

# **Oracle® Solaris Cluster Data Service for Oracle Grid Engine Guide**

Copyright © 2000, 2010, 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 software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 which 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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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

# Contents

---

<b>Preface</b> .....	5
<b>1 Installing and Configuring HA for Oracle Grid Engine</b> .....	11
HA for Oracle Grid Engine Overview .....	12
Overview of Installing and Configuring HA for Oracle Grid Engine .....	13
Planning the HA for Oracle Grid Engine Installation and Configuration .....	14
Configuration Restrictions .....	14
Configuration Requirements .....	14
Configuration Considerations .....	17
Configuration Planning Questions .....	17
Preparing the Nodes and Disks .....	18
▼ How to Prepare the Nodes and Disks .....	19
Installing and Configuring Oracle Grid Engine .....	20
▼ How to Install and Configure Oracle Grid Engine .....	20
▼ How to Enable Oracle Grid Engine to Run in a Cluster .....	22
Verifying the Installation and Configuration of Oracle Grid Engine .....	23
▼ How to Verify the Installation and Configuration of Oracle Grid Engine .....	23
Installing the HA for Oracle Grid Engine Packages .....	24
▼ How to Install the HA for Oracle Grid Engine Packages .....	24
Configuring the HAStoragePlus Resource Type to Work With HA for Oracle Grid Engine ....	26
▼ How to Register and Configure an HAStoragePlus Resource .....	26
Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine .....	27
▼ How to Configure Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine .....	27
Registering and Configuring HA for Oracle Grid Engine .....	28
Specifying Configuration Parameters for HA for Oracle Grid Engine Resources .....	29
▼ How to Create and Enable HA for Oracle Grid Engine Resources .....	31
Setting HA for Oracle Grid Engine Extension Properties .....	32

Verifying the HA for Oracle Grid Engine Installation and Configuration .....	32
▼ How to Verify the HA for Oracle Grid Engine Installation and Configuration .....	33
Tuning the HA for Oracle Grid Engine Fault Monitors .....	33
Debugging HA for Oracle Grid Engine .....	34
▼ How to Activate Debugging for HA for Oracle Grid Engine .....	34
<b>A Files for Configuring and Removing HA for Oracle Grid Engine Resources .....</b>	<b>37</b>
Listing of sge_config .....	37
Listing of sge_register .....	38
Listing of sge_remove .....	39
<b>Index .....</b>	<b>41</b>

# Preface

---

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

---

**Note** – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

---

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

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

## Using UNIX Commands

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

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

## Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

## Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

## Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at <http://docs.sun.com>.

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

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

## Related Third-Party Web Site References

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

---

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

---

## Documentation, Support, and Training

See the following web sites for additional resources:

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

## Oracle Welcomes Your Comments

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

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

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

## Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Oracle Solaris Operating System (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

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

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

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



# Installing and Configuring HA for Oracle Grid Engine

---

This chapter explains how to install and configure HA for Oracle Grid Engine.

---

**Note** – Oracle Grid Engine was formerly known as Sun ONE Grid Engine. In this book, references to Oracle Grid Engine also apply to Sun ONE Grid Engine unless this book explicitly states otherwise.

---

This chapter contains the following sections.

- “HA for Oracle Grid Engine Overview” on page 12
- “Overview of Installing and Configuring HA for Oracle Grid Engine” on page 13
- “Planning the HA for Oracle Grid Engine Installation and Configuration” on page 14
- “Preparing the Nodes and Disks” on page 18
- “Installing and Configuring Oracle Grid Engine” on page 20
- “Verifying the Installation and Configuration of Oracle Grid Engine” on page 23
- “Installing the HA for Oracle Grid Engine Packages” on page 24
- “Configuring the HAStoragePlus Resource Type to Work With HA for Oracle Grid Engine” on page 26
- “Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine” on page 27
- “Registering and Configuring HA for Oracle Grid Engine” on page 28
- “Verifying the HA for Oracle Grid Engine Installation and Configuration” on page 32
- “Tuning the HA for Oracle Grid Engine Fault Monitors” on page 33
- “Debugging HA for Oracle Grid Engine” on page 34

## HA for Oracle Grid Engine Overview

Oracle Grid Engine is a distributed resource management program, which runs jobs in parallel on multiple machines. To minimize the loss of work that a failure of a machine might cause, nodes in the management tier must be protected against failure. However, protection of individual execution nodes in the grid against failure is not required. Failure of an individual execution node in a grid causes only a minor loss of work.

