Online
	pclus3	Online	Online
	pclus4	Online	Online

```
# clresource status -g rac_server_proxy-rg +
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	State	Status Message
rac_server_proxy-rs	pclus1	Online	Online - Oracle instance UP
	pclus2	Online	Online - Oracle instance UP
	pclus3	Online	Online - Oracle instance UP
	pclus4	Online	Online - Oracle instance UP

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	<code>/var/cluster/vucmm/vucmm_reconf.log.0 (0,1,...)</code>
Log file for the current userland cluster membership monitor (UCMM) reconfiguration	<code>/var/cluster/ucmm/ucmm_reconf.log</code>
Log files for previous UCMM reconfigurations	<code>/var/cluster/ucmm/ucmm_reconf.log.0 (0,1,...)</code>

The directory `/var/opt/SUNWscor/oracle_server/proxyresource` contains log files for the resource that represents the Support for 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 Support for 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 a Support for Oracle RAC Framework Resource Group” on page 133](#)
- [“Failure of a Multiple-Owner Volume-Manager Framework Resource Group” on page 136](#)

- [“SUNW.qfs Registration Fails Because the Registration File Is Not Found” on page 138](#)
- [“Node Panic Caused by a Timeout” on page 138](#)
- [“Failure of Services that Depend on a SUNW.ScalDeviceGroup Resource” on page 139](#)
- [“Failure of a SUNW.rac_framework or SUNW.vucmm_framework Resource to Start” on page 139](#)
- [“SUNW.rac_framework Failure-to-Start Status Messages” on page 140](#)
- [“SUNW.vucmm_framework Failure-to-Start Status Messages” on page 140](#)
- [“How to Recover From the Timing Out of the START Method” on page 141](#)
- [“Failure of a Resource to Stop” on page 142](#)

Failure of a Support for Oracle RAC Framework Resource Group

This section describes problems that can affect the Support for Oracle RAC framework resource group.

- [“Node Panic During Initialization of Support for Oracle RAC” on page 133](#)
- [“Failure of the ucmm Daemon to Start” on page 134](#)
- [“How to Recover From a Failure of the ucmm Daemon or a Related Component” on page 134](#)

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 "ucmm" 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 138](#).

Solution: For instructions to correct the problem, see [“How to Recover From a Failure of the ucmm Daemon or a Related Component” on page 134.](#)

Note - When the node is a global-cluster 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 ucmm Daemon to Start

The UCMM daemon, ucmm, 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 ucmm 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 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 ucmm Daemon or a Related Component” on page 134.](#)

▼ How to Recover From a Failure of the ucmm 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 133](#)
- [“Failure of the ucmm Daemon to Start” on page 134](#)

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 132.](#)

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 138](#).

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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

4. On the node where the problem occurred, take offline and bring online the Support for Oracle RAC framework resource group.

This step refreshes the resource group with the configuration changes you made.

- a. **Assume the root role or assume a role that provides `solaris.cluster.admin` RBAC authorization.**
- b. **Type the command to take offline the Support for 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 Support for Oracle RAC framework resource group and its resources.**

```
# clresourcegroup online -eM -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 136](#)
- [“Failure of the vucmmd Daemon to Start” on page 136](#)
- [“How to Recover From a Failure of the vucmmd Daemon or a Related Component” on page 137](#)

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 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 137](#).

▼ 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 136](#)
- [“Failure of the vucmmd Daemon to Start” on page 136](#)

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 132](#).

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 *Administering an Oracle Solaris Cluster 4.4 Configuration*](#).

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. Assume the root role 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 -eM -n node vucmm-fmwk-rg
```

SUNW.qfs Registration Fails Because the Registration File Is Not Found

Oracle Solaris Cluster resource-type registration files are located in the `/opt/cluster/lib/rgm/rtreg/` or `/usr/cluster/lib/rgm/rtreg/` directory. The `SUNW.qfs` resource-type registration file is located in the `/opt/SUNWsamfs/sc/etc/` directory.

If Oracle Solaris Cluster software is already installed when you install StorageTek QFS software, the necessary mapping to the `SUNW.qfs` registration file is automatically created. But if Oracle Solaris Cluster software is not already installed when you install StorageTek QFS software, the necessary mapping to the `SUNW.qfs` registration file is not made, even when Sun Cluster software is later installed. Attempts to register the `SUNW.qfs` resource type therefore fail because the Oracle Solaris Cluster software is unaware of the location of its registration file.

To enable Oracle Solaris Cluster software to locate the `SUNW.qfs` resource type, create a symbolic link to the directory:

```
# cd /usr/cluster/lib/rgm/rtreg
# ln -s /opt/SUNWsamfs/sc/etc/SUNW.qfs SUNW.qfs
```

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 116](#).

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 Support for Oracle RAC framework resource group on the node that panicked.

Failure of Services that Depend on a SUNW.ScalDeviceGroup Resource

In a configuration that uses Oracle ASM with Solaris Volume Manager mirrored logical volumes, if availability is lost for a service that depends on a SUNW.ScalDeviceGroup resource, the problem might be a failure of the SUNW.ScalDeviceGroup probe. This problem can be mitigated by increasing the value of the IOTimeout property of the SUNW.ScalDeviceGroup resource type.

```
# clresource set -p IOTimeout=timeout-value-in-seconds resource-name
```

The default of the IOTimeout property is 30 seconds. This value can be set much higher and should be determined for each system individually. Start by increasing the value to 240 seconds, then tune the property further as needed to achieve optimal results.

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 121.

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 - ucmmmd is not running

Description: The ucmmmd daemon is not running on the node where the resource resides.

Solution: For information about how to correct this problem, see [“Failure of the ucmmmd Daemon to Start” on page 134](#).

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 ucmmmd Daemon or a Related Component” on page 134](#).
- If the message does not indicate a problem, no action is required.

Online

Description: Reconfiguration of Support for 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 141](#).

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 136](#).

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 137](#).
- If the message does not indicate a problem, no action is required.

Online

Description: Reconfiguration of Support for Oracle RAC was not completed until after the START method of the SUNW.vucmm_f 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 141](#).

▼ How to Recover From the Timing Out of the START Method

1. **Assume the root role 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 a Support for 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 Support for Oracle RAC framework resource group.

If the Support for Oracle RAC framework resource group was created by using the `clsetup` utility, the name of the resource group is `rac-f 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 -eM 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 *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.

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” on page 143](#)
- [“Modifying Online the Resource for a Scalable Device Group” on page 144](#)
- [“Extending an Existing Configuration of Support for Oracle RAC” on page 144](#)
- [“Removing an Oracle Grid Infrastructure Resource” on page 153](#)
- [“Removing Support for Oracle RAC” on page 155](#)

Overview of Tasks for Modifying an Existing Configuration of Support for Oracle RAC

Table 18, [“Tasks for Modifying an Existing Configuration of Support for Oracle RAC,” on page 143](#) summarizes the administration tasks for Support for Oracle RAC.

Perform these tasks whenever they are required.

TABLE 18 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 144
Extend an existing configuration of Support for Oracle RAC.	“Extending an Existing Configuration of Support for Oracle RAC” on page 144
Remove Support for Oracle RAC.	“Removing Support for Oracle RAC” on page 155

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. Assume the `root` role 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 145.
- You are adding a volume manager. See [“How to Add a Volume Manager Resource to the `SUNW.vucmm_framework` Resource Group”](#) on page 151.

▼ 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 Support for Oracle RAC framework resource group
- Resource groups for scalable device group resources
- Resource groups that contain resources for the StorageTek QFS metadata server
- Resource groups for logical hostname resources
- The resource group for the Support for Oracle RAC database

Note - During this procedure, some resources might report that they cannot come online, most likely because of missing software that cannot yet be installed. These failures are expected and are safe to ignore until all procedure steps are completed.

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 39](#).
- Ensure that the node that you add is connected to the shared storage that is used by the Support for Oracle RAC configuration.

1. **Assume the root role 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 Support for 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 Database files.

If you are not using any scalable device groups for Oracle Database 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. Add the nodes to any resource groups that contain resources for the StorageTek QFS metadata server.

If no resource groups that contain resources for the StorageTek QFS metadata server are configured, omit this step.

For each resource group to which you are adding nodes, run the following command:

```
# clresourcegroup add-node -n nodelist qfs-mds-rg
```

-n nodelist

Specifies a comma-separated list of cluster nodes to which you are adding Support for Oracle RAC.

qfs-mds-rg

Specifies the name of the resource group to which you are adding nodes.

9. 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 -eM scal-dg-rg
```

scal-dg-rg

Specifies the name of the resource group that you are bringing online.

10. Start Oracle Grid Infrastructure.

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

Startup will be queued to init within 30 seconds.

11. Add the nodes to the resource group for each [Support for 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.

12. 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 243.

13. Modify each resource for [Support 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 Support for Oracle RAC database resource that you are disabling.

b. Enable the resource.

```
# clresource enable rac-db-rs
```

rac-db-rs

Specifies the name of the Support for Oracle RAC database resource that you are enabling.

14. Bring online each resource group for Support for Oracle RAC databases.

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

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

rac-db-rg

Specifies the name of the resource group that you are bringing online.

Example 9 Adding Support for Oracle RAC to Selected Nodes

This example shows the sequence of operations that is required to add Support for Oracle RAC to nodes `pc1us3` and `pc1us4` of a four-node cluster.

The configuration of Support for Oracle RAC in this example is as follows:

- StorageTek QFS shared file system on Solaris Volume Manager for Sun Cluster is used to store Oracle Database files.
- The mount points of file systems that are used for Oracle Database files are as follows:
 - Oracle Database files: `/db_qfs/OraData`
 - Oracle Database binary files and related files: `/db_qfs/OraHome`
- The `oradg` disk set is used only by the Support for Oracle RAC database.
- The name of the Support for Oracle RAC database is `swb`.
- The StorageTek QFS shared file systems use a Solaris Volume Manager for Sun Cluster multi-owner disk set that is named `oradg`. The creation of this disk set is shown in [Example](#)

3, “Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster,” on page 76.

- The configuration uses a multiple-owner volume-manager framework resource group.

The configuration of resource groups in this example is shown in the following table.

Resource Group	Purpose
vucmm-framework-rg	Multiple-owner volume-manager resource group.
rac-framework-rg	Support for Oracle RAC framework resource group.
scaldg-rg	Resource group for scalable device-group resources.
qfsmds-rg	Resource group for StorageTek QFS metadata server resources.
scalmnt-rg	Resource group for scalable file-system mount-point resources.
rac_server_proxy-rg	Support for Oracle RAC database resource group.

The resource groups that are required for this configuration are shown in [Figure 2, “Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System,”](#) on page 181.

1. To add the nodes to the resource group that contains scalable file-system mount-point resources, the following command is run:

```
# clresourcegroup add-node -S -n pclus3,pclus4 scalmnt-rg
```

2. To add the nodes to the multiple-owner volume-manager framework resource group, the following command is run:

```
# clresourcegroup add-node -S -n pclus3,pclus4 vucmm-framework-rg
```

3. To add the nodes to the Support for Oracle RAC framework resource group, the following command is run:

```
# clresourcegroup add-node -S -n pclus3,pclus4 rac-framework-rg
```

4. To add the nodes to the Solaris Volume Manager for Sun Cluster multi-owner disk set oradg, the following command is run:

```
# metaset -s oradg -M -a -h pclus3 pclus4
```

5. To add the nodes to the resource group that contains scalable device group resources, the following command is run:

```
# clresourcegroup add-node -S -n pclus3,pclus4 scaldg-rg
```

6. To mount the shared file systems that are to be accessed from the nodes that are being added, the following commands are run:

```
# mount /db_qfs/OraData
```

```
# mount /db_qfs/OraHome
```

7. To add the nodes to the resource group that contains resources for the StorageTek QFS metadata server, the following command is run:

```
# clresourcegroup add-node -n pclus3,pclus4 qfsmds-rg
```

8. To bring online the resource group that contains scalable device group resources, the following command is run:

```
# clresourcegroup online -eM scaldg-rg
```

9. To start Oracle Clusterware and to verify the correct startup of Oracle Clusterware, the following commands are run:

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

Startup will be queued to init within 30 seconds.

```
# /db_qfs/OraHome/crs/bin/crsctl check crs
```

CSS appears healthy

CRS appears healthy

EVM appears healthy

10. To add the nodes to the resource group for the Support for Oracle RAC database, the following command is run:

```
# clresourcegroup add-node -S -n pclus3,pclus4 rac_server_proxy-rg
```

After the addition of the nodes to the resource group for the Support for Oracle RAC database, the required Oracle Clusterware resources are created. The creation of these Oracle Clusterware resources is beyond the scope of this example.

11. To bring online the resource group for the Support for Oracle RAC database, the following command is run:

```
# clresourcegroup online -eM rac_server_proxy-rg
```

▼ 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 Support for 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 Support for Oracle RAC is to run.

1. **Assume the root role 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 -M 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 71.

Removing an Oracle Grid Infrastructure Resource

This section contains the following procedures to remove an Oracle Grid Infrastructure resource:

- [“How to Remove a Dependency”](#) on page 153
- [“How to Delete the sun.resource Resource”](#) on page 154

▼ How to Remove a Dependency

This procedure shows how to set the offline-restart dependency to remove a dependency.

1. **Assume the root role.**
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 Support for 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. **Assume the root role.**
2. **Ensure that the dependency is removed, as described in [“How to Remove a Dependency” on page 153](#), 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 155.
- **Selected nodes in a cluster.** See [“How to Remove Support for Oracle RAC From Selected Nodes”](#) on page 162.

▼ 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 Support for Oracle RAC databases are running, perform this task to remove a Support for Oracle RAC database from the cluster. The remaining Support for 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 Support for Oracle RAC database
- Resource groups for logical hostname resources
- Resource groups for scalable file-system mount-point resources
- Resource groups that contain resources for the StorageTek QFS metadata server
- Resource groups for scalable device group resources
- The Support for Oracle RAC framework resource group
- The multiple-owner volume-manager framework resource group, if used



Caution - You might perform this task to remove a Support for Oracle RAC database from a cluster where multiple Support for Oracle RAC databases are running. In this situation, do *not* remove any resource group on whose resources the remaining Support for 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 Support for 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, assume the root role.**
2. **Remove the resource group for each Support for Oracle RAC database that you are removing.**

For each Support for 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 Database utilities to remove from the cluster each Support for Oracle RAC database that you no longer require.**
4. **If you are removing Support for Oracle RAC entirely, use Oracle Database utilities to remove the following items from all nodes in a cluster:**
 - Support for Oracle RAC software
 - 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 resources for the StorageTek QFS metadata server.

If no resource groups that contain resources for the StorageTek QFS metadata server are configured, omit this step.

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

```
# clresourcegroup delete -F qfs-mds-rg
```

```
qfs-mds-rg
```

Specifies the resource group that you are removing.

8. Remove the StorageTek QFS shared file systems that were represented by resources in the resource group that you deleted in [Step 7](#).

9. 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.

10. Destroy any scalable device groups that were affected by the removal of resource groups in [Step 9](#).

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/dN`, 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 a Support for Oracle RAC database from a cluster where multiple Support for Oracle RAC databases are running, omit the remaining steps in this procedure.

11. Remove the Support for Oracle RAC framework resource group.

```
# clresourcegroup delete -F rac-fmwk-rg
```

rac-fwk-rg

Specifies the resource group that you are removing.

12. Remove the multiple-owner volume-manager framework resource group, if used.

```
# clresourcegroup delete -F vucmm-fwk-rg
```

vucmm-fwk-rg

Specifies the resource group that you are removing.

13. 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 112](#).

14. (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
```

15. Reboot each node in the cluster.

Example 10 Removing Support for Oracle RAC From a Cluster

This example shows the sequence of operations that is required to remove Support for Oracle RAC from all nodes of a four-node cluster. The nodes in this cluster are named `pclus1`, `pclus2`, `pclus3`, and `pclus4`. Only one Support for Oracle RAC database is configured on the cluster.

The configuration of Support for Oracle RAC in this example is as follows:

- StorageTek QFS shared file system on Solaris Volume Manager for Sun Cluster is used to store Oracle Database files.
- The mount points of file systems that are used for Oracle Database files are as follows:
 - Oracle Database files: `/db_qfs/OraData`
 - Oracle Database binary files and related files: `/db_qfs/OraHome`
- The `oradg` disk set is used only by the Support for Oracle RAC database.
- The name of the Support for Oracle RAC database is `swb`.

- The StorageTek QFS shared file systems use a Solaris Volume Manager for Sun Cluster multi-owner disk set that is named `oradg`. The creation of this disk set is shown in [Example 3, “Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster,”](#) on page 76.
- The configuration uses a multiple-owner volume-manager framework resource group.

The resource groups that are required for this configuration are shown in [Figure 2, “Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System,”](#) on page 181.

Resource Group	Purpose
<code>vucmm-framework-rg</code>	Multiple-owner volume-manager resource group.
<code>rac-framework-rg</code>	Support for Oracle RAC framework resource group.
<code>scaldg-rg</code>	Resource group for scalable device-group resources.
<code>qfsmds-rg</code>	Resource group for StorageTek QFS metadata server resources.
<code>scalmnt-rg</code>	Resource group for scalable file-system mount-point resources.
<code>rac_server_proxy-rg</code>	Support for Oracle RAC database resource group.

The resource groups that are required for this configuration are shown in [Figure 2, “Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System,”](#) on page 181.

1. To remove the resource group for the Support for Oracle RAC database, the following command is run:

```
# clresourcegroup delete -F rac_server_proxy-rg
```

After the removal of this resource group, Oracle utilities are used to remove the following items:

- The Support for Oracle RAC database
- The Support for Oracle RAC software
- The Oracle Clusterware software

The removal of these items is beyond the scope of this example.

2. To disable the Oracle Clusterware framework resource, the following command is run:

```
# clresource disable crs_framework-rs
```

3. To remove the resource group that contains scalable file-system mount-point resources, the following command is run:

```
# clresourcegroup delete -F scalmnt-rg
```


4. To remove the resource group that contains resources for the StorageTek QFS metadata server, the following command is run:

```
# clresourcegroup delete -F qfsmds-rg
```

After the removal of this resource group, StorageTek QFS utilities are used to remove the StorageTek QFS shared file systems that are used for Oracle Database files. The removal of these file systems is beyond the scope of this example.

5. To remove the resource group that contains scalable device group resources, the following command is run:

```
# clresourcegroup delete -F scaldg-rg
```

6. To destroy the Solaris Volume Manager for Sun Cluster multi-owner disk set oradg, the following commands are run:

```
# metaclear -s oradg -a
# metaset -s oradg -d \
-f /dev/did/dsk/d8 /dev/did/dsk/d9 /dev/did/dsk/d15 /dev/did/dsk/d16
# metaset -s oradg -d -h pclus1 pclus2 pclus3 pclus4
```

The following global devices are removed from the disk set:

- /dev/did/dsk/d8
- /dev/did/dsk/d9
- /dev/did/dsk/d15
- /dev/did/dsk/d16

7. To remove the Support for Oracle RAC framework resource group, the following command is run:

```
# clresourcegroup delete -F rac-framework-rg
```

8. To remove the multiple-owner volume-manager framework resource group, the following command is run:

```
# clresourcegroup delete -F vucmm-framework-rg
```

9. To unregister the resource type of each resource that was removed, the following command is run:

```
# clresourcetype unregister \
SUNW.scalable_rac_server_proxy,\
SUNW.ScalMountPoint,\
SUNW.qfs,\
SUNW.ScalDeviceGroup,\
SUNW.rac_svm,\
SUNW.crs_framework,\
```

`SUNW.rac_framework`

After the unregistration of these resource types, the following operations are performed:

- The removal of the Support for Oracle RAC software packages
- The reboot of each node in the cluster

These operations are beyond the scope of this example.

▼ 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 Support for Oracle RAC databases are running, perform this task to remove a Support for Oracle RAC database from selected nodes. The Support for Oracle RAC database that you remove continues to run on the other cluster nodes. The remaining Support for 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 Support for Oracle RAC database
- Resource groups for logical hostname resources
- Resource groups for scalable file-system mount-point resources
- Resource groups that contain resources for the StorageTek QFS metadata server
- Resource groups for scalable device group resources
- The Support for Oracle RAC framework resource group



Caution - You might perform this task to remove a Support for Oracle RAC database from selected nodes of a cluster where multiple Support for Oracle RAC databases are running. In this situation, do *not* remove the nodes from any resource group on whose resources the remaining Support for 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 Support for Oracle RAC framework resource group, do *not* remove the nodes from this resource group.

1. **Assume the root role.**

2. Remove the nodes from the resource group for each Support for Oracle RAC database that you are removing.

For each Support for Oracle RAC database that you are removing, perform the following steps:

a. Take offline the resource group for the Support for 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 Support for 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 Database 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 Support for Oracle RAC database
- Oracle Grid Infrastructure

6. Switch any resource groups that contain resources for the StorageTek QFS metadata server to a node from which you are *not* removing Support for Oracle RAC.

If no configured resource groups contain resources for the StorageTek QFS metadata server, omit this step.

```
# clresourcegroup switch -n node-to-stay qfs-mds-rg
```

node-to-stay Specifies the node to which you are switching the resource group. This node must be a node from which you are *not* removing Support for Oracle RAC.

qfs-mds-rg Specifies the name of the resource group that you are switching to another node.

7. 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.

8. **Remove the nodes from the node list of any resource groups that contain resources for the StorageTek QFS metadata server.**

If no resource groups that contain resources for the StorageTek QFS metadata server are configured, omit this step.

The resource groups to modify are the resource groups that you switched to another node in [Step 6](#).

```
# clresourcegroup remove-node -n nodelist qfs-mds-rg
```

-n nodelist

Specifies a comma-separated list of cluster nodes that you are removing from the resource group.

qfs-mds-rg

Specifies the name of the resource group from which you are removing nodes.

9. **Remove the configuration of the StorageTek QFS shared file systems from the nodes.**

10. 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.

11. Remove the nodes from any scalable device groups that were affected by the removal of nodes from resource groups in [Step 10](#).

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 a Support for Oracle RAC database from selected nodes of a cluster where multiple Support for Oracle RAC databases are running, omit the remaining steps in this procedure.

12. Remove the nodes from the Support for 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.

13. 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.

14. (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
```

15. Reboot each node from which you removed Support for Oracle RAC.

Example 11 Removing Support for Oracle RAC From Selected Nodes

This example shows the sequence of operations that is required to remove Support for Oracle RAC from nodes `pc1us3` and `pc1us4` of a four-node cluster.

The configuration of Support for Oracle RAC in this example is as follows:

- StorageTek QFS shared file system on Solaris Volume Manager for Sun Cluster is used to store Oracle Database files.
- The mount points of file systems that are used for Oracle Database files are as follows:
 - Oracle Database files: `/db_qfs/OraData`
 - Oracle Database binary files and related files: `/db_qfs/OraHome`
- The `oradg` disk set is used only by the Support for Oracle RAC database.
- The name of the Support for Oracle RAC database is `swb`.
- The StorageTek QFS shared file systems use a Solaris Volume Manager for Sun Cluster multi-owner disk set that is named `oradg`. The creation of this disk set is shown in [Example](#)

3, “Creating a Multi-Owner Disk Set in Solaris Volume Manager for Sun Cluster,” on page 76.

- The configuration uses a multiple-owner volume-manager framework resource group.

The resource groups that are required for this configuration are shown in [Figure 2](#), “Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System,” on page 181.

Resource Group	Purpose
vucmm-framework-rg	Multiple-owner volume-manager resource group.
rac-framework-rg	Support for Oracle RAC framework resource group.
scaldg-rg	Resource group for scalable device-group resources.
qfsmds-rg	Resource group for StorageTek QFS metadata server resources.
scalmnt-rg	Resource group for scalable file-system mount-point resources.
rac_server_proxy-rg	Support for Oracle RAC database resource group.

The resource groups that are required for this configuration are shown in [Figure 2](#), “Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System,” on page 181.

1. To remove nodes pclus3 and pclus4 from the resource group for the Support for Oracle RAC database, the following commands are run:

```
# clresourcegroup offline -n pclus3,pclus4 rac_server_proxy-rg
# clresourcegroup remove-node -n pclus3,pclus4 rac_server_proxy-rg
```

2. To remove nodes pclus3 and pclus4 from the node list of the Oracle Clusterware resource for the Support for Oracle RAC database, the following commands are run:

```
# /db_qfs/OraHome/crs/bin/crs_register ora.swb.swb3.inst \
-update -r "ora.pclus3.vip"
# /db_qfs/OraHome/crs/bin/crs_register ora.swb.swb4.inst \
-update -r "ora.pclus4.vip"
```

3. To remove from nodes pclus3 and pclus4 the Oracle Clusterware resources that represent Oracle Solaris Cluster resources, the following commands are run:

```
Removal of resource for Oracle database files from node pclus3
# /db_qfs/OraHome/crs/bin/crs_stop sun.pclus3.scaloramnt-OraData-rs
# /db_qfs/OraHome/crs/bin/crs_unregister sun.pclus3.scaloramnt-OraData-rs
# /db_qfs/OraHome/crs/bin/crs_profile -delete sun.pclus3.scaloramnt-OraData-rs \
-dir /var/cluster/ucmm/profile
```

Removal of resource for Oracle binary files from node pclus3

```
# /db_qfs/OraHome/crs/bin/crs_stop sun.pclus3.scaloramnt-OraHome-rs
# /db_qfs/OraHome/crs/bin/crs_unregister sun.pclus3.scaloramnt-OraHome-rs
# /db_qfs/OraHome/crs/bin/crs_profile -delete sun.pclus3.scaloramnt-OraHome-rs \
-dir /var/cluster/ucmm/profile
```

Removal of resource for Oracle database files from node pclus4

```
# /db_qfs/OraHome/crs/bin/crs_stop sun.pclus4.scaloramnt-OraData-rs
# /db_qfs/OraHome/crs/bin/crs_unregister sun.pclus4.scaloramnt-OraData-rs
# /db_qfs/OraHome/crs/bin/crs_profile -delete sun.pclus4.scaloramnt-OraData-rs \
-dir /var/cluster/ucmm/profile
```

Removal of resource for Oracle binary files from node pclus4

```
# /db_qfs/OraHome/crs/bin/crs_stop sun.pclus4.scaloramnt-OraHome-rs
# /db_qfs/OraHome/crs/bin/crs_unregister sun.pclus4.scaloramnt-OraHome-rs
# /db_qfs/OraHome/crs/bin/crs_profile -delete sun.pclus4.scaloramnt-OraHome-rs \
-dir /var/cluster/ucmm/profile
```

The commands remove Oracle Clusterware resources that represent the following Oracle Solaris Cluster resources:

- `scaloramnt-OraData-rs` – A resource of type `SUNW.ScalMountPoint` that represents the mount point of the file system for database files
- `scaloramnt-OraHome-rs` – A resource of type `SUNW.ScalMountPoint` that represents the mount point of the file system for binary files and associated files

After the resource is removed from nodes `pclus3` and `pclus4`, Oracle Database utilities are used to remove the following items from these nodes:

- The Support for Oracle RAC database
- The Support for Oracle RAC software
- The Oracle Clusterware software

The removal of these items is beyond the scope of this example.

4. To switch the resource group that contains resources for the StorageTek QFS metadata server to node `pclus1`, the following command is run:

```
# clresourcegroup switch -n pclus1 qfsmds-rg
```

5. To remove nodes `pclus3` and `pclus4` from the resource group that contains scalable file-system mount-point resources, the following commands are run:

```
# clresourcegroup offline -n pclus3,pclus4 scalmnt-rg
# clresourcegroup remove-node -n pclus3,pclus4 scalmnt-rg
```

6. To remove nodes `pclus3` and `pclus4` from the node list of the resource group that contains resources for the StorageTek QFS metadata server, the following command is run:

```
# clresourcegroup remove-node -n pclus3,pclus4 qfsmgs-rg
```

After nodes pclus3 and pclus4 are removed from the node list, the configuration of the StorageTek QFS shared file systems is removed from these nodes. This operation is beyond the scope of this example.

- To remove nodes pclus3 and pclus4 from the resource group that contains scalable device group resources, the following commands are run:

```
# clresourcegroup offline -n pclus3,pclus4 scaldg-rg
# clresourcegroup remove-node -n pclus3,pclus4 scaldg-rg
```

- To remove nodes pclus3 and pclus4 from the Solaris Volume Manager for Oracle Solaris Cluster multi-owner disk set oradg, the following command is run:

```
# metaset -s oradg -d -h pclus3 pclus4
```

- To remove nodes pclus3 and pclus4 from the Support for Oracle RAC framework resource group, the following commands are run:

```
# clresourcegroup offline -n pclus3,pclus4 rac-framework-rg
# clresourcegroup remove-node -n pclus3,pclus4 rac-framework-rg
```

- To remove nodes pclus3 and pclus4 from the multiple-owner volume-manager framework resource group, the following commands are run:

```
# clresourcegroup offline -n pclus3,pclus4 vucmm-framework-rg
# clresourcegroup remove-node -n pclus3,pclus4 vucmm-framework-rg
```

After the removal of nodes pclus3 and pclus4 from the framework resource groups, the Support for Oracle RAC software packages can optionally be removed from the nodes.

After the removal is complete, the status of resource groups and resources is as follows:

```
# clresourcegroup status
```

```
=== Cluster Resource Groups ===
```

Group Name	Node Name	Suspended	Status
rac-framework-rg	pclus1	No	Online
	pclus2	No	Online
vucmm-framework-rg	pclus1	No	Online
	pclus2	No	Online
scaldg-rg	pclus1	No	Online
	pclus2	No	Online

```

qfsmds-rg          pclus1      No      Online
                   pclus2      No      Offline

scalmnt-rg         pclus1      No      Online
                   pclus2      No      Online

rac_server_proxy-rg pclus1      No      Online
                   pclus2      No      Online
    
```

clresource status

=== Cluster Resources ===

Resource Name	Node Name	State	Status Message
-----	-----	-----	-----
rac-framework-rs	pclus1	Online	Online
	pclus2	Online	Online
crs_framework-rs	pclus1	Online	Online
	pclus2	Online	Online
vucmm-svm-rs	pclus1	Online	Online
	pclus2	Online	Online
scaloradg-rs	pclus1	Online	Online - Diskgroup online
	pclus2	Online	Online - Diskgroup online
qfs-mds-rs	pclus1	Online	Online - Service is online.
	pclus2	Offline	Offline
scaloramnt-OraData-rs	pclus1	Online	Online
	pclus2	Online	Online
scaloramnt-OraHome-rs	pclus1	Online	Online
	pclus2	Online	Online
rac_server_proxy-rs	pclus1	Online	Online - Oracle instance UP
	pclus2	Online	Online - Oracle instance UP

Upgrading Support for Oracle RAC

This chapter explains how to upgrade a configuration of Oracle Solaris Cluster Oracle Real Application Clusters (Oracle RAC).

If you upgrade Oracle Solaris Cluster core software, you must also upgrade the Support for Oracle RAC software. For more information, see [Upgrading Your Oracle Solaris Cluster 4.4 Environment](#).

Your existing configuration of Support for Oracle RAC might not include the Support for Oracle RAC framework resource group. In this situation, you *must* register and configure the Support for Oracle RAC framework resource group after upgrading Oracle Solaris Cluster software. Otherwise, Support for Oracle RAC cannot run with Oracle Solaris Cluster software. For more information, see [“Registering and Configuring the Support for Oracle RAC Framework Resource Group”](#) on page 63.

Information about how to upgrade an existing configuration is provided in the subsections that follow:

- [“Upgrading Resources in Support for Oracle RAC”](#) on page 173
- [“Adding Storage Resources for Oracle Database Files”](#) on page 176
- [“Adding Resources for Interoperation With Oracle Grid Infrastructure”](#) on page 177

Upgrading Resources in Support for Oracle RAC

Upgrade these resource types if you are upgrading from an earlier version of Support for Oracle RAC. For general instructions that explain how to upgrade a resource type, see [“Upgrading a Resource Type”](#) in [Planning and Administering Data Services for Oracle Solaris Cluster 4.4](#).

If you upgraded the cluster from the Oracle Solaris Cluster 4.0 or 4.1 release and are using the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type, you must re-create that

resource type and any resources of that type before you upgrade a resource type to the new Oracle Solaris Cluster version.

This section contains the following information:

- [“Re-creating the Oracle Grid Infrastructure `sun.storage_proxy.type` Resource Type and Resources” on page 174](#)
- [“Information for Registering the New Versions of Support for Oracle RAC Resource Types” on page 175](#)

Re-creating the Oracle Grid Infrastructure `sun.storage_proxy.type` Resource Type and Resources

Starting in the Oracle Solaris Cluster 4.2 release, the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type is changed. If the cluster is upgraded from the Oracle Solaris Cluster 4.0 or 4.1 release and is using the `sun.storage_proxy.type` resource type, this resource type and any resources of this type must be re-created before you upgrade a resource type to the new Oracle Solaris Cluster version.

▼ How to Re-create the Oracle Grid Infrastructure `sun.storage_proxy.type` Resource Type and Resources

If you upgraded the cluster from the Oracle Solaris Cluster 4.0 or 4.1 release, perform this procedure re-create the Oracle Grid Infrastructure `sun.storage_proxy.type` resource type and any resources of this type before you upgrade a resource type to the new Oracle Solaris Cluster version. This is necessary to accommodate the changes made in the new resource type version.

1. Remove any offline-restart dependencies.

Follow procedures in [“How to Remove a Dependency” on page 153](#) to remove offline-restart dependencies for all Oracle Solaris Cluster `SUNW.scalable_rac_server_proxy` or `SUNW.scalable_asm_diskgroup_proxy` resources that use the `SUNW.ScalMountPoint`, `SUNW.ScalDeviceGroup`, or `SUNW.scalable_acfs_proxy` resource type.

Removing these dependencies allows the Oracle Solaris Cluster resource to remain online when the underlying Oracle Grid Infrastructure resource is stopped.

2. Remove the Oracle Grid Infrastructure `sun.resource` resource.

Follow procedures in [“How to Delete the `sun.resource` Resource” on page 154](#).

3. Remove the Oracle Grid Infrastructure resource type.

```
# Grid_home/bin/crsctl delete type sun.storage_proxy
```

4. Recreate the Oracle Grid Infrastructure resource type and resources.

- If you are using SUNW.scalable_rac_server_proxy resources, follow procedures in [“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”](#) on page 243.
- If you are using SUNW.scalable_asm_diskgroup_proxy resources, follow procedures in [“How to Register and Configure HA for Oracle Database With Oracle Grid Infrastructure for a Cluster \(CLI\)”](#) in *Oracle Solaris Cluster Data Service for Oracle Database Guide*.