To eliminate single points of failure in the management tier of an Oracle Grid Engine system, HA for Oracle Grid Engine provides fault monitoring and automatic fault recovery for the following Oracle Grid Engine daemons:

- Queue master daemon
- Scheduling daemon

You must configure HA for Oracle Grid Engine as a failover service.

For conceptual information about failover data services and scalable data services, see [Oracle Solaris Cluster Concepts Guide](#).

Because the management tier relies on the Oracle Grid Engine file system, the NFS server that exports this file system must also be protected against failure. To eliminate single points of failure in the NFS server, use the Oracle Solaris Cluster HA for NFS data service. For more information about this data service, see [Oracle Solaris Cluster Data Service for Network File System \(NFS\) Guide](#).

Each component of Oracle Grid Engine has a data service that protects the component when the component is configured in Oracle Solaris Cluster. See the following table.

TABLE 1-1 Protection of Oracle Grid Engine Components by Oracle Solaris Cluster Data Services

Oracle Grid Engine Component	Data Service
Oracle Grid Engine daemons:	HA for Oracle Grid Engine
▪ Queue master daemon (sge_qmaster)	The resource type is SUNW.gds.
▪ Scheduling daemon (sge_schedd)	
NFS server	Oracle Solaris Cluster HA for NFS
	The resource type is SUNW.nfs.

# Overview of Installing and Configuring HA for Oracle Grid Engine

The following table summarizes the tasks for installing and configuring HA for Oracle Grid Engine and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order that they are listed in the table.

TABLE 1-2 Tasks for Installing and Configuring HA for Oracle Grid Engine

Task	Instructions
Plan the installation	<a href="#">“HA for Oracle Grid Engine Overview” on page 12</a> <a href="#">“Planning the HA for Oracle Grid Engine Installation and Configuration” on page 14</a>
Prepare the nodes and disks	<a href="#">“Preparing the Nodes and Disks” on page 18</a>
Install and configure Oracle Grid Engine	<a href="#">“Installing and Configuring Oracle Grid Engine” on page 20</a>
Verify HA for Oracle Grid Engine installation and configuration	<a href="#">“Verifying the Installation and Configuration of Oracle Grid Engine” on page 23</a>
Install HA for Oracle Grid Engine Packages	<a href="#">“Installing the HA for Oracle Grid Engine Packages” on page 24</a>
Configure the HAStoragePlus resource type to work with HA for Oracle Grid Engine	<a href="#">“Configuring the HAStoragePlus Resource Type to Work With HA for Oracle Grid Engine” on page 26</a>
Configure Oracle Solaris Cluster HA for NFS for use with HA for Oracle Grid Engine	<a href="#">“Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine” on page 27</a>
Register and Configure HA for Oracle Grid Engine	<a href="#">“Registering and Configuring HA for Oracle Grid Engine” on page 28</a>
Verify HA for Oracle Grid Engine installation and configuration	<a href="#">“Verifying the HA for Oracle Grid Engine Installation and Configuration” on page 32</a>
Tune HA for Oracle Grid Engine fault monitors	<a href="#">“Tuning the HA for Oracle Grid Engine Fault Monitors” on page 33</a>
Debug HA for Oracle Grid Engine	<a href="#">“Debugging HA for Oracle Grid Engine” on page 34</a>

# Planning the HA for Oracle Grid Engine Installation and Configuration

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

---

**Note** – Before you begin, consult your Oracle Grid Engine documentation for configuration restrictions and requirements that are not imposed by Oracle Solaris Cluster software.

---

## Configuration Restrictions

The configuration restrictions in the subsections that follow apply only to HA for Oracle Grid Engine.