Information for Registering the New Versions of Support for Oracle RAC Resource Types

Note - If the cluster is upgraded from the Oracle Solaris Cluster 4.0 or 4.1 release and is using the Oracle Grid Infrastructure sun.storage_proxy.type resource type, that resource type and any resources of that type must be re-created before you upgrade a Support for Oracle RAC resource type to the new Oracle Solaris Cluster version. Follow procedures in [“Re-creating the Oracle Grid Infrastructure sun.storage_proxy.type Resource Type and Resources”](#) on page 174.

The following table lists the name of each Support for Oracle RAC resource type and its Resource Type Registration (RTR) file name.

TABLE 19 Support for Oracle RAC Resource Types

Resource Type	RTR File
SUNW.crs_framework	/usr/cluster/lib/rgm/rtreg/SUNW.crs_framework
SUNW.qfs [†]	/usr/cluster/lib/rgm/rtreg/SUNW.qfs
SUNW.rac_framework	/usr/cluster/lib/rgm/rtreg/SUNW.rac_framework
SUNW.rac_svm	/usr/cluster/lib/rgm/rtreg/SUNW.rac_svm
SUNW.scalable_acfs_proxy	/opt/cluster/lib/rgm/rtreg/SUNW.scalable_acfs_proxy
SUNW.scalable_asm_diskgroup_proxy	/opt/cluster/lib/rgm/rtreg/SUNW.scalable_asm_diskgroup_proxy
SUNW.scalable_asm_instance	/opt/cluster/lib/rgm/rtreg/SUNW.scalable_asm_instance

Resource Type	RTR File
SUNW.scalable_asm_instance_proxy	/opt/cluster/lib/rgm/rtreg/SUNW.scalable_asm_instance_proxy
SUNW.scalable_rac_server_proxy	/opt/cluster/lib/rgm/rtreg/SUNW.scalable_rac_server_proxy
SUNW.ScalDeviceGroup	/opt/cluster/lib/rgm/rtreg/SUNW.ScalDeviceGroup
SUNW.ScalMountPoint	/opt/cluster/lib/rgm/rtreg/SUNW.ScalMountPoint
SUNW.vucmm_framework	/usr/cluster/lib/rgm/rtreg/SUNW.vucmm_framework
SUNW.vucmm_svm	/usr/cluster/lib/rgm/rtreg/SUNW.vucmm_svm
SUNW.wait_zc_boot	/usr/cluster/lib/rgm/rtreg/SUNW.wait_zc_boot

[†]Delivered in the StorageTek QFS product.

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

```
# clresourcetype show resource-type
```

resource-type

Specifies the resource type whose version you are determining.

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

```
# grep -i RT_VERSION /path/RTRfilename
```

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

Note - In the global zone, the Resource Group option of `clsetup` locates for you the available resource type versions you can upgrade to.

Adding Storage Resources for Oracle Database Files

Oracle Solaris Cluster software includes resource types that 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 Database files, add storage resources to manage the availability of the storage on which the Oracle Database software depends.

Before adding storage resources for Oracle Database files, ensure that you have upgraded resources in the Support for Oracle RAC framework resource group. For more information, see [“Upgrading Resources in Support for Oracle RAC”](#) on page 173.

For detailed instructions for adding storage resources for Oracle Database files, see [“Registering and Configuring Storage Resources for Oracle Database Files” on page 78](#).

Follow these instructions only if you are adding storage resources to an existing configuration of storage management schemes. If you are extending an existing configuration of Support for Oracle RAC by supporting additional storage management schemes, see the following sections:

- [Chapter 2, “Configuring Storage for Oracle Database Files”](#)
- [“How to Add a Volume Manager Resource to the SUNW.vucmm_framework Resource Group” on page 151](#)

Adding Resources for Interoperation With Oracle Grid Infrastructure

Oracle Solaris Cluster software includes resource types that enable Oracle Solaris Cluster software and Oracle Grid Infrastructure to interoperate. These resource types also enable Support for Oracle RAC database instances to be administered from Oracle Solaris Cluster.

Before adding resources for interoperation with Oracle Grid Infrastructure, ensure that you have performed the tasks in the following sections:

- [“Upgrading Resources in Support for Oracle RAC” on page 173](#)
- [“Adding Storage Resources for Oracle Database Files” on page 176](#)

For detailed instructions for adding resources for interoperation with Oracle Grid Infrastructure, see [“Configuring Resources for Support for Oracle RAC Database Instances” on page 97](#).

◆◆◆ **A P P E N D I X 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 Support for Oracle RAC Configurations in the Global Cluster” on page 180](#)
- [“Sample Support for Oracle RAC Configurations in a Zone Cluster” on page 186](#)

Sample Support for Oracle RAC Configurations in the Global Cluster

FIGURE 1 Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster




FIGURE 2 Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System




FIGURE 3 Configuration of Support for Oracle RAC With StorageTek QFS Shared File System and Hardware RAID




FIGURE 4 Configuration of Support for Oracle RAC With a NAS Device

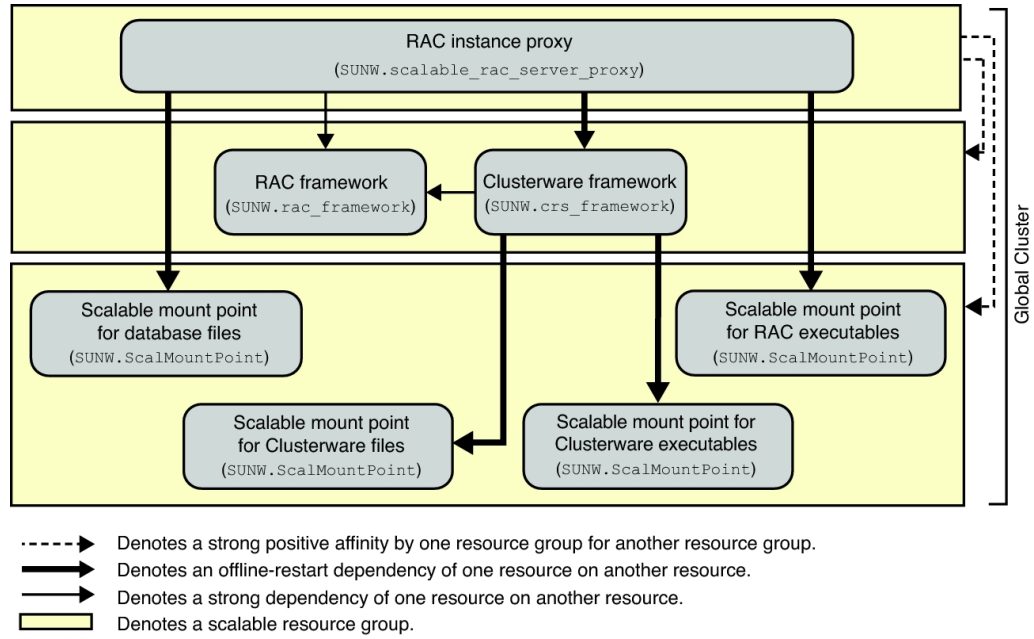


FIGURE 5 Configuration of Support for Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster

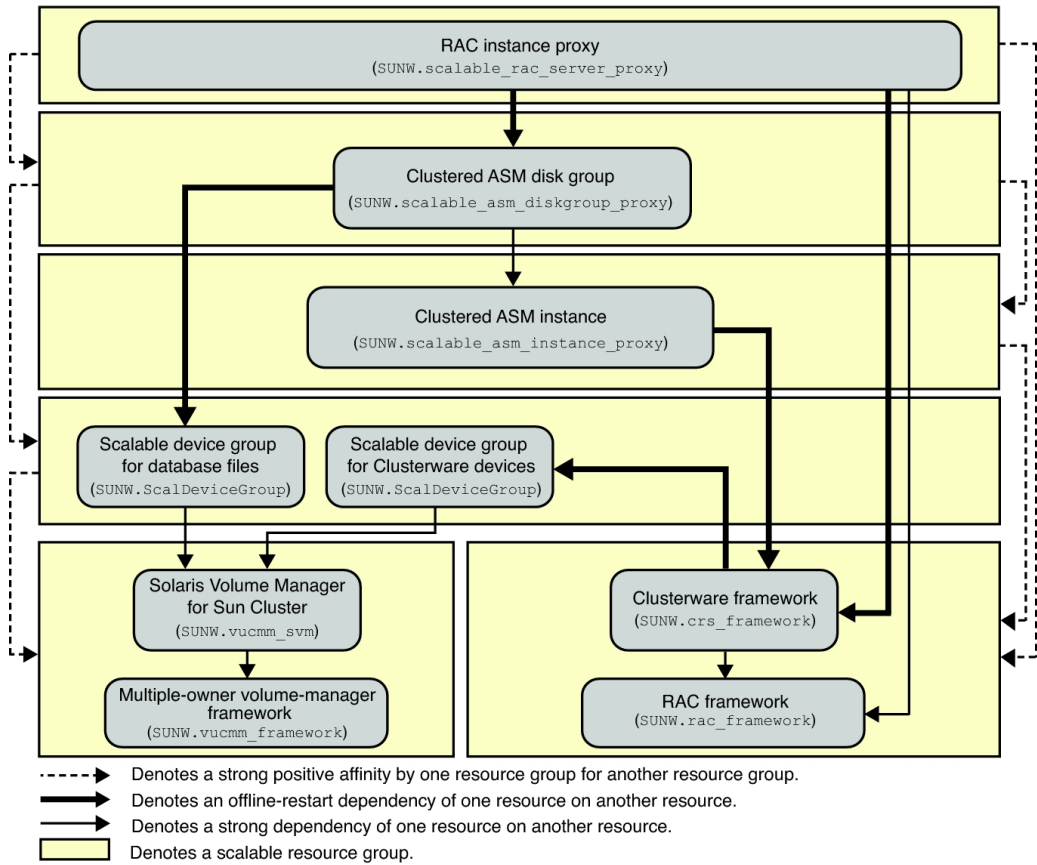



FIGURE 6 Configuration of Support for Oracle RAC With Oracle ASM and Hardware RAID



Sample Support for Oracle RAC Configurations in a Zone Cluster

FIGURE 7 Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster in a Zone Cluster

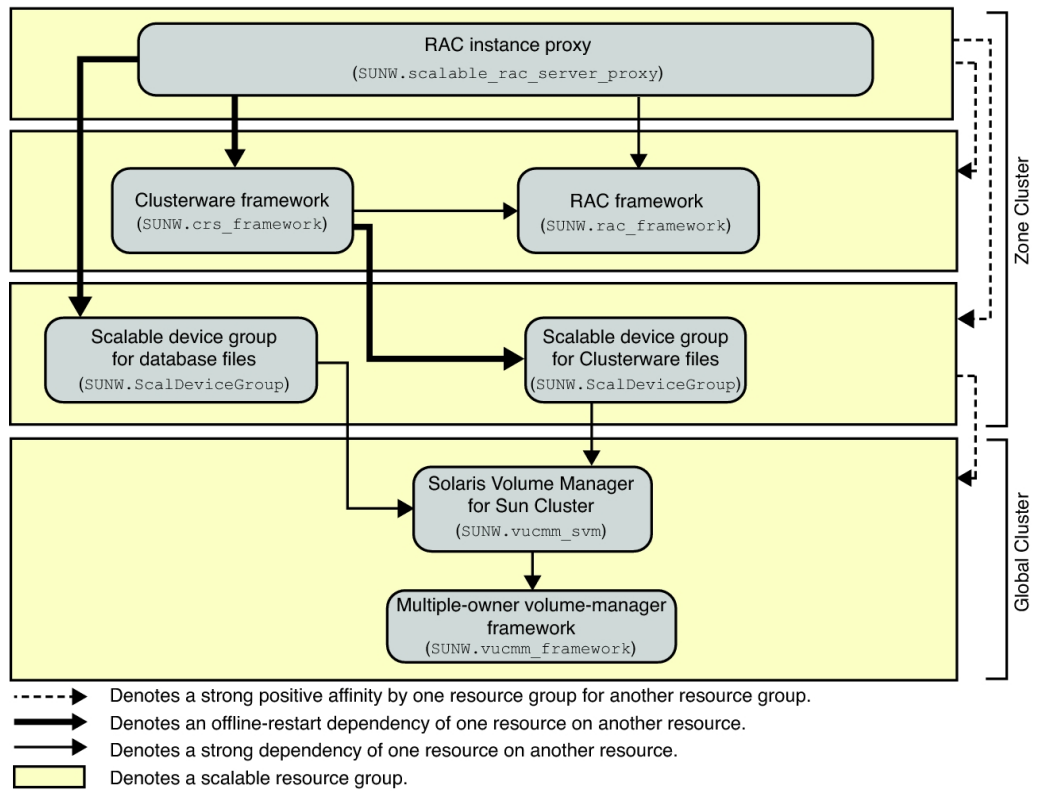


FIGURE 8 Configuration of Support for Oracle RAC With Solaris Volume Manager for Sun Cluster and StorageTek QFS Shared File System in a Zone Cluster

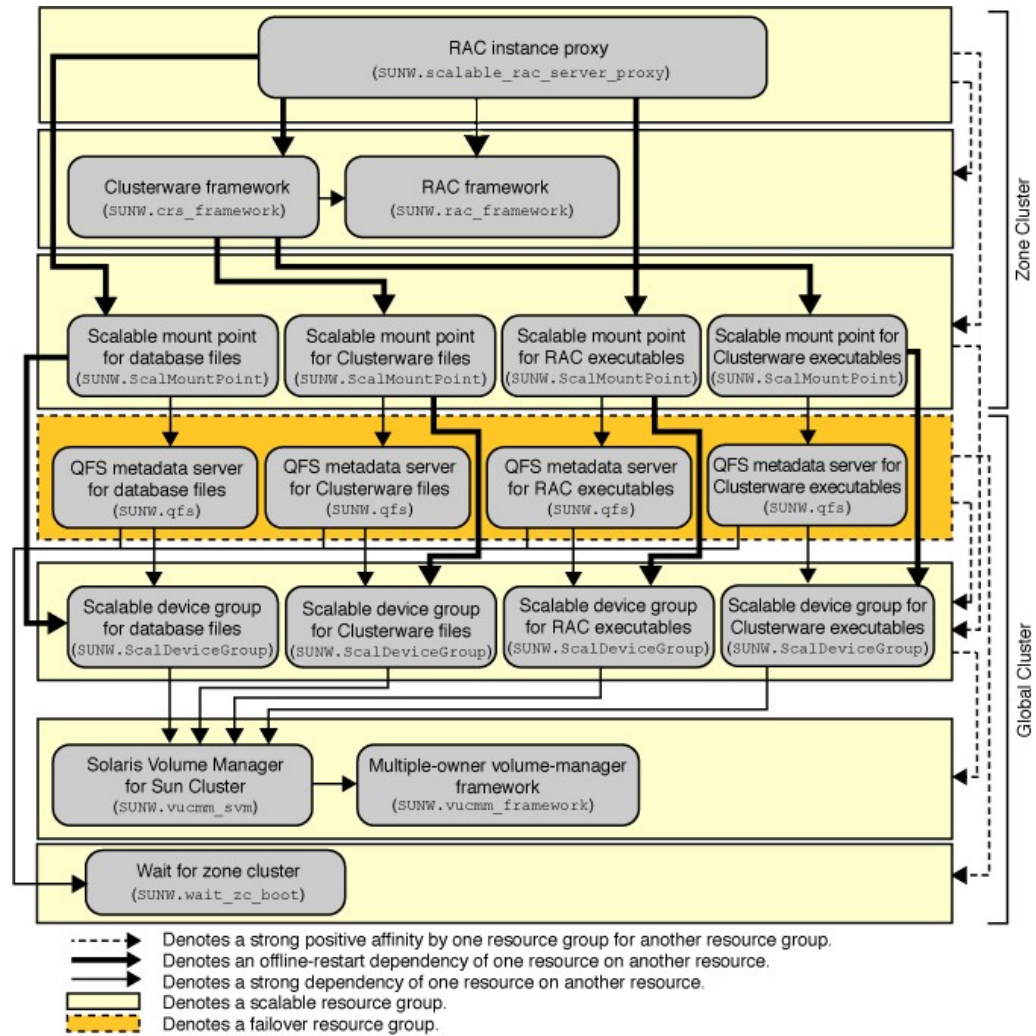


FIGURE 9 Configuration of Support for Oracle RAC With StorageTek QFS Shared File System and Hardware RAID in a Zone Cluster