---

**Caution** – Your data service configuration might not be supported if you do not observe these restrictions.

---

### Oracle Grid Engine Shadow Daemon

Do *not* use the Oracle Grid Engine shadow daemon. The Oracle Grid Engine shadow daemon provides an optional mechanism for recovery from failures. This mechanism interferes with the automatic fault recovery that Oracle Solaris Cluster provides.

### Oracle Grid Engine Berkley DB spooling server

Do *not* choose the option to use a Berkley DB spooling server. Either choose the Classic spooling method or the local Berkley DB spooling method. Currently it is not possible to configure the Berkley DB spooling server in a highly available way within the Oracle Solaris Cluster framework.

### Start at Boot Option

Do *not* choose the start at boot option when installing Oracle Grid Engine. To ensure that HA for Oracle Grid Engine can provide fault monitoring and automatic fault recovery, Oracle Grid Engine must be started *only* by Oracle Solaris Cluster.

## Configuration Requirements

The configuration requirements in this section apply only to HA for Oracle Grid Engine.



**Caution** – If your data service configuration does not conform to these requirements, the data service configuration might not be supported.

## Oracle Grid Engine Software Version Requirements

Use Oracle Grid Engine version 6.0 or 6.1. Make sure to apply the most recent available Patches to the Oracle Grid Engine software.

## Operating System for the Oracle Grid Engine Management Tier

The Oracle Grid Engine management tier must run on Oracle Solaris Cluster nodes. Because Oracle Solaris Cluster runs only on the Solaris Operating System, the Oracle Grid Engine management tier must also run on the Solaris Operating System. However, Oracle Grid Engine supports other operating systems. Therefore, this requirement applies only to the management tier, *not* to individual execution nodes in the grid.

## Memory Requirements

Ensure that enough free memory is available on the cluster nodes where you plan to run the Oracle Grid Engine master.

The amount of free memory that is required on each cluster node depends on the number of jobs that are running on the grid. For example:

- If 100 jobs are running, 10 Mbytes of free memory are required.
- If 10,000 jobs are running, 1 Gbyte of free memory is required.

## Disk Space Requirements

Ensure that you have enough disk space in the Oracle Grid Engine file system and on the local disk of each node.

The disk space requirements for each type of file or directory in the Oracle Grid Engine file system are listed in the following table.

File Type or Directory Type	Required Disk Space
Binary files	15 Mbytes for each architecture
Spool directories	30–200 Mbytes
Installation tar file	40 Mbytes

On the local disk of each node, 10–20 Mbytes of disk space are required. If you are installing the Oracle Grid Engine software on the local disk of a node, 15 Mbytes of disk space are additionally required for the binary files.

## HA for Oracle Grid Engine Configuration Requirements

Configure HA for Oracle Grid Engine as a failover data service. You cannot configure HA for Oracle Grid Engine as a scalable data service. For more information, see:

- [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#)
- [“Registering and Configuring HA for Oracle Grid Engine” on page 28](#)

---

**Note** – If you are using the Solaris 10 OS, install and configure this data service to run only in the global zone. At publication of this document, this data service is not supported in non-global zones. For updated information about supported configurations of this data service, contact your Oracle service representative.

---

## NFS Configuration for the Oracle Grid Engine File System

The Oracle Grid Engine file system must reside on a multihost disk. This disk must be available to the other nodes in the cluster that will be used for the Oracle Grid Engine administrative services,

You must use NFS to export the Oracle Grid Engine file system to the noncluster nodes. The NFS server that exports this file system must also be protected against failure. To protect the NFS server against failure, use the Oracle Solaris Cluster HA for NFS data service. For more information about this data service, see [Oracle Solaris Cluster Data Service for Network File System \(NFS\) Guide](#).

## Oracle Solaris Cluster HA for NFS Configuration Requirements

Configure the resources for the Oracle Grid Engine management tier in the same resource group as the resource for NFS. For more information, see [“Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine” on page 27](#).

## Dependencies Between Oracle Grid Engine Components

The dependencies between Oracle Grid Engine components are shown in the following table.