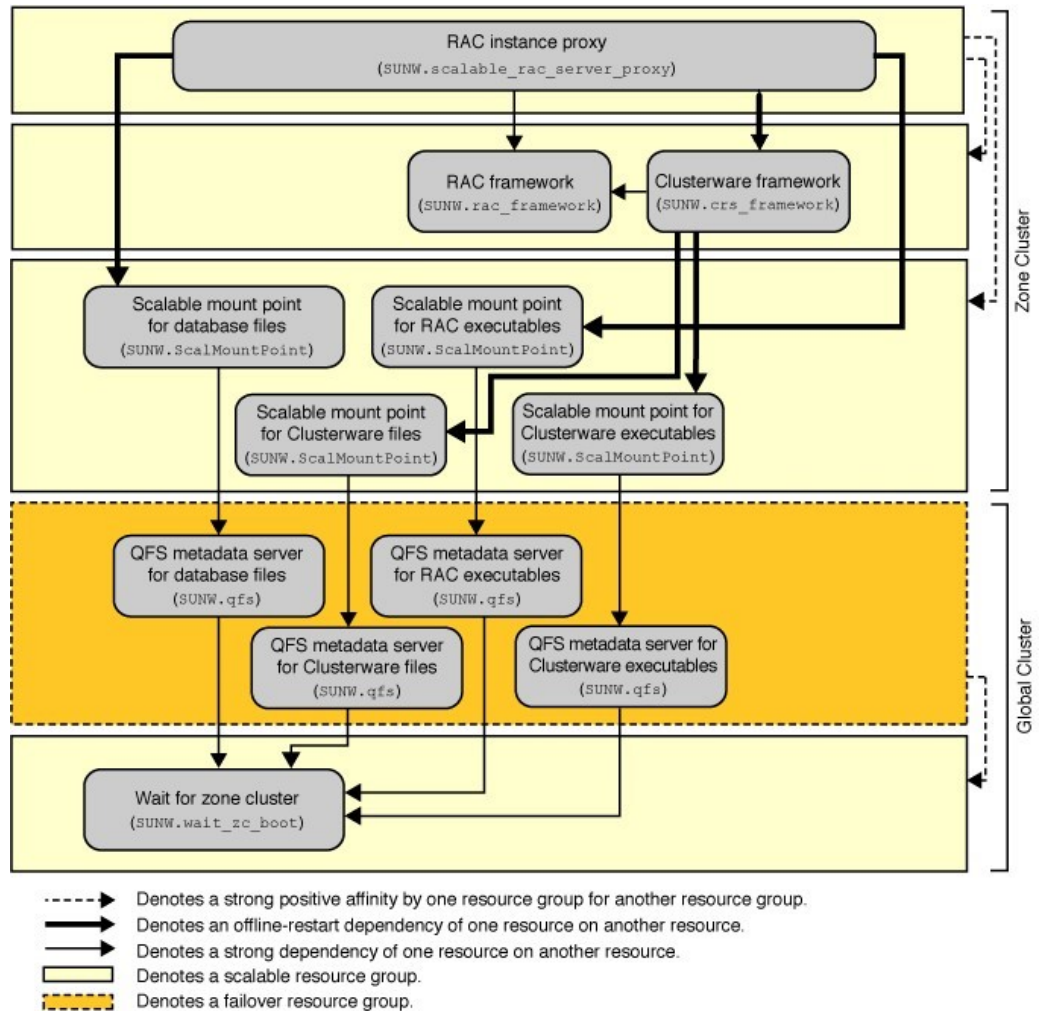


FIGURE 10 Configuration of Support for Oracle RAC With a NAS Device in a Zone Cluster




FIGURE 11 Configuration of Support for Oracle RAC With Oracle ASM and Solaris Volume Manager for Sun Cluster in a Zone Cluster

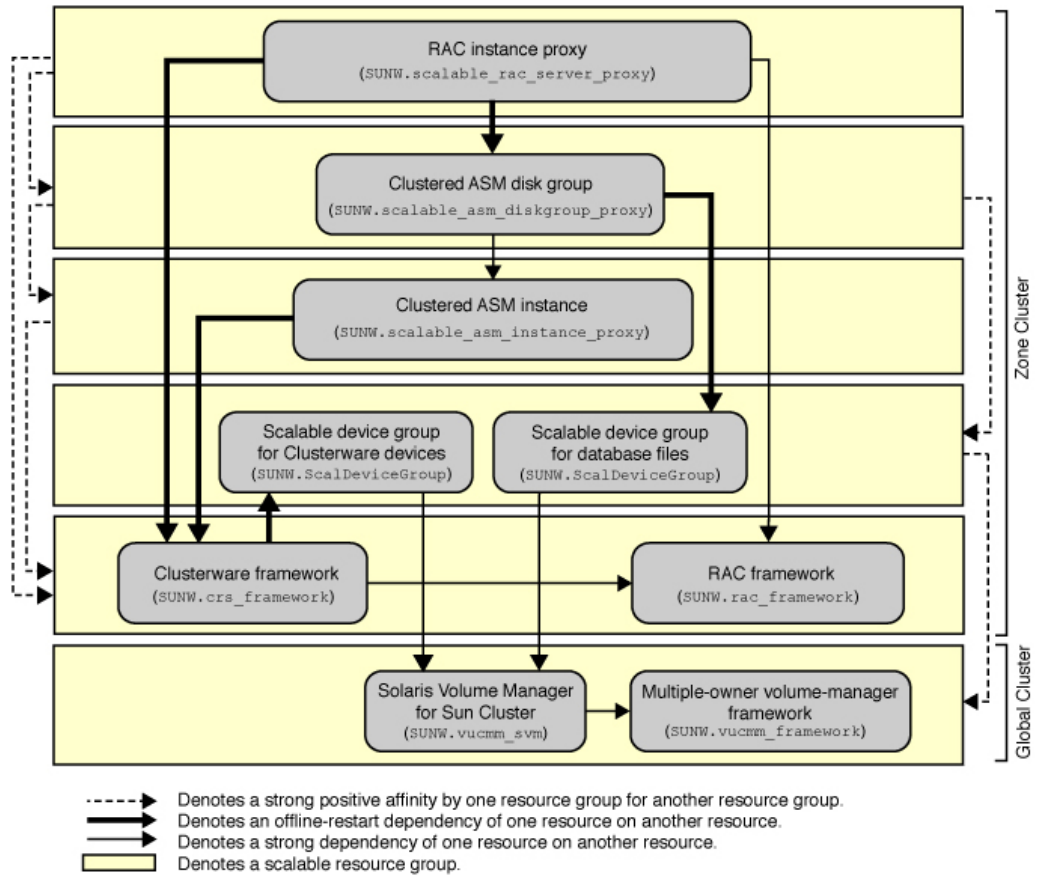
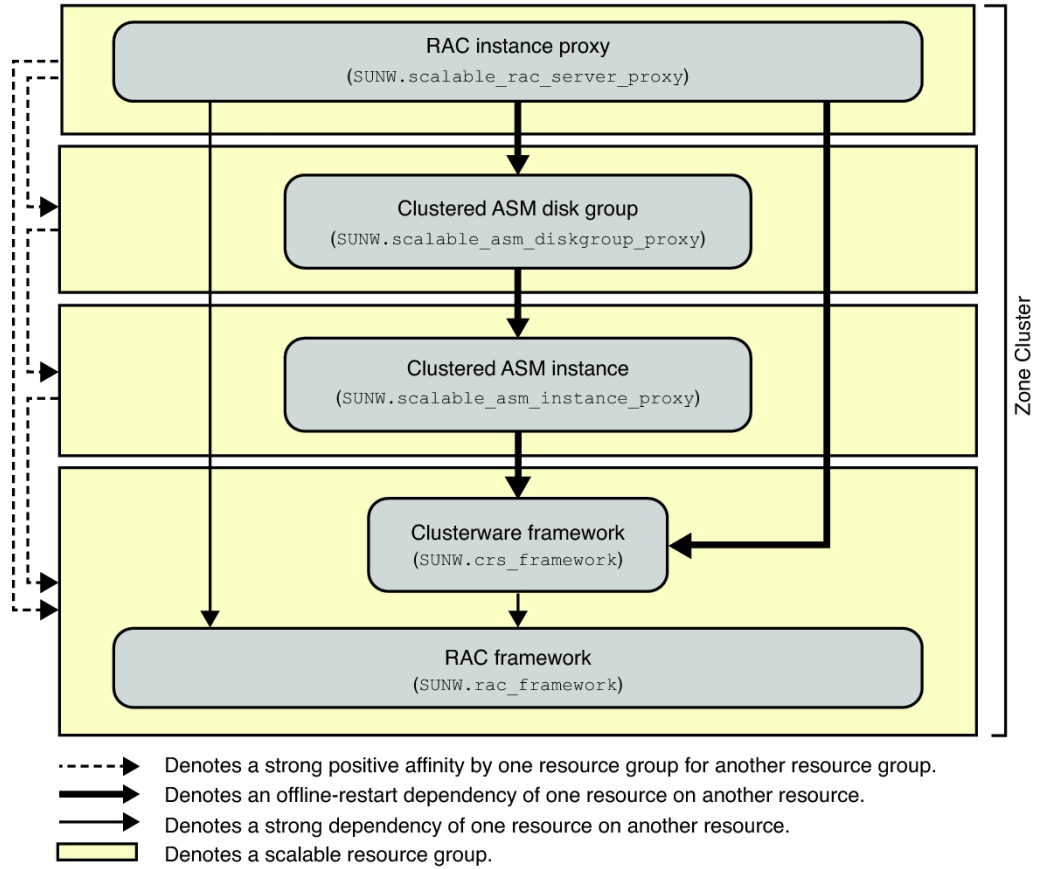


FIGURE 12 Configuration of Support for Oracle RAC With Oracle ASM and Hardware RAID in a Zone Cluster



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 20, “Preset Actions for DBMS Errors,” on page 193.](#)
- Logged alerts for which an action is preset are listed in [Table 21, “Preset Actions for Logged Alerts,” on page 199.](#)

TABLE 20 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
471	RESTART	*	di	Oracle background process died
472	RESTART	*	di	Oracle background process died

Error Number	Action	Connection State	New State	Message
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
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

Error Number	Action	Connection State	New State	Message
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	*	*	osnfpn: cannot create shared memory segment
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	

Error Number	Action	Connection State	New State	Message
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	
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	*	*	

Error Number	Action	Connection State	New State	Message
7324	RESTART	*	*	
7325	RESTART	*	*	
7351	RESTART	*	*	
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	*	*	
9757	RESTART	*	*	
9756	RESTART	*	*	
9758	RESTART	*	*	
9761	RESTART	*	*	
9765	RESTART	*	*	

Error Number	Action	Connection State	New State	Message
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	*	*	
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.

Error Number	Action	Connection State	New State	Message
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 21 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)

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:

- “[SUNW.crs_framework Extension Properties](#)” on page 201
- “[SUNW.rac_framework Extension Properties](#)” on page 202
- “[SUNW.scalable_asm_diskgroup_proxy Extension Properties](#)” on page 203
- “[SUNW.scalable_asm_instance_proxy Extension Properties](#)” on page 205
- “[SUNW.scalable_rac_server_proxy Extension Properties](#)” on page 206
- “[SUNW.ScalDeviceGroup Extension Properties](#)” on page 209
- “[SUNW.ScalMountPoint Extension Properties](#)” on page 211
- “[SUNW.vucmm_framework Extension Properties](#)” on page 214
- “[SUNW.vucmm_svm Extension Properties](#)” on page 214
- “[SUNW.wait_zc_boot Extension Properties](#)” on page 217

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 [rg_properties\(7\)](#) 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_acfs_proxy Extension Properties

acfs_mountpoint

This property specifies the mount point of an Oracle ACFS file system.

Data Type	String
Default	No default defined
Minimum length	1
Tunable	When disabled

Debug_level

Note - All SQL*Plus messages that the Oracle ACFS 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 Oracle ACFS 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

Proxy_probe_interval

This property specifies the interval, in seconds, between probes of the Oracle ACFS resource for which this resource is acting as a proxy.

Data Type	Integer
Range	5–300
Default	30
Tunable	Any time

Proxy_probe_timeout

This property specifies the timeout value, in seconds, that the proxy monitor uses when checking the status of the Oracle ACFS resource for which this resource is acting as a proxy.

Data Type	Integer
Range	5–120
Default	60
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 Database 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

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 Support for 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 Support for Oracle RAC database is specified during the installation of Support for 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 Support for 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 Support for 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 Database home directory. The Oracle Database home directory contains the binary files, log files, and parameter files for the Oracle Database software.

Data Type: String

Range: Not applicable

Default: No default defined

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 a Support for 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 Database 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. If you are using Solaris Volume Manager for Sun Cluster in a zone cluster and are not using all the devices in the metaset, you must set this property. Otherwise, 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.

s-qfs Specifies that the file system is a StorageTek QFS shared file system.

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 StorageTek QFS shared file system, these options are obtained from the `/etc/opt/SUNWsamfs/samfs.cmd` file.
- 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 StorageTek QFS shared file system, set this property to the name that was assigned to the file system when the file system was created. The file system must be correctly configured. For more information, see your StorageTek QFS shared file system documentation.

- 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 [Managing Network-Attached Storage Devices in an Oracle Solaris Cluster 4.4 Environment](#).

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 Oracle Solaris 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

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” on page 219](#)
- [“Registering and Configuring the Framework Resource Groups by Using Oracle Solaris Cluster Maintenance Commands” on page 220](#)
- [“Registering and Configuring Oracle ASM Resource Groups \(CLI\)” on page 226](#)
- [“Creating Storage Management Resources by Using Oracle Solaris Cluster Maintenance Commands” on page 229](#)
- [“Creating Resources for Interoperation With Oracle Grid Infrastructure by Using Oracle Solaris Cluster Maintenance Commands” on page 240](#)

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 115](#)
- [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 Support for Oracle RAC Framework Resource Group \(clsetup\)” on page 64](#). 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 220](#)
- [“How to Register and Configure the Framework Resource Groups in the Global Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 221](#)
- [“How to Register and Configure the Support for Oracle RAC Framework Resource Group in a Zone Cluster by Using Oracle Solaris Cluster Maintenance Commands” on page 224](#)

Overview of the Framework Resource Groups

This section describes the following framework resource groups:

- [“Support for Oracle RAC Framework Resource Group” on page 220](#)
- [“Multiple-Owner Volume-Manager Framework Resource Group” on page 221](#)

Support for Oracle RAC Framework Resource Group

The Support for Oracle RAC framework resource group enables Support for 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 Support for Oracle RAC to be managed by using Oracle Solaris Cluster commands

Note - The resource types that are defined for the Support for 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 Support for 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. **Assume the root role or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**
2. **Create a scalable Support for Oracle RAC resource group.**

Note - If you are performing the steps in this procedure to register and configure the Support for Oracle RAC resource framework in a zone cluster and Support for Oracle RAC support is not also required in the global cluster, you do not need to also create a Support for Oracle RAC framework resource group in the global cluster. In that case, skip this step and proceed to [Step 6](#).

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-fwk-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 Support for 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. Bring online and in a managed state the Support for Oracle RAC framework resource group and its resources.

```
# clresourcegroup online -eM rac-fwk-rg]
```

rac-fwk-rg

Specifies that the Support for Oracle RAC resource group that you created in [Step 2](#) is to be moved to the MANAGED state and brought online.

If you intend to create a multiple-owner volume-manager framework resource group, proceed to [Step 6](#). Otherwise, creation of the Support for Oracle RAC framework resource group is completed.

6. Create a scalable multiple-owner volume-manager framework resource group, if one does not already exist.

```
# clresourcegroup create -n nodelist -S vucmm-fwk-rg
```

-n nodelist=nodelist

Specifies the same node list that you configured for the scalable Support for Oracle RAC resource group.

vucmm-fwk-rg

Specifies the name that you are assigning to the multiple-owner volume-manager framework resource group.

7. Register the SUNW.vucmm_framework resource type.

```
# clresourcetype register SUNW.vucmm_framework
```

8. Add an instance of the SUNW.vucmm_framework resource type to the resource group that you created in [Step 6](#).

```
# clresource create -g vucmm-fwk-rg -t SUNW.vucmm_framework vucmm-fwk-rs
```

-g vucmm-fwk-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 6](#).

vucmm-fwk-rs

Specifies the name that you are assigning to the SUNW.vucmm_framework resource.

9. Register and add an instance of the resource type that represents the Solaris Volume Manager for Sun Cluster volume manager that you are using for Oracle Database files.

a. Register the resource type.

```
# clresourcetype register SUNW.vucmm_svm
```

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 vucmm-fwk-rg \  
-t SUNW.vucmm_svm \  
-p resource_dependencies=vucmm-fwk-rs \  
vucmm-svm-rs
```

```
-g vucmm-fwk-rg
```

Specifies the resource group that you created in [Step 6](#).

```
-p resource_dependencies=vucmm-fwk-rs
```

Specifies that this instance depends on the framework resource that you created in [Step 8](#).

```
vucmm-svm-rs
```

Specifies the name that you are assigning to the SUNW.vucmm_svm resource.

10. Bring online and in a managed state the multiple-owner volume-manager framework resource group and its resources.

```
# clresourcegroup online -eM vucmm-fwk-rg
```

```
vucmm-fwk-rg
```

Specifies that the multiple-owner volume-manager framework resource group that you created in [Step 6](#) is to be moved to the MANAGED state and brought online.

▼ **How to Register and Configure the Support for Oracle RAC Framework Resource Group in a**

Zone Cluster by Using Oracle Solaris Cluster Maintenance Commands

Perform the steps in this procedure to register and configure the Support for Oracle RAC framework resource group in a zone cluster for the StorageTek QFS shared file system with Solaris Volume Manager.

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. **Assume the root role or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**
2. **Create a scalable Support for Oracle RAC 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 -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
```

3. **Register the `SUNW.rac_framework` resource type.**

```
# clresourcetype register -Z zcname 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 -Z zcname -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. **Bring online and in a managed state the Support for Oracle RAC framework resource group and its resources.**

```
# clresourcegroup online -Z zcname -eM rac-fmwk-rg
```

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 \(CLI\)” on page 226](#)
- [“How to Register and Configure Oracle ASM Resource Groups in a Zone Cluster \(CLI\)” on page 228](#)

▼ 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 Support for Oracle RAC framework resource and the Oracle Grid Infrastructure framework resource.
 - Ensure that the Support for 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-fwk-rg* by *asm-inst-rg*.**

```
# clresourcegroup set -p rg_affinities=++rac-fwk-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.**

```
# clresource create -g asm-inst-rg \  
-t SUNW.scalable_asm_instance_proxy \  
-p ORACLE_HOME=Grid_home \  
-p CRS_HOME=Grid_home \  
-p resource_dependencies_offline_restart=crs-fwk-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[,storage-rs \  
-d asm-dg-rs
```

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

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

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

```
# clresourcegroup online -M 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 Support for Oracle RAC framework resource and the Oracle Grid Infrastructure framework resource.
 - Ensure that the Support for 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.

```
# clresource create -Z zcname -g asm-inst-rg \  
-t SUNW.scalable_asm_instance_proxy \  
-p ORACLE_HOME=Grid_home \  
-p ORACLE_HOME=Grid_home \  
-p ORACLE_HOME=Grid_home \  
-p ORACLE_HOME=Grid_home
```

```
-p CRS_HOME=Grid_home \  
-p resource_dependencies_offline_restart=crs-fwk-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.

```
-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[,storage-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 -M asm-inst-rg
```

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

```
# clresourcegroup online -Z zcname -M 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 Database Files \(clsetup\)”](#) on page 79. The following information is in this section:

- [“Resources for Scalable Device Groups and Scalable File-System Mount Points”](#) on page 230

- [“Resources for the StorageTek QFS Metadata Server” on page 231](#)
- [“How to Create a Resource for a Scalable Device Group in the Global Cluster” on page 231](#)
- [“How to Create a Resource for a Scalable Device Group in a Zone Cluster” on page 232](#)
- [“How to Register and Configure Resources for the StorageTek QFS Metadata Server in the Global Cluster” on page 233](#)
- [“How to Register and Configure Resources for the StorageTek QFS Metadata Server for a Zone Cluster” on page 235](#)
- [“How to Create a Resource for a File-System Mount Point in the Global Cluster” on page 236](#)
- [“How to Create a Resource for a File-System Mount Point in Zone Cluster” on page 238](#)

The following resources to represent storage for Oracle files are required:

- Resources for scalable device groups and scalable file-system mount points
- Resources for the StorageTek QFS metadata server

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 StorageTek QFS or 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 StorageTek QFS shared file system or NFS file system on a qualified NAS device that you are using for Oracle files.

The resource that represents a StorageTek QFS shared file system can start only if the file system's StorageTek QFS metadata server is running. Similarly, the resource that represents a StorageTek QFS shared file system can stop only if the file system's StorageTek QFS metadata server is stopped. To meet this requirement, configure a resource for each StorageTek QFS metadata server. For more information, see [“Resources for the StorageTek QFS Metadata Server” on page 231](#).

Resources for the StorageTek QFS Metadata Server

If you are using the StorageTek QFS shared file system, create one resource for each StorageTek QFS metadata server. The configuration of resource groups for these resources depends on the number of file systems in your configuration.

- If your configuration contains a small number of file systems, create one resource group for all resources for the StorageTek QFS metadata server.
- If your configuration contains a large number of file systems, configure resources for the StorageTek QFS metadata server in several resource groups as follows:
 - Distribute resources among the resource groups to ensure optimum load balancing.
 - Select a different primary node for each resource group to prevent all the resource groups from being online on the same node simultaneously.

▼ 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. **Assume the root role 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 Support for 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=++vucmm-fwk-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 Support for 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=global:vucmm-svm-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 -M 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. Assume the `root` role 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 Support for 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=++vucmm-fwk-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 Database 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 Support for 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=global:vucmm-svm-rs{local_node} \
-p diskgroupname=disk-group \
-p logicaldeviceList="dN[,dX...]"
scal-dg-rs
```

-p logicaldeviceList="dN[,dX...]" Specifies a comma-separated list of Solaris Volume Manager for Sun Cluster devices. Set this property if the devices that you want to use are a subset of all the devices in the metaset. When this property is set, storage validation is skipped for devices in the metaset that are not specified to this property. Otherwise, storage validation might fail.

5. Bring online and in a managed state the resource group that you created in [Step 2](#).

```
# clresourcegroup online -Z zcname -M scal-dg-rg
```

▼ How to Register and Configure Resources for the StorageTek QFS Metadata Server in the Global Cluster

Perform this task *only* if you are using the StorageTek QFS shared file system.

Perform this procedure on only one node of the cluster.

1. **Assume the `root` role or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**

2. **Create a failover resource group to contain the resources for the StorageTek QFS metadata server.**

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 231.

```
# clresourcegroup create -n nodelist \  
[-p rg_affinities=++scal-dg-rg] \  
[-p rg_description="description"] \  
qfs-mds-rg
```

3. **Register the `SUNW.qfs` resource type.**

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

4. **For each StorageTek QFS shared file system that you are using, add an instance of the `SUNW.qfs` resource type to the resource group that you created in [Step 2](#).**

Each instance of `SUNW.qfs` represents the metadata server of the file system.

If you are also using a volume manager, set a strong dependency by the instance of `SUNW.qfs` 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 231.

```
# clresource create -t SUNW.qfs -g qfs-mds-rg \  
-p qfsfilesystem=path \  
[-p resource_dependencies=scal-dg-rs] \  
qfs-mds-rs
```

5. **Bring online and in a managed state the resource group that you created in [Step 2](#).**

```
# clresourcegroup online -M qfs-mds-rg
```

▼ How to Register and Configure Resources for the StorageTek QFS Metadata Server for a Zone Cluster

Perform the steps in this procedure to register and configure resources for the StorageTek QFS metadata server for a zone cluster.

Perform these steps in the global cluster.

1. **Assume the root role or assume a role that provides `solaris.cluster.admin` and `solaris.cluster.modify` RBAC authorizations.**
2. **Create a scalable resource group to contain the `SUNW.wait_zc_boot` resource in the global cluster.**

```
# clresourcegroup create -n nodelist \
-p rg_mode=Scalable \
-p maximum primaries=num-in-list \
-p desired primaries=num-in-list \
zc-wait-rg
```

3. **Register the `SUNW.wait_zc_boot` resource type.**

```
# clresourcetype register SUNW.wait_zc_boot
```

4. **Add an instance of the `SUNW.wait_zc_boot` resource type to the resource group that you created in [Step 2](#).**

```
# clresource create -g zc-wait-rg -t SUNW.wait_zc_boot \
-p ZCName=zcname zc-wait-rs
```

5. **Bring online and in a managed state the resource group that you created in [Step 2](#).**

```
# clresourcegroup online -M zc-wait-rg
```

6. **Create a failover resource group to contain the resources for the StorageTek QFS metadata server.**

Set a strong positive affinity by the resource group for the resource group that contains the `SUNW.wait_zc_boot` resource that is configured for the zone cluster.

If you also use 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 231.

```
# clresourcegroup create -n nodelist \  
-p rg_affinities=++wait-zc-rg[,++scal-dg-rg] \  
[-p rg_description="description"] \  
qfs-mds-rg
```

7. Register the SUNW.qfs resource type.

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

8. Add an instance of the SUNW.qfs resource type to the resource group that you created in [Step 6](#) for each StorageTek QFS shared file system that you are using.

Each instance of SUNW.qfs represents the metadata server of the file system.

Set a strong dependency by the instance of SUNW.qfs on the SUNW.wait_zc_boot resource that is configured for the zone cluster.

If you are also using a volume manager, set a strong dependency by the instance of SUNW.qfs 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 231.

```
# clresource create -t SUNW.qfs -g qfs-mds-rg \  
-p qfsfilesystem=path \  
\-p resource_dependencies=zc-wait-rs[,scal-dg-rs] \  
qfs-mds-rs
```

9. Bring online and in a managed state the resource group that you created in [Step 6](#).

```
# clresourcegroup online -M qfs-mds-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. Assume the root role 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 231.

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 StorageTek QFS shared file system, type the following command:

Set a strong dependency by the instance of `SUNW.ScaLMountPoint` on the resource for the StorageTek QFS metadata server for the file system. The resource for the StorageTek QFS metadata server set is created in [“How to Register and Configure Resources for the StorageTek QFS Metadata Server in the Global Cluster”](#) on page 233.

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 231.

```
# clresource create -t SUNW.ScaLMountPoint -g scal-mp-rg \
-p resource_dependencies=qfs-mds-rs \
[-p resource_dependencies_offline_restart=scal-dg-rs] \
-p mountpointdir=mp-path \
-p filesystemtype=s-qfs \
-p targetfilesystem=fs-name qfs-mp-rs
```

- For each file system on a qualified NAS device, type the following command:

```
# clresource create -t SUNW.ScalMountPoint -g scal-mp-rg \  
-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 -eM scal-mp-rg
```

▼ How to Create a Resource for a File-System Mount Point in Zone Cluster

Perform the steps in this procedure to create a resource for a file-system mount point in a zone cluster. For Oracle RAC configurations with the StorageTek QFS shared file system on Solaris Volume Manager for Sun Cluster and the StorageTek QFS shared file system on hardware RAID, you create a scalable resource group to contain all the scalable mount point resources in a zone cluster.

Note - The node list is the list of zone-cluster nodes.

Perform this procedure from the global cluster.

1. Assume the `root` role 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 in zone cluster.

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 231.

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 zcnodelist \
-p desired_primaries=num-in-list \
-p maximum_primaries=num-in-list \
[-p rg_affinities==+global:scal-dg-rg] \
[-p rg_description="description"] \
-p rg_mode=Scalable scal-mp-rg
```

3. Register the `SUNW.ScalMountPoint` resource type.

```
# clresourcetype register -Z zcname SUNW.ScalMountPoint
```

4. For each StorageTek QFS 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](#).

```
# clresource create -Z zcname -t SUNW.ScalMountPoint -d -g scal-mp-rg \
-p resource_dependencies=global:qfs-mds-rs \
[-y resource_dependencies_offline_restart=global:scal-dg-rs \]
-p mountpointdir=mp-path \
-p filesystemtype=s-qfs \
-p targetfilesystem=fs-name qfs-mp-rs
```

- Set a strong dependency by the instance of `SUNW.ScalMountPoint` on the resource for the StorageTek QFS metadata server for the file system.

The resource for the StorageTek QFS metadata server set is created in [“How to Register and Configure Resources for the StorageTek QFS Metadata Server in the Global Cluster”](#) on page 233.

- 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 231.

5. Bring online and in a managed state the resource group that you created in [Step 2](#).

```
# clresourcegroup online -Z zcname -M 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 98. The following information is in this section:

- [“How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster”](#) on page 243
- [“How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure”](#) on page 244
- [“How to Create Oracle Solaris Cluster Resources in a Zone Cluster for Interoperation With Oracle Grid Infrastructure”](#) on page 247

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 Support for 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 `pc1us1` that represents the Oracle Solaris Cluster resource `scal-dg-rs` must be as follows:

```
sun.pc1us1.scal-dg-rs
```

FIGURE 13 Proxy Resources for Configurations With a Volume Manager

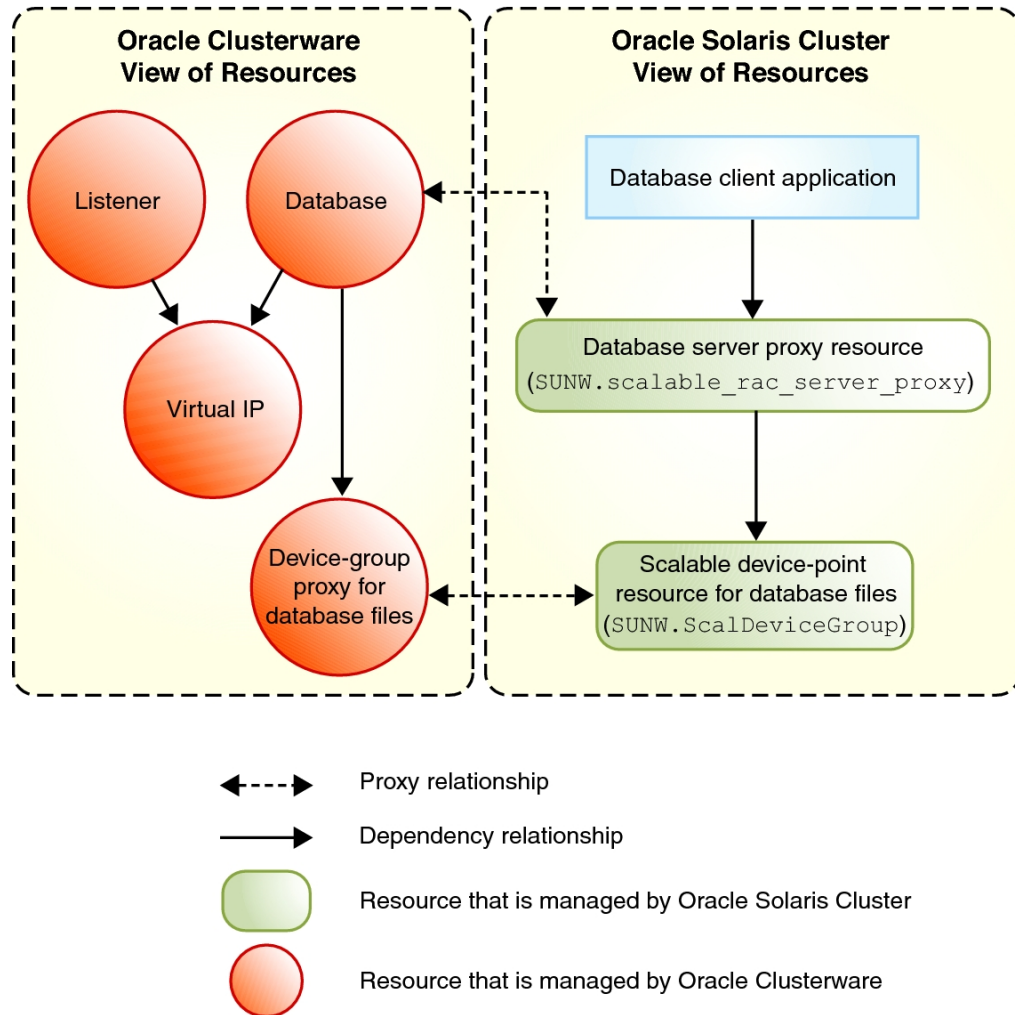



FIGURE 14 Proxy Resources for Configurations With a Shared File System



▼ How to Create an Oracle Grid Infrastructure Resource for Interoperation With Oracle Solaris Cluster

For Support 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, SUNW.scalable_acfs_proxy, 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, SUNW.scalable_acfs_proxy, or SUNW.ScalMountPoint resource is online. And the SUNW.ScalDeviceGroup, SUNW.scalable_acfs_proxy, or SUNW.ScalMountPoint resource only comes online if the actual volume-manager disk set or disk group or the mount point is online.

1. **Assume the root role 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 local_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, SUNW.scalable_acfs_proxy, 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 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 153.

▼ How to Create Oracle Solaris Cluster Resources in the Global Cluster for Interoperation With Oracle Grid Infrastructure

Note - If you are configuring Support for 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 247.

Perform this procedure on only one node of the cluster.

1. **Assume the root role 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 Support for Oracle RAC Framework Resource Group” on page 63](#).

Set a strong dependency by the instance of `SUNW.crs_framework` on the instance of `SUNW.rac_framework` in the Support for 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 231](#).
- 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 236](#).

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 236](#).

```
# 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 Support for Oracle RAC database server.**

Set a strong positive affinity by the scalable resource group for the Support for 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 231.
- 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 236.

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-fwk-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 Support for 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 231](#).
- 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 236](#).

```
# clresource create -g rac-db-rg \  
-t SUNW.scalable_rac_server_proxy \  
-p resource_dependencies=rac-fwk-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 \  
rac-srvr-proxy-rs
```

7. Bring online the resource group that you created in [Step 4](#).

```
# clresourcegroup online -M 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 Support for 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 244](#).

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. Assume the root role 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 Support for Oracle RAC Framework Resource Group” on page 63.](#)

Set a strong dependency by the instance of `SUNW.crs_framework` on the instance of `SUNW.rac_framework` in the Support for 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 231.](#)
- 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 Zone Cluster” on page 238.](#)

You might have configured a storage resource for the file system that you are using for Oracle Clusterware 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 Zone Cluster” on page 238.](#)

```
# 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 Support for Oracle RAC database server.

Set a strong positive affinity by the scalable resource group for the Support for 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 231.
- 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 Zone Cluster”](#) on page 238.

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-fwk-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 Support for 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 231.