TABLE 1-3 Dependencies Between Oracle Grid Engine Components

Oracle Grid Engine Component	Dependency
Oracle Grid Engine queue master daemon (sge_qmaster)	SUNW.HAStoragePlus resource
Oracle Grid Enginescheduling daemon (sge_schedd)	Oracle Grid Engine queue master daemon (sge_qmaster) resource

These dependencies are set when you register and configure HA for Oracle Grid Engine. For more information, see [“Registering and Configuring HA for Oracle Grid Engine” on page 28](#).

## Configuration Considerations

The configuration considerations in the subsections that follow affect the installation and configuration of HA for Oracle Grid Engine.

### Location of the Oracle Grid Engine Binary Files

You can install Oracle Grid Engine on one of the following locations:

- A highly available local file system
- The cluster file system

For the advantages and disadvantages of placing the Oracle Grid Engine binary files on a highly available local file system and the cluster file system, see “[Configuration Guidelines for Oracle Solaris Cluster Data Services](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

---

**Tip** – To enable the type of file system to be identified from the mount point, use a prefix that indicates the type of file system as follows:

- For mount points on a highly available local file system, use the `/local` prefix.
  - For mount points on the cluster file system, use the `/global` prefix.
- 

### File Systems for Spool Directories and Binary Files

The optimum distribution of spool directories and binary files among file systems depends on the grid configuration. See the following table.

Grid Configuration	File System Configuration
The execution tier contains fewer than 200 hosts.	Use a single shared NFS file system under the root of the Oracle Grid Engine file system for the spool directories and binary files.
The execution tier contains about 200 hosts, or the applications are disk intensive.	Use a separate area on an NFS file system for the spool directories.
The execution tier contains more than 200 hosts, or NFS performance is likely to be a problem.	See the Oracle Grid Engine documentation for alternate grid configurations.

## Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of HA for Oracle Grid Engine. Write the answers to these questions in the space that is provided on the data service worksheets in “[Configuration Worksheets](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

- Which resource group will you use for the following resources:
  - Logical host name resource
  - HAStoragePlus resource
  - NFS resource
  - Oracle Grid Engine application resources

Use the answer to this question when you perform the following procedures:

- [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#)
  - [“Configuring the HAStoragePlus Resource Type to Work With HA for Oracle Grid Engine” on page 26](#)
  - [“Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine” on page 27](#)
  - [“Specifying Configuration Parameters for HA for Oracle Grid Engine Resources” on page 29](#)
- What is the logical host name for the Oracle Grid Engine resource? Clients access the data service through this logical host name.

Use the answer to this question when you perform the procedure [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#).

- Which resources will you use for the components of Oracle Grid Engine?

You require one resource for each component in the following list:

- Queue master daemon
- Scheduling daemon

Use the answer to this question when you perform the procedure [“Specifying Configuration Parameters for HA for Oracle Grid Engine Resources” on page 29](#).

- Where will the system configuration files reside?

See [“Configuration Guidelines for Oracle Solaris Cluster Data Services”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for the advantages and disadvantages of using the local file system instead of the cluster file system.

## Preparing the Nodes and Disks

Preparing the nodes and disks modifies the configuration of the operating system to enable HA for Oracle Grid Engine to eliminate single points of failure in an Oracle Grid Engine system.

Before you begin, ensure that the requirements in the following sections are met:

- [“Memory Requirements” on page 15](#)
- [“Disk Space Requirements” on page 15](#)

## ▼ How to Prepare the Nodes and Disks

**1 Become superuser on all the cluster nodes where you are installing Oracle Grid Engine.**

**2 Create an administrative user account for Oracle Grid Engine on all those cluster nodes.**

Either select an existing user account other than root for the grid administration, or create an account specifically for grid administration.

---

**Tip** – For consistency with the Oracle Grid Engine documentation, name the account `sgeadmin`.

---

**3 Create a directory for the root of Oracle Grid Engine file system.**

```
# mkdir sge-root-dir
```

---

**Note** – The `sge-root-dir` must reside in the cluster filesystem. Refer to “[Configuring the HASToragePlus Resource Type to Work With HA for Oracle Grid Engine](#)” on page 26 for more details.

---

**4 Change the owner of the root of the Oracle Grid Engine file system to the administrative user whose account you created in Step 2.**

```
# chown sge-admin sge-root-dir
```

**5 Set the mode of the root of Oracle Grid Engine file system to `drwxr-xr-x`.**

```
# chmod 755 sge-root-dir
```

**6 Specify the port number and protocol for the `sge_qmaster` and `sge_execd` services.**

Choose an unused port number below 1024. The `sge_qmaster` and `sge_execd` services are to be provided through Transmission Control Protocol (TCP).

To specify the port number and protocol, add the following line to the `/etc/services` file.

```
sge_qmaster    port-no/tcp
sge_execd     port-no/tcp
```

**7 For each type of host in the grid, create a plain text file that contains the names of all hosts of that type in the grid.**

The `install_qmaster` script uses these files when you install Oracle Grid Engine. Create a separate file for each type of host in the grid:

- Execution hosts
- Administrative hosts
- Submit hosts

**Example 1-1** Preparing the Nodes and Disks for the Installation of Oracle Grid Engine

This example shows how to prepare the nodes and disks for an Oracle Grid Engine installation that is to be configured as follows:

- The root of Oracle Grid Engine file system is the `/global/gridmaster` directory. This directory resides in the cluster file system.
- The account for grid administration is named `sgeadmin`.
- The `sge_qmaster` service is to be provided through port 536 and TCP.
- The `sge_execd` service is to be provided through port 537 and TCP.

The sequence of operations for preparing the nodes and disks for the installation of Oracle Grid Engine is as follows:

1. To create the `/global/gridmaster` directory for the root of Oracle Grid Engine file system, the following command is run:

```
# mkdir /global/gridmaster
```

2. To change the owner of the `/global/gridmaster` directory to the `sgeadmin` user, the following command is run:

```
# chown sgeadmin /global/gridmaster
```

3. To set the mode of the `/global/gridmaster` directory to `drwxr-xr-x`, the following command is run:

```
# chmod 755 /global/gridmaster
```

4. To specify that the `sge_qmaster` service is to be provided through port 536 and TCP, and that the `sge_execd` service is to be provided through port 537 and TCP, the following line is added to the `/etc/services` file:

```
sge_qmaster    536/tcp
sge_execd     537/tcp
```

## Installing and Configuring Oracle Grid Engine

The procedure that follows explains only the special requirements for installing Oracle Grid Engine for use with HA for Oracle Grid Engine. For complete information about installing and configuring Oracle Grid Engine, see your Oracle Grid Engine documentation.

To enable Oracle Grid Engine to run in a cluster, you must modify Oracle Grid Engine to use a logical host name.

### ▼ How to Install and Configure Oracle Grid Engine

Before you begin, ensure that you have the host names of all hosts in the grid. Create a separate list of host names for each type of host in the grid:

- Execution hosts
- Administrative hosts
- Submit hosts

- 1 Become superuser of the cluster node where you are installing Oracle Grid Engine.**
- 2 Install the Oracle Grid Engine distribution files. You have to choose between the tar .gz format and the pkgadd format.**

Follow the instructions outlined in “How to Load the Distribution Files On a Workstation” in the *N1 Grid Engine 6 Installation Guide*.

---

**Note** – If you choose the pkgadd format, you need to make sure to install Patches for the Oracle Grid Engine software on exactly the same node the Oracle Grid Engine packages are registered on.

---

- 3 Set the SGE\_ROOT environment variable to the directory for the root of Oracle Grid Engine file system that you created in “Preparing the Nodes and Disks” on page 18.**

```
# SGE_ROOT=sge-root-dir
# export SGE_ROOT
```

- 4 Go to the directory for the root of Oracle Grid Engine file system.**

```
# cd sge-root-dir
```

- 5 Start the script that installs the Oracle Grid Engine master host.**

```
# ./install_qmaster
```

- 6 Follow the prompts on screen to provide or confirm the following information:**

- The name of the Oracle Grid Engine administrative user
- The value of the SGE\_ROOT environment variable
- The TCP port number
- The name of the Oracle Grid Engine cell to be configured
- The path to the spool directory
- The setup for the correct file permissions
- Details of your domain name service (DNS) domains

- 7 When you are asked whether you want to use classic spooling or Berkley DB, do not choose to use a Berkely DB spooling Server.**

Either choose the classic spooling method, or choose Berkley DB with local spooling.