- 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 Zone Cluster”](#) on page 238.

```
# clresource create -Z zcname -g rac-db-rg \  
-t SUNW.scalable_rac_server_proxy \  
-p resource_dependencies=rac-fwk-rs \  
-p resource_dependencies_offline_restart=crs-fwk-rs \  
[, db-storage-rs, bin-storage-rs] \  
-p oracle_home=ora-home \  
-p crs_home=Grid_home \  
-p db_name=db-name \  
rac-srvr-proxy-rs
```

7. Bring online the resource group that you created in [Step 4](#).

```
# clresourcegroup online -Z zcname -M rac-db-rg
```

Index

A

abort step timeout

Solaris Volume Manager for Sun Cluster, 215

acfs_mountpoint extension property, 202

actions

file-system mount-point fault monitor, 119

preset for fault monitor, 193

scalable device-group fault monitor, 118

adding

resources to multiple-owner volume-manager

framework resource group, 151

resources to Support for Oracle RAC framework

resource group, 151

Support for Oracle RAC to nodes, 145

volume manager resources, 151

volumes to be monitored, 144

administering

Support for Oracle RAC, 111

affinities

examples, 180

file-system mount-point resource groups, 237

scalable device-group resource groups, 231, 232

StorageTek QFS resource groups, 234

Support for Oracle RAC server resource

groups, 245, 248

alert files

in cluster file system, 60

in Oracle ACFS file system, 60

file system options

StorageTek QFS shared file system, 54

UNIX file system, 62

storage management schemes, 25

in StorageTek QFS shared file system, 52

archived redo log files

in cluster file system, 60

in Oracle ACFS file system, 61

file system options

StorageTek QFS shared file system, 54

UNIX file system, 62

optimum I/O performance, 61

storage management schemes, 25

in StorageTek QFS shared file system, 52

arrays

disk, 49, 49

See also redundant array of independent disks

(RAID)

ASM *See* Oracle Automatic Storage Management (Oracle ASM)

asm_diskgroups extension property

description, 203, 205

ASM_DISKSTRING parameter, 59

automating

database startup and shutdown

clsetup utility, 98

with Oracle Solaris Cluster maintenance

commands, 244, 247

B

binary files

in cluster file system, 60

in Oracle ACFS file system, 60

determining location, 27

file system options

StorageTek QFS shared file system, 54

UNIX file system, 62

storage management schemes, 25, 26

in StorageTek QFS shared file system, 52

boot
 verifying, 109

C

caution notice

- adding volume manager resources to framework resource group, 152
- restriction for Oracle ASM with Solaris Volume Manager, 27
- Support for Oracle RAC verification, 110

changing *See* modifying

checking *See* verifying

CLI, 64 *See* Oracle Solaris Cluster maintenance

commands

See also `clsetup` utility

`Client_retries` extension property, 206

`Client_retry_interval` extension property, 207

`clsetup` utility

- comparison with Oracle Solaris Cluster maintenance commands, 64, 68, 79, 97
- database resources, 98, 98
- multiple-owner volume-manager framework resources, 68
- names of resources created by, 112
- Oracle ASM resources, 84
- Oracle Grid Infrastructure resources, 93
- proxy resources, 98
- storage resources, 79
- Support for Oracle RAC framework resources, 64

cluster file system

- creating, 59, 61
- installation of Oracle Database files, 28
- mounting, 59
- Oracle Grid Infrastructure file types supported, 26
- Oracle RAC file types supported, 25
- requirements, 60
- task summary, 47

Cluster Ready Services *See* Oracle Grid Infrastructure

command-line interface, 64 *See* Oracle Solaris Cluster maintenance commands

See also `clsetup` utility

commands

 license verification, 24

components

 validation failures, 134, 136

concatenating

 slices, 73

configuration files

 in cluster file system, 60

 in Oracle ACFS file system, 60

 determining location, 27

 file system options

 StorageTek QFS shared file system, 54

 UNIX file system, 62

 storage management schemes, 25

 in StorageTek QFS shared file system, 52

configuration information

 multiple-owner volume-manager framework resource group, 132

 UCMM, 132

configuring

 DID, 49, 57

 file-system mount-point resources

 with `clsetup` utility, 79

 with Oracle Solaris Cluster maintenance commands, 236

 hardware RAID, 49

 interoperation with Oracle Grid Infrastructure, 243

 multiple-owner volume-manager framework

 resource group, 67

 for clusters, 68, 221

 Oracle ASM, 56

 resource group, 83, 84

 qualified NAS devices, 48

 resources after Oracle Solaris Cluster upgrade, 177

 scalable device-group resources

 with `clsetup` utility, 79

 with Oracle Solaris Cluster maintenance commands, 231, 232

 Solaris Volume Manager for Sun Cluster, 49

 StorageTek QFS resources

 with `clsetup` utility, 79

 with Oracle Solaris Cluster maintenance commands, 233

- with Oracle Solaris Cluster maintenance commands for a zone cluster, 235
- Support for Oracle RAC
 - examples, 179
 - for selected nodes, 145
- Support for Oracle RAC framework resource group
 - for clusters, 64, 221
 - for zone clusters, 225
 - overview, 63
 - planning, 220
- Support for Oracle RAC server
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 244, 247
- UFS, 62
- control files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 54
 - storage management schemes, 25
 - in StorageTek QFS shared file system, 52
- core file fault monitors, 120
- creating
 - cluster file system, 59
 - device groups, 71
 - file-system mount-point resources
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 236
 - global device groups, 71
 - LUN, 49
 - multi-owner disk sets, 71
 - multiple-owner volume-manager framework resource group
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 221
 - Oracle ASM instance, 92
 - Oracle ASM resources
 - with `clsetup` utility, 84
 - Oracle Grid Infrastructure resources
 - with `clsetup` utility, 93
 - proxy resources
 - with `clsetup` utility, 98
 - scalable device-group resources
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 231, 232
 - slices, 49
 - StorageTek QFS resources
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 233
 - with Oracle Solaris Cluster maintenance commands for a zone cluster, 235
 - StorageTek QFS shared file system, 54
 - Support for Oracle RAC framework resource group
 - with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 221
 - with Oracle Solaris Cluster maintenance commands in a zone cluster, 225
 - Support for Oracle RAC server resource group
 - with `clsetup` utility, 98, 98
 - with Oracle Solaris Cluster maintenance commands, 245, 248
 - Support for Oracle RAC server resources
 - with `clsetup` utility, 98, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
- CRS *See* Oracle Grid Infrastructure
- `crs_framework` resource type
 - dependencies, 101, 245, 248
 - extension properties, 201
 - instantiating
 - with `clsetup` utility, 93
 - with Oracle Solaris Cluster maintenance commands, 245, 248
 - names of instances created by `clsetup`, 112
 - registering
 - with `clsetup` utility, 93
 - with Oracle Solaris Cluster maintenance commands, 245, 248
- `crs_home` extension property, 207
 - description, 205
- cylinders

- restrictions, 50
- D**
- daemons
 - ucmmd
 - failure to start, 134
 - unexpected termination, 133
 - vucmmd
 - failure to start, 136
 - unexpected termination, 136
- data files
 - file system options, 54
 - location on shared file system, 96
 - storage management schemes, 25
- Data Guard *See* Oracle Data Guard
- database administrator (DBA)
 - creating, 31
 - granting access to file systems, 55
 - granting access to volumes, 74
- database management system (DBMS)
 - errors
 - preset actions, 193
 - timeouts, 120
- database resources
 - creating with `clsetup` utility, 98
 - extension properties, 206
- databases
 - automating startup and shutdown
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 244, 247
 - creating, 96
 - file systems for, 52
 - resources
 - creating with `clsetup` utility, 97
 - creating with Oracle Solaris Cluster maintenance commands, 240
 - verifying resources, 107
- DAU (disk allocation unit), 54
- `db_name` extension property, 207
- DBA (database administrator)
 - creating, 31
 - granting access to file systems, 55
 - granting access to volumes, 74
- dbca command, 96
- DBMS (database management system)
 - errors
 - preset actions, 193
 - timeouts, 120
- Debug_level extension property
 - description, 204
 - scalable_asm_instance_proxy resource type, 202, 205
 - scalable_rac_server_proxy resource type, 207
 - ScalDeviceGroup resource type, 209
 - ScalMountPoint resource type, 211
 - SUNW.scalable_asm_instance_proxy resource type, 202, 205
 - SUNW.scalable_rac_server_proxy resource type, 207
 - SUNW.ScalDeviceGroup resource type, 209
 - SUNW.ScalMountPoint resource type, 211
 - SUNW.vucmm_svm resource type, 214
- Degraded - reconfiguration in progress message, 140, 141
- deleting *See* modifying *See* removing *See* uninstalling
- dependencies
 - crs_framework resource type, 101, 245, 248
 - offline-restart, 103
 - oracle_asm_diskgroup resource type, 87
 - qfs resource type, 82, 87, 234
 - rac_framework resource type, 67
 - scalable_oracle_asm_instance_proxy resource type, 87
 - scalable_rac_server_proxy resource type, 86, 102, 246, 249
 - ScalDeviceGroup resource type, 82, 88, 232, 233
 - ScalMountPoint resource type, 82, 82, 87, 237
 - SUNW.crs_framework resource type, 101, 245, 248
 - SUNW.oracle_asm_diskgroup resource type, 87
 - SUNW.qfs resource type, 82, 87, 234
 - SUNW.rac_framework resource type, 67
 - SUNW.scalable_oracle_asm_instance_proxy resource type, 87

-
- SUNW.scalable_rac_server_proxy resource type, 86, 102, 246, 249
 - SUNW.ScalDeviceGroup resource type, 82, 88, 232, 233
 - SUNW.ScalMountPoint resource type, 82, 82, 87, 237
 - SUNW.vucmm_framework resource type, 70
 - SUNW.vucmm_svm resource type, 70, 224
 - vucmm_framework resource type, 70
 - vucmm_svm resource type, 70, 224
 - device groups, 71, 71
 - See also* multi-owner disk sets
 - See also* shared-disk groups
 - creating, 71
 - extension properties, 209
 - monitoring additional volumes, 144
 - optimum I/O performance, 61
 - resource group affinities, 231, 232
 - storage resources for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - device identity (DID)
 - configuring, 49, 57
 - devices
 - adding to multi-owner disk sets, 72
 - raw, 49, 49
 - See also* redundant array of independent disks (RAID)
 - removing from multi-owner disk sets, 158
 - diagnostic information, 132
 - DID (device identity)
 - configuring, 49, 57
 - directories
 - Oracle Database home, 206, 208
 - /var/opt, 34
 - disabling
 - Oracle RAC server resources, 112
 - disk allocation unit (DAU), 54
 - disk groups
 - Oracle ASM, 203, 205
 - disk sets
 - multi-owner, 71
 - Diskgroupname extension property, 210
 - disks
 - arrays, 49, 49
 - See also* redundant array of independent disks (RAID)
 - effect of number on reservation timeout, 116
 - installation of Oracle Database files, 28
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - slices, 49
 - soft partitions, 73
 - stripe width, 54
 - E**
 - editing *See* modifying
 - enabling, 21
 - See also* starting
 - Support for Oracle RAC server resource group, 247, 250
 - environment variables, 209
 - errors
 - DBMS
 - preset actions for, 193
 - timeouts
 - core-file creation, 120
 - /etc/group file, 32
 - /etc/opt/SUNWsamfs/samfs.cmd file, 54
 - /etc/passwd file, 32
 - /etc/shadow file, 32
 - /etc/system file, 36
 - /etc/vfstab file
 - StorageTek QFS, 54
 - UNIX file system, 62
 - /etc/nsswitch.conf file, 30
 - event logs, 132
 - examples
 - configuration, 179
 - executing *See* starting
 - extension properties
 - crs_framework resource type, 201
 - hardware RAID, 202

- Oracle ASM, 202, 202, 203
 - qualified network-attached storage (NAS) devices, 202
 - rac_framework resource type, 202
 - RAID, 202
 - scalable_asm_diskgroup_proxy resource type, 202, 203
 - scalable_asm_instance_proxy resource type, 205
 - scalable_rac_server_proxy resource type, 206
 - ScalDeviceGroup resource type, 209
 - ScalMountPoint resource type, 211
 - setting, 219
 - Solaris Volume Manager for Sun Cluster, 214
 - StorageTek QFS shared file system, 202
 - SUNW.crs_framework resource type, 201
 - SUNW.rac_framework resource type, 202
 - SUNW.scalable_asm_diskgroup_proxy resource type, 202, 203
 - SUNW.scalable_asm_instance_proxy resource type, 205
 - SUNW.scalable_rac_server_proxy resource type, 206
 - SUNW.ScalDeviceGroup resource type, 209
 - SUNW.ScalMountPoint resource type, 211
 - SUNW.vucmm_framework resource type, 214
 - SUNW.vucmm_svm resource type, 214
 - SUNW.wait_zc_boot resource type, 217
 - vucmm_framework resource type, 214
 - vucmm_svm resource type, 214
 - wait_zc_boot resource type, 217
- F**
- Failfast: Aborting because "ucmmd" died message, 133
 - Failfast: Aborting because "vucmmd" died message, 136
 - failures
 - component validations, 134, 136
 - panics
 - during initialization, 133, 136
 - timeouts, 138
 - public networks, 30
 - rac_framework resource
 - startup, 139
 - reconfiguration of resources, 140, 141
 - START method timeout, 140, 141
 - startup of resources, 139
 - stopping of resources, 142
 - SUNW.rac_framework resource, 139
 - SUNW.vucmm_framework resource, 139
 - ucmmd daemon
 - startup, 134
 - unexpected termination, 133
 - vucmm_framework resource
 - startup, 139
 - vucmmd daemon
 - startup, 136
 - unexpected termination, 136
 - fault monitoring, 22, 22
 - fault monitors
 - actions
 - file-system mount-point fault monitor, 119
 - scalable device-group fault monitor, 118
 - core-file creation, 120
 - device groups
 - resource type for, 118
 - file systems
 - resource type for, 118
 - mount points
 - resource type for, 118
 - preset actions, 193
 - tuning, 117
 - Faulted - ucmmd is not running message, 139
 - file systems
 - extension properties, 211
 - installation of Oracle Database files, 28
 - restrictions, 72
 - storage resources for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - UFS options, 62

- file-system mount points
 - resource groups for
 - affinities, 237
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 236
 - planning, 230
 - resources
 - state information, 132
 - `syslog()` function, 132
 - resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 236
 - planning, 230
 - file-system mount-point fault monitor, 119
 - files
 - core fault monitors, 120
 - determining location, 25
 - diagnostic information, 132
 - `/etc/group`, 32
 - `/etc/opt/SUNWsamfs/samfs.cmd`, 54
 - `/etc/passwd`, 32
 - `/etc/shadow`, 32
 - `/etc/system`, 36
 - `/etc/vfstab`
 - StorageTek QFS, 54
 - UNIX file system, 62
 - `/etc/nsswitch.conf`, 30
 - group, 32
 - `nsswitch.conf`, 30
 - Oracle Database configuration, 27
 - `passwd`, 32
 - `samfs.cmd`, 54
 - `shadow`, 32
 - storage management schemes for, 25
 - Support for Oracle RAC, 52, 60, 60
 - system, 36
 - `vfstab`
 - StorageTek QFS, 54
 - UNIX file system, 62
 - Filesystemtype extension property, 212
 - flashback log files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 62
 - storage management schemes, 25
 - in StorageTek QFS shared file system, 52
 - framework *See* Oracle Real Application Clusters (Oracle RAC) framework resource group
- ## G
- global device groups, 71, 71
 - See also* multi-owner disk sets
 - See also* shared-disk groups
 - creating, 71
 - monitoring additional volumes, 144
 - storage resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - global devices
 - adding to multi-owner disk sets, 72
 - removing from multi-owner disk sets, 158
 - group database
 - `nsswitch.conf` file, 30
 - group file, 32
 - groups
 - creating, 31
- ## H
- hardware redundant array of independent disks (RAID)
 - configuring, 49
 - extension properties, 202
 - installing, 49
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - reservation step timeout, 116
 - task summary, 45
 - hardware requirements, 23
 - home directory
 - Oracle Database, 206, 208

I

- I/O (input/output) performance, 61
- identifiers
 - user, 30
- input/output (I/O) performance, 61
- installing
 - hardware RAID, 49
 - Oracle ASM, 56
 - Oracle RAC
 - overview, 90
 - verifying installation, 91
 - qualified NAS devices, 48
 - Solaris Volume Manager for Sun Cluster, 49
 - storage management software, 47
 - Support for Oracle RAC, 39
 - verifying installation, 103
- interfaces
 - network, 91
- Iotimeout extension property, 212

L

- license requirements, 24
- load
 - effect on reservation timeout, 116
- local disks
 - installation of Oracle Database files, 28
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
- location
 - diagnostic information files, 132
 - log files, 132
 - Oracle Database files, 25
 - redo log files, 61
- location of data files, 96
- log files, 132
 - locations, 132
 - Support for Oracle RAC server proxy, 207
 - use in troubleshooting, 132
- logical unit number (LUN)
 - creating, 49
- logical-hostname resources
 - zone clusters, 38

- Logicaldevicelist extension property, 210
- logicaldevicelist extension property, 233
- LogicalHostname resource type
 - names of instances created by clsetup, 112
- LUN (logical unit number)
 - creating, 49

M

- memory
 - shared, 35, 36
- messages
 - debug, 202, 205, 207
 - panics, 133, 136
 - startup failure, 139
- metadata server
 - resource groups for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - planning, 231
 - resources for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - planning, 231
- modifying
 - extension properties
 - command for, 115
 - Support for Oracle RAC
 - by removing nodes from, 162
- Monitor_probe_interval extension property, 208
- Monitor_retry_count extension property
 - ScalDeviceGroup resource type, 211
 - ScalMountPoint resource type, 212
 - SUNW.ScalDeviceGroup resource type, 211
 - SUNW.ScalMountPoint resource type, 212
- Monitor_retry_interval extension property
 - ScalDeviceGroup resource type, 211
 - ScalMountPoint resource type, 212
 - SUNW.ScalDeviceGroup resource type, 211
 - SUNW.ScalMountPoint resource type, 212

- monitoring
 - Support for Oracle RAC, 121
 - mount options
 - UFS, 62
 - mount points
 - extension properties, 211
 - mounting
 - cluster file system, 59
 - StorageTek QFS shared file system, 54
 - Mountoptions extension property, 213
 - Mountpointdir extension property, 213
 - multi-owner disk sets
 - creating, 71
 - Oracle Automatic Storage Management (Oracle ASM) (Oracle ASM), 75
 - storage resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - multiple-owner volume-manager framework resource group
 - adding resources to, 151
 - configuration information, 132
 - configuring, 67
 - configuring for clusters, 68
 - creating
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 221
 - failure to start, 136
 - overview, 221
 - unexpected termination, 136
 - verifying, 104
- N**
- name service
 - bypassing, 30
 - database user entries in, 32
 - NAS devices *See* qualified network-attached storage (NAS) devices
 - necessary privileges
 - zone clusters, 38
 - Network Information Service (NIS)
 - bypassing, 30
 - database user entries in, 32
 - network interfaces, 91
 - networks
 - private, 91
 - public
 - failures, 30
 - installation options, 91
 - NIS (Network Information Service)
 - bypassing, 30
 - database user entries in, 32
 - nodes
 - adding Support for Oracle RAC to, 145
 - failure
 - public networks, 30
 - panics
 - during initialization, 133, 136
 - timeouts, 138
 - preparing for Support for Oracle RAC, 29
 - removing Support for Oracle RAC from, 162
 - uninstalling
 - Support for Oracle RAC software packages from, 168
 - nsswitch.conf file, 30
 - NTP
 - preventing Oracle Clusterware time synchronization in zone clusters, 38
- O**
- OCR (Oracle cluster registry) files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 62
 - storage management schemes, 26
 - in StorageTek QFS shared file system, 52
 - offline-restart dependencies, 103
 - removing from Oracle Grid Infrastructure resource, 153
 - oinstall group, 31
 - Online message, 140, 141

- online redo log files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 54
 - storage management schemes, 25
 - in StorageTek QFS shared file system, 52
- oper group, 31
- Oracle
 - error numbers, 193
- Oracle ACFS
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
- Oracle ACFS file system
 - creating, 62
 - requirements, 60
 - task summary, 47
- Oracle ASM
 - device isolation, 56
 - restriction with Solaris Volume Manager, 27
 - software requirements, 24
- Oracle Automatic Storage Management (Oracle ASM)
 - configuring, 56
 - creating an instance, 92
 - device groups, 75
 - disk groups, 203, 205
 - extension properties, 202, 202, 203
 - installing, 56
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - reservation step timeout, 116
 - resource group
 - configuring, 83
 - creating with `clsetup` utility, 84
 - task summary, 45
 - zone clusters, 57
- Oracle cluster registry (OCR) files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 62
 - storage management schemes, 26
 - in StorageTek QFS shared file system, 52
- Oracle Clusterware
 - file system options
 - StorageTek QFS shared file system, 54
 - preventing time synchronization in zone clusters, 38
- Oracle Clusterware framework resource method timeout
 - guidelines for setting, 117
- Oracle Data Guard, 29
- Oracle Database binary files
 - location, 27
- Oracle Database configuration files
 - location, 27
- Oracle Database files
 - binary file location, 27
 - configuration file location, 27
 - disks, 28
 - local disks, 28
 - storage management schemes for, 25
 - storage resources for
 - creating with `clsetup` utility, 79, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - StorageTek QFS shared file system, 28
- Oracle Database RDBMS (relational database management system)
 - file system options
 - StorageTek QFS shared file system, 54
- Oracle Database relational database management system (RDBMS)
 - file system options
 - StorageTek QFS shared file system, 54
- Oracle Grid Infrastructure
 - configuring interoperation, 243
 - extension properties, 201
 - overriding networking defaults, 91
 - re-creating `sun.storage_proxy.typ` resources after upgrade, 174
 - removing a resource, 153
 - removing offline-restart dependency, 153
 - removing `sun.resource`, 154
 - resources
 - creating with `clsetup` utility, 93
 - state changes, 113

- software requirements, 24
- storage management schemes, 26
- Oracle RDBMS (relational database management system)
 - file system options
 - UNIX file system, 62
 - storage management schemes, 25
- Oracle Real Application Clusters (Oracle RAC)
 - installing
 - overview, 90
 - verifying installation, 91
 - verifying installation and configuration, 91
- Oracle Real Application Clusters (Oracle RAC) framework resource group
 - configuring
 - for clusters, 220
 - creating
 - with Oracle Solaris Cluster maintenance commands, 221
 - with Oracle Solaris Cluster maintenance commands in a zone cluster, 225
 - overview, 220
 - planning, 220
 - upgrading resources in, 173
- Oracle Real Application Clusters (Oracle RAC) server
 - configuring
 - with Oracle Solaris Cluster maintenance commands, 247
 - configuring with `clsetup` utility, 98, 98
 - resource groups
 - creating with `clsetup` utility, 98, 98
 - creating with Oracle Solaris Cluster maintenance commands, 245
 - resources
 - creating with `clsetup` utility, 98, 98
 - disabling, 113
- Oracle Real Application Clusters (Support for Oracle RAC)
 - preparing nodes for, 29
- Oracle Real Application Clusters (Support for Oracle RAC) framework resource group
 - adding resources to, 151
 - configuring
 - for clusters, 64
 - overview, 63
 - creating
 - with `clsetup` utility, 64
 - verifying, 104
- Oracle Real Application Clusters (Support for Oracle RAC) server
 - configuring
 - with Oracle Solaris Cluster maintenance commands, 244
 - resource groups
 - creating with Oracle Solaris Cluster maintenance commands, 248
 - enabling, 247, 250
 - resources
 - creating with Oracle Solaris Cluster maintenance commands, 246, 249
 - enabling, 247, 250
- Oracle relational database management system (RDBMS)
 - file system options
 - UNIX file system, 62
 - storage management schemes, 25
- Oracle Solaris Cluster
 - framework, 23
- Oracle Solaris Cluster maintenance commands
 - comparison with `clsetup` utility, 64, 68, 79, 97
 - multiple-owner volume-manager framework resource group
 - creating, 221
 - storage resources, 230, 231, 232
 - Support for Oracle RAC framework resource group
 - creating, 221, 225
 - tuning extension properties with, 115
- Oracle Solaris Cluster software
 - publisher, 40, 40
- Oracle Solaris Cluster Support for Oracle Real Application Clusters (Support for Oracle RAC)
 - files related to, 52
 - overview, 22
- `oracle` user, 31
 - granting access to file systems, 55
 - granting access to volumes, 74

- oracle_asm_diskgroup resource type
 - dependencies, 87
 - names of instances created by clsetup, 112
- Oracle_home extension property
 - scalable_asm_instance_proxy resource type, 206
 - scalable_rac_server_proxy resource type, 208
 - SUNW.scalable_asm_instance_proxy resource type, 206
 - SUNW.scalable_rac_server_proxy resource type, 208
- overriding networking defaults, 91
- overview
 - Support for Oracle RAC, 22
 - Support for Oracle RAC framework resource group, 220

P

- package, 39
- panics
 - during initialization, 133, 136
 - timeouts, 138
- partitions
 - restrictions, 50
 - soft, 73
- passwd file, 32
- prerequisites *See* requirements
- preset actions
 - fault monitor, 193
- primary groups, 31
- private network interfaces, 91
- projects
 - prerequisite, 84
- proxy resources
 - creating with clsetup utility, 98
 - examples, 241, 242
- Proxy_probe_interval extension property, 203
 - description, 204
- Proxy_probe_timeout extension property, 203, 206, 208
 - description, 204
- public networks

- failure, 30
- installation options, 91
- publisher
 - Oracle Solaris Cluster software, 40, 40

Q

- QFS file system *See* StorageTek QFS shared file system
- qfs resource type
 - dependencies, 82, 87, 234
 - instantiating
 - with clsetup utility, 79
 - with Oracle Solaris Cluster maintenance commands, 234
 - names of instances created by clsetup, 112
 - registering
 - with clsetup utility, 79
 - with Oracle Solaris Cluster maintenance commands, 234
- qualified network-attached storage (NAS) devices
 - configuring, 48
 - extension properties, 202
 - installing, 48
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - storage resources for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - supported with Support for Oracle RAC, 78
 - task summary, 46

R

- RAC *See* Oracle Solaris Cluster Support for Oracle Real Application Clusters (Support for Oracle RAC))
- rac_framework resource type
 - dependencies, 67
 - extension properties, 202
 - failure of instance to start, 139
 - instantiating

- with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 225
- monitoring instances of, 121
- names of instances created by `clsetup`, 112
- purpose, 220
- registering
 - with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 225
- START method timeout, 140
- upgrading, 173
- `rac_svm` resource type
 - upgrading, 173
- RAID (redundant array of independent disks)
 - configuring, 49
 - extension properties, 202
 - installing, 49
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - reservation step timeout, 116
 - task summary, 45
- raw devices, 49, 49
 - See also* redundant array of independent disks (RAID)
- RDBMS (relational database management system)
 - file system options
 - StorageTek QFS shared file system, 54
 - UNIX file system, 62
 - storage management schemes, 25
- reconfiguration failures
 - `SUNW.rac_framework`, 140
 - `SUNW.vucmm_framework`, 141
- reconfiguration timeouts
 - Oracle ASM, 202
 - reservation step, 202, 214
 - Solaris Volume Manager for Sun Cluster definition, 215
- recovery files
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - storage management schemes, 26
 - in StorageTek QFS shared file system, 52
- redo log files *See* archived redo log files *See* online redo log files
- redundant array of independent disks (RAID)
 - configuring, 49
 - extension properties, 202
 - installing, 49
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - reservation step timeout, 116
 - task summary, 45
- registering
 - `crs_framework` resource type
 - with `clsetup` utility, 93
 - with Oracle Solaris Cluster maintenance commands, 245, 248
 - `qfs` resource type
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 234
 - `rac_framework` resource type
 - with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 222, 225, 225
 - `scalable_rac_server_proxy` resource type
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
 - `ScalDeviceGroup` resource type
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
 - `ScalMountPoint` resource type
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
 - `SUNW.crs_framework` resource type
 - with `clsetup` utility, 93
 - with Oracle Solaris Cluster maintenance commands, 245, 248
 - `SUNW.qfs` resource type
 - with `clsetup` utility, 79

- with Oracle Solaris Cluster maintenance commands, 234
- SUNW.rac_framework resource type
 - with clsetup utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 225
- SUNW.scalable_rac_server_proxy resource type
 - with clsetup utility, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
- SUNW.ScalDeviceGroup resource type
 - with clsetup utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
- SUNW.ScalMountPoint resource type
 - with clsetup utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
- SUNW.vucmm_framework resource type
 - with clsetup utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
- SUNW.vucmm_svm resource type
 - with clsetup utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
- SUNW.wait_zc_boot resource type
 - with Oracle Solaris Cluster maintenance commands, 235
- vucmm_framework resource type
 - with clsetup utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
- vucmm_svm resource type
 - with clsetup utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
- relational database management system (RDBMS)
 - file system options
 - StorageTek QFS shared file system, 54
 - UNIX file system, 62
 - storage management schemes, 25
- removing, 159
 - See also* uninstalling
 - offline-restart dependency, 153
 - Oracle Grid Infrastructure resource, 153
 - Oracle Grid Infrastructure sun.resource, 154
 - Support for Oracle RAC
 - from clusters, 155
 - from selected nodes, 162
- requirements
 - hardware, 23
 - multi-owner disk sets, 71
 - Oracle Database files, 25
 - software, 23
- reservation step timeout
 - description, 202, 214
 - guidelines for setting, 116
- Reservation_timeout extension property
 - description, 202, 214
 - guidelines for setting, 116
- resource group affinities, 245, 248
- Resource Group Manager (RGM)
 - restrictions, 220
- resource groups
 - file-system mount point
 - affinities, 237
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 236
 - planning, 230
 - multiple-owner volume-manager framework
 - creating with clsetup utility, 68
 - creating with Oracle Solaris Cluster maintenance commands, 221
 - overview, 221
 - Oracle 10g RAC server
 - creating with clsetup utility, 98
 - Oracle ASM
 - creating with clsetup utility, 84
 - sample configurations, 179
 - scalable device group
 - affinities, 231, 232
 - creating with clsetup utility, 79

- creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
- StorageTek QFS metadata server
 - affinities, 234
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - creating with Oracle Solaris Cluster maintenance commands for a zone cluster, 235
 - planning, 231
- Support for Oracle RAC framework
 - creating with `clsetup` utility, 64
 - creating with Oracle Solaris Cluster maintenance commands, 221
 - planning, 220
- Support for Oracle RAC server
 - affinities, 245, 248
 - creating with `clsetup` utility, 98
 - creating with Oracle Solaris Cluster maintenance commands, 245, 248
 - enabling, 247, 250
- resource types
 - crs_framework
 - dependencies, 101, 245, 248
 - extension properties, 201
 - instantiating with `clsetup` utility, 93
 - instantiating with Oracle Solaris Cluster maintenance commands, 245, 248
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 93
 - registering with Oracle Solaris Cluster maintenance commands, 245, 248
 - fault monitors, 118
 - LogicalHostname
 - names of instances created by `clsetup`, 112
 - names of instances created by `clsetup`, 112
 - oracle_asm_diskgroup
 - dependencies, 87
 - names of instances created by `clsetup`, 112
 - qfs
 - dependencies, 82, 87, 234
 - instantiating with `clsetup` utility, 79
 - instantiating with Oracle Solaris Cluster maintenance commands, 234
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 79
 - registering with Oracle Solaris Cluster maintenance commands, 234
 - rac_framework
 - dependencies, 67
 - extension properties, 202
 - failure of instance to start, 139
 - instantiating with `clsetup` utility, 64
 - instantiating with Oracle Solaris Cluster maintenance commands, 222, 225
 - monitoring instances of, 121
 - names of instances created by `clsetup`, 112
 - purpose, 220
 - registering with `clsetup` utility, 64
 - registering with Oracle Solaris Cluster maintenance commands, 222, 225
 - START method timeout, 140
 - upgrading, 173
 - rac_svm
 - upgrading, 173
 - restrictions, 220
 - sample configurations, 179
 - scalable_asm_diskgroup_proxy
 - extension properties, 202, 203
 - scalable_asm_instance_proxy
 - extension properties, 205
 - scalable_oracle_asm_instance_proxy
 - dependencies, 87
 - names of instances created by `clsetup`, 112
 - scalable_rac_server_proxy
 - dependencies, 86, 102, 246, 249
 - extension properties, 206
 - instantiating with `clsetup` utility, 98
 - names of instances created by `clsetup`, 112, 112
 - registering with `clsetup` utility, 98
 - registering with Oracle Solaris Cluster maintenance commands, 246, 249

- ScalDeviceGroup
 - dependencies, 82, 88, 232, 233
 - extension properties, 209
 - instantiating with `clsetup` utility, 79
 - instantiating with Oracle Solaris Cluster maintenance commands, 232, 233
 - modifying instances of, 144
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 79
 - registering with Oracle Solaris Cluster maintenance commands, 232, 233
- ScalMountPoint
 - dependencies, 82, 82, 87, 237
 - extension properties, 211
 - instantiating with `clsetup` utility, 79
 - instantiating with Oracle Solaris Cluster maintenance commands, 237
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 79
 - registering with Oracle Solaris Cluster maintenance commands, 237
- SUNW.crs_framework
 - dependencies, 101, 245, 248
 - extension properties, 201
 - instantiating with `clsetup` utility, 93
 - instantiating with Oracle Solaris Cluster maintenance commands, 245, 248
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 93
 - registering with Oracle Solaris Cluster maintenance commands, 245, 248
- SUNW.LogicalHostname
 - names of instances created by `clsetup`, 112
- SUNW.oracle_asm_diskgroup
 - dependencies, 87
 - names of instances created by `clsetup`, 112
- SUNW.qfs
 - dependencies, 82, 87, 234
 - instantiating with `clsetup` utility, 79
 - instantiating with Oracle Solaris Cluster maintenance commands, 234
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 79
 - registering with Oracle Solaris Cluster maintenance commands, 234
- SUNW.rac_framework
 - dependencies, 67
 - extension properties, 202
 - failure of instance to start, 139
 - instantiating with `clsetup` utility, 64
 - instantiating with Oracle Solaris Cluster maintenance commands, 222, 225
 - monitoring instances of, 121
 - names of instances created by `clsetup`, 112
 - purpose, 220
 - registering with `clsetup` utility, 64
 - registering with Oracle Solaris Cluster maintenance commands, 222, 225
 - START method timeout, 140
 - upgrading, 173
- SUNW.rac_svm
 - upgrading, 173
- SUNW.scalable_asm_diskgroup_proxy
 - extension properties, 202, 203
- SUNW.scalable_asm_instance_proxy
 - extension properties, 205
- SUNW.scalable_oracle_asm_instance_proxy
 - dependencies, 87
 - names of instances created by `clsetup`, 112
- SUNW.scalable_rac_server_proxy
 - dependencies, 86, 102, 246, 249
 - extension properties, 206
 - instantiating with `clsetup` utility, 98
 - instantiating with Oracle Solaris Cluster maintenance commands, 246, 249
 - names of instances created by `clsetup`, 112, 112
 - registering with `clsetup` utility, 98
 - registering with Oracle Solaris Cluster maintenance commands, 246, 249
- SUNW.ScalDeviceGroup
 - dependencies, 82, 88, 232, 233
 - extension properties, 209
 - instantiating with `clsetup` utility, 79

- instantiating with Oracle Solaris Cluster maintenance commands, 232, 233
- modifying instances of, 144
- names of instances created by `clsetup`, 112
- registering with `clsetup` utility, 79
- registering with Oracle Solaris Cluster maintenance commands, 232, 233
- SUNW.ScalMountPoint
 - dependencies, 82, 82, 87, 237
 - extension properties, 211
 - instantiating with `clsetup` utility, 79
 - instantiating with Oracle Solaris Cluster maintenance commands, 237
 - names of instances created by `clsetup`, 112
 - registering with `clsetup` utility, 79
 - registering with Oracle Solaris Cluster maintenance commands, 237
- SUNW.vucmm_framework
 - dependencies, 70
 - extension properties, 214
 - failure of instance to start, 139
 - instantiating with `clsetup` utility, 68
 - instantiating with Oracle Solaris Cluster maintenance commands, 223
 - overview, 221
 - registering with `clsetup` utility, 68
 - registering with Oracle Solaris Cluster maintenance commands, 223
 - START method timeout, 141
- SUNW.vucmm_svm
 - dependencies, 70, 224
 - extension properties, 214
 - instantiating with `clsetup` utility, 68
 - instantiating with Oracle Solaris Cluster maintenance commands, 224
 - overview, 221
 - registering with `clsetup` utility, 68
 - registering with Oracle Solaris Cluster maintenance commands, 224
- SUNW.wait_zc_boot
 - extension properties, 217
- instantiating with Oracle Solaris Cluster maintenance commands, 235
- registering with Oracle Solaris Cluster maintenance commands, 235
- vucmm_framework
 - dependencies, 70
 - extension properties, 214
 - failure of instance to start, 139
 - instantiating with `clsetup` utility, 68
 - instantiating with Oracle Solaris Cluster maintenance commands, 223
 - registering with `clsetup` utility, 68
 - registering with Oracle Solaris Cluster maintenance commands, 223
 - START method timeout, 141
- vucmm_svm
 - dependencies, 70, 224
 - extension properties, 214
 - instantiating with `clsetup` utility, 68
 - instantiating with Oracle Solaris Cluster maintenance commands, 224
 - registering with `clsetup` utility, 68
 - registering with Oracle Solaris Cluster maintenance commands, 224
- wait_zc_boot
 - dependencies, 82
 - extension properties, 217
- resources
 - adding to multiple-owner volume-manager framework resource group, 151
 - adding to Support for Oracle RAC framework resource group, 151
 - database
 - creating with `clsetup` utility, 97
 - creating with Oracle Solaris Cluster maintenance commands, 240
 - file-system mount point
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 236
 - planning, 230
 - multiple-owner volume-manager framework resource group

- creating with `clsetup` utility, 68
 - creating with Oracle Solaris Cluster maintenance commands, 221
 - Oracle ASM resource group
 - creating with `clsetup` utility, 84
 - Oracle Grid Infrastructure
 - creating with `clsetup` utility, 93
 - removing, 153
 - state changes, 113
 - Oracle Solaris Cluster, 113
 - proxy
 - creating with `clsetup` utility, 98
 - examples, 241, 242
 - sample configurations, 179
 - scalable device group
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - storage
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 229, 231, 232
 - planning, 230
 - StorageTek QFS metadata server
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - planning, 231
 - Support for Oracle RAC framework resource group
 - creating with `clsetup` utility, 64
 - creating with Oracle Solaris Cluster maintenance commands, 221
 - planning, 220
 - Support for Oracle RAC server
 - creating with `clsetup` utility, 98, 98
 - creating with Oracle Solaris Cluster maintenance commands, 246, 249
 - disabling, 113
 - enabling, 247, 250
 - restrictions
 - cylinders, 50
 - file systems, 72
 - multi-owner disk sets, 71
 - partitions, 50
 - resource types, 220
 - RGM, 220
 - return step timeout
 - Solaris Volume Manager for Sun Cluster, 215
 - RGM (Resource Group Manager)
 - restrictions, 220
 - running *See* starting
- S**
- `samfs.cmd` file, 54
 - sample configurations, 179
 - sample configurations in a zone cluster, 186
 - sample configurations in the global cluster, 180
 - scalable device groups
 - fault monitor, 118
 - modifying resources, 144
 - resource group affinities, 231, 232
 - resource groups for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - state information, 132
 - `syslog()` function, 132
 - `scalable_asm_diskgroup_proxy` resource type
 - extension properties, 202, 203
 - `scalable_asm_instance_proxy` resource type
 - extension properties, 205
 - `scalable_oracle_asm_instance_proxy` resource type
 - dependencies, 87
 - names of instances created by `clsetup`, 112
 - `scalable_rac_server_proxy` resource type
 - dependencies, 86, 102, 246, 249
 - extension properties, 206
 - instantiating

- with `clsetup` utility, 98
- with Oracle Solaris Cluster maintenance commands, 246, 249
- names of instances created by `clsetup`, 112, 112
- registering
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
- `ScalDeviceGroup` resource type
 - dependencies, 82, 88, 232, 233
 - extension properties, 209
 - instantiating
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
 - modifying instances of, 144
 - names of instances created by `clsetup`, 112
 - registering
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
- `ScalMountPoint` resource type
 - dependencies, 82, 82, 87, 237
 - extension properties, 211
 - instantiating
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
 - names of instances created by `clsetup`, 112
 - registering
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
- secondary groups, 31
- server
 - extension properties, 206
 - resource groups
 - creating with `clsetup` utility, 98
 - creating with Oracle Solaris Cluster maintenance commands, 245, 248
 - enabling, 247, 250
 - resource groups for Oracle 10g
 - creating with `clsetup` utility, 98
 - resources
 - creating with `clsetup` utility, 98, 98
 - creating with Oracle Solaris Cluster maintenance commands, 246, 249
 - disabling, 113
 - enabling, 247, 250
 - server fault monitor
 - preset actions, 193
 - shadow file, 32
 - shared memory, 35
 - zone clusters, 36
 - shared-disk groups
 - Oracle Automatic Storage Management (Oracle ASM), 75
 - storage resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - shutdown
 - verifying, 109
 - shutting down
 - databases, 113
 - slices
 - concatenating, 73
 - disk, 49
 - soft partitions, 73
 - software package, 39
 - software requirements, 23
 - Oracle Grid Infrastructure, 24
 - software update requirements, 24
 - Solaris Volume Manager, 28
 - restriction with Oracle ASM, 27
 - Solaris Volume Manager for Sun Cluster
 - configuring, 49
 - extension properties, 214
 - installing, 49
 - multi-owner disk sets, 71
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - storage resources for
 - creating with `clsetup` utility, 79

- creating with Oracle Solaris Cluster maintenance commands, 231, 232
- planning, 230
- task summary, 43
- SPFILE file
 - in cluster file system, 60
 - in Oracle ACFS file system, 60
 - file system options, 62
 - storage management schemes, 25
 - in StorageTek QFS shared file system, 52
- sqlplus command, 96
- standard properties
 - effect on fault monitors, 118
- Start failed state, 139
- START method, 140, 141
- start step timeout
 - Solaris Volume Manager for Sun Cluster, 215
- starting
 - databases, 113
 - failure by resources, 139
- Startup_wait_count extension property, 208
- state information
 - database resources, 112
 - file-system mount-point resources, 132
 - scalable device-group resources, 132
- status information
 - Support for Oracle RAC, 121
- stopping
 - failure by resources, 142
- storage management schemes
 - choosing, 25
 - installing software for, 47
 - zone clusters, 28
- storage resources
 - configuring after Oracle Solaris Cluster upgrade, 176
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 229, 231, 232
 - planning, 230
 - supported NAS devices, 78
 - verifying, 105
- StorageTek QFS metadata server
 - resource groups for
 - affinities, 234
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - planning, 231
 - resources for
 - creating with clsetup utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 233
 - planning, 231
- StorageTek QFS shared file system
 - creating, 54
 - extension properties, 202
 - installation of Oracle Database files, 28
 - mounting, 54
 - Oracle Grid Infrastructure file types supported, 26
 - Oracle RAC file types supported, 25
 - requirements, 52
 - storage resources for
 - creating with clsetup utility, 79, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - task summary, 41
- stripes, 54
 - adding to volume, 73
- sun.resource removing, 154
- sun.storage_proxy.type resource type
 - re-creating after upgrade, 174
- SUNW.crs_framework resource
 - setting method timeouts, 117
- SUNW.crs_framework resource type
 - dependencies, 101, 245, 248
 - extension properties, 201
 - instantiating
 - with clsetup utility, 93
 - with Oracle Solaris Cluster maintenance commands, 245, 248
 - names of instances created by clsetup, 112
 - registering
 - with clsetup utility, 93

- with Oracle Solaris Cluster maintenance commands, 245, 248
- SUNW.LogicalHostname resource type
 - names of instances created by `clsetup`, 112
- SUNW.oracle_asm_diskgroup resource type
 - dependencies, 87
 - names of instances created by `clsetup`, 112
- SUNW.qfs resource type
 - dependencies, 82, 87, 234
 - instantiating
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 234
 - names of instances created by `clsetup`, 112
 - registering
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 234
- SUNW.rac_framework resource type
 - dependencies, 67
 - extension properties, 202
 - failure of instance to start, 139
 - instantiating
 - with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 225
 - monitoring instances of, 121
 - names of instances created by `clsetup`, 112
 - purpose, 220
 - registering
 - with `clsetup` utility, 64
 - with Oracle Solaris Cluster maintenance commands, 222, 225
 - START method timeout, 140
 - upgrading, 173
- SUNW.rac_svm resource type
 - upgrading, 173
- SUNW.scalable_asm_diskgroup_proxy resource type
 - extension properties, 202, 203
- SUNW.scalable_asm_instance_proxy resource type
 - extension properties, 205
- SUNW.scalable_oracle_asm_instance_proxy resource type
 - dependencies, 87
 - names of instances created by `clsetup`, 112
- SUNW.scalable_rac_server_proxy resource type
 - dependencies, 86, 102, 246, 249
 - extension properties, 206
 - instantiating
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
 - names of instances created by `clsetup`, 112, 112
 - registering
 - with `clsetup` utility, 98
 - with Oracle Solaris Cluster maintenance commands, 246, 249
- SUNW.ScalDeviceGroup resource type
 - dependencies, 82, 88, 232, 233
 - extension properties, 209
 - instantiating
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
 - names of instances created by `clsetup`, 112
 - registering
 - modifying instances of, 144
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 232, 233
- SUNW.ScalMountPoint resource type
 - dependencies, 82, 82, 87, 237
 - extension properties, 211
 - instantiating
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
 - names of instances created by `clsetup`, 112
 - registering
 - with `clsetup` utility, 79
 - with Oracle Solaris Cluster maintenance commands, 237
- SUNW.vucmm_framework resource type

- dependencies, 70
- extension properties, 214
- failure of instance to start, 139
- instantiating
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
- overview, 221
- registering
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
- START method timeout, 141
- SUNW.vucmm_svm resource type
 - dependencies, 70, 224
 - extension properties, 214
 - instantiating
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
 - overview, 221
 - registering
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
- SUNW.wait_zc_boot resource type
 - extension properties, 217
 - instantiating
 - with Oracle Solaris Cluster maintenance commands, 235
 - registering
 - with Oracle Solaris Cluster maintenance commands, 235, 235
- Support for Oracle RAC
 - administering, 111
 - configuring
 - for selected nodes, 145
 - examples, 179
 - fault monitors, 117
 - installing, 39
 - verifying installation, 103
 - log file locations, 132
 - modifying
 - by removing nodes from, 162
 - monitoring, 121
 - multi-owner disk sets for, 71
 - Oracle Data Guard, 29
 - overview, 22
 - removing
 - from clusters, 155
 - from selected nodes, 162
 - sample configurations, 179
 - software package, installing, 39
 - software packages
 - uninstalling from clusters, 159
 - uninstalling from selected nodes, 168
 - status information, 121
 - troubleshooting, 121
 - tuning, 115
 - upgrading, 173
- Svm_abort_step_timeout extension property, 215
- Svm_return_step_timeout extension property, 215
- Svm_start_step_timeout extension property, 215
- Svm_step1_timeout extension property, 215
- Svm_step2_timeout extension property, 216
- Svm_step3_timeout extension property, 216
- Svm_step4_timeout extension property
 - definition, 216
- Svm_stop_step_timeout extension property, 216
- syslog() function, 132
- syslog messages, 202, 205, 207
- system configuration files *See* Oracle Database configuration files
- system file, 36
- system messages file, 132
- system panics *See* panics
- system parameter file
 - in cluster file system, 60
 - in Oracle ACFS file system, 60
 - file system options, 62
 - storage management schemes, 25
 - in StorageTek QFS shared file system, 52

T

Targetfilesystem extension property, 213

timeouts

 core-file creation, 120

 log files for, 132

 Oracle ASM, 202

 Oracle Clusterware framework resource method, 117

 panics, and, 138

 reservation step, 116, 202, 214

 Solaris Volume Manager for Sun Cluster definition, 215

topology requirements, 24

trace files

 in cluster file system, 60

 in Oracle ACFS file system, 60

 file system options

 StorageTek QFS shared file system, 54

 UNIX file system, 62

 storage management schemes, 25

 in StorageTek QFS shared file system, 52

troubleshooting

 resources that cannot come online, 145

 Support for Oracle RAC, 121

tuning

 fault monitors, 117

 Support for Oracle RAC, 115

turning off *See* disabling

turning on *See* starting

U

UCMM (userland cluster membership monitor)

 configuration information, 132

 failure to start, 134

 unexpected termination, 133

ucmm_reconf.log file, 132

ucmmd daemon

 failure to start, 134

 unexpected termination, 133

UFS (UNIX file system)

 configuring, 62

uninstalling

 Support for Oracle RAC software packages

 from clusters, 159

 from selected nodes, 168

UNIX file system (UFS)

 configuring, 62

upgrade

 re-creating sun.storage_proxy.type

 resources, 174

upgrading

 Support for Oracle RAC, 173

user identifiers, 30

User_env extension property

 scalable_rac_server_proxy resource type, 209

 SUNW.scalable_rac_server_proxy resource

 type, 209

userland cluster membership monitor (UCMM)

 configuration information, 132

 failure to start, 134

 unexpected termination, 133

users

 creating, 31

 granting access to file systems, 55

 granting access to volumes, 74

V

validation failures

 components, 134, 136

/var/cluster/ucmm directory, 132

/var/opt directory, 34

verifying

 cluster boot, 109

 cluster shutdown, 109

 database resources, 107

 installation, 91

 multiple-owner volume-manager framework

 resource group, 104

 Oracle RAC, 91

 storage resources, 105

 Support for Oracle RAC framework resource

 group, 104

vfstab file

 StorageTek QFS, 54

- UNIX file system, 62
 - volume managers, 28, 28
 - See also* Solaris Volume Manager for Sun Cluster
 - adding to existing configuration, 151
 - storage resources for
 - creating with `clsetup` utility, 79
 - creating with Oracle Solaris Cluster maintenance commands, 231, 232
 - planning, 230
 - volumes, 210
 - adding to multi-owner disk sets, 72
 - monitoring, 144
 - removing from multi-owner disk sets, 158
 - voting disk
 - in cluster file system, 60
 - in Oracle ACFS file system, 61
 - file system options, 62
 - storage management schemes, 26
 - in StorageTek QFS shared file system, 52
 - VUCMM *See* multiple-owner volume-manager framework resource group
 - vucmm_framework resource type
 - dependencies, 70
 - extension properties, 214
 - failure of instance to start, 139
 - instantiating
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
 - registering
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 223
 - START method timeout, 141
 - vucmm_reconf.log file, 132
 - vucmm_svm resource type
 - dependencies, 70, 224
 - extension properties, 214
 - instantiating
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
 - registering
 - with `clsetup` utility, 68
 - with Oracle Solaris Cluster maintenance commands, 224
 - vucmmd daemon
 - failure to start, 136
 - unexpected termination, 136
- ## W
- wait_zc_boot resource type
 - dependencies, 82
 - extension properties, 217
 - width
 - disk stripes, 54
 - wizards
 - multiple-owner volume-manager framework resource group, 68
 - Oracle ASM resource group, 83
 - storage resources, 79
 - Support for Oracle RAC database instances, 97
 - Support for Oracle RAC framework resource group, 64
- ## Z
- Zcname extension property
 - definition, 217
 - zone clusters
 - configuring shared memory, 36
 - creating a file-system mount point resource, 238
 - creating a Support for Oracle RAC framework resource group, 225
 - creating resources, 247
 - creating storage resources, 232
 - creating StorageTek QFS resources, 235
 - logical-hostname resources, 38
 - necessary privileges, 38
 - Oracle ASM, 57
 - preparing for Support for Oracle RAC, 29
 - preventing Oracle Clusterware time synchronization, 38
 - sample configurations, 186
 - storage management schemes, 28