- 8 When you are prompted, specify the range of group IDs for Oracle Grid Engine to use.**

To ensure that you allocate enough group IDs, specify a range of approximately 100 group IDs, for example, 20000-20100.

**9 Follow the prompts on screen to provide or confirm the following information:**

- The path to the spooling directory for the execution daemon
- The email address of the user who should receive problem reports
- Confirm the configuration parameters

**10 When you are asked if you want to install the script that starts Oracle Grid Engine at boot time, reply no.**

You are asked if you want to install the script that starts Oracle Grid Engine at boot time.

```
We can install the startup script that will
start qmaster/scheduler at machine boot (y/n) [y] >> n
```

To ensure that HA for Oracle Grid Engine can provide fault monitoring and automatic fault recovery, Oracle Grid Engine must be started *only* by Oracle Solaris Cluster.

**11 Follow the prompts on screen to provide or confirm the following information:**

- Specify the list of execution, admin and submit hosts
- Do not use a shadow host
- Select a scheduler profile

## ▼ How to Enable Oracle Grid Engine to Run in a Cluster

**1 Become superuser of a node in the cluster that will host Oracle Grid Engine.****2 Create a failover resource group to contain the HA for Oracle Grid Engine resources.**

Use the resource group that you identified when you answered the questions in “[Configuration Planning Questions](#)” on page 17.

```
# clresourcegroup create -p Pathprefix=sge-root-dir sge-rg
```

```
-p Pathprefix= sge-root-dir
```

Specifies a directory on a cluster file system that Oracle Solaris Cluster HA for NFS uses to maintain administrative and status information. This directory must be the directory that you created for the root of the Oracle Grid Engine file system in “[Preparing the Nodes and Disks](#)” on page 18.

```
sge-rg
```

Specifies that the resource group that you are creating is named *sge-rg*.

**3 Add a resource for the Oracle Grid Engine logical host name to the failover resource group that you created in [Step 2](#).**

```
# clreslogicalhostname create \
-g sge-rg \
-h hostlist \
sge-lh-rs
```

- g *sge-rg*  
Specifies that the logical host name resource is to be added to the failover resource group that you created in [Step 2](#)
- h *hostlist*  
Specifies a comma-separated list of host names that are to be made available by this logical host name resource
- sge-lh-rs*  
Specifies that the resource that you are creating is named *sge-lh-rs*

## Verifying the Installation and Configuration of Oracle Grid Engine

Before you install the HA for Oracle Grid Engine packages, verify that the Oracle Grid Engine software is correctly installed and configured to run in a cluster. This verification does *not* verify that the Oracle Grid Engine application is highly available because the HA for Oracle Grid Engine data service is not yet installed.

---

**Note** – If any step in this procedure fails, see your Oracle Grid Engine documentation for more information about how to verify the Oracle Grid Engine installation.

---

### ▼ How to Verify the Installation and Configuration of Oracle Grid Engine

You verify the installation and configuration of Oracle Grid Engine by submitting a dummy job and checking that the required processes are running.

- 1 **Log in to the master host as the administrative user whose account you created in “[Preparing the Nodes and Disks](#)” on page 18.**
- 2 **Set the `SGE_ROOT` environment variable to the directory for the root of Oracle Grid Engine file system that you created in “[Preparing the Nodes and Disks](#)” on page 18.**

```
$ SGE_ROOT=sge-root-dir
$ export SGE_ROOT
```

- 3 **Start the script that modifies your environment to enable Oracle Grid Engine to run.**

```
$ . $SGE_ROOT/default/common/settings.sh
```

- 4 **Submit a dummy job to Oracle Grid Engine.**

```
$ qsub $SGE_ROOT/examples/jobs/sleeper.sh
your job 1 (*Sleeper*) has been submitted
```

**5 On the master host, confirm that these processes are running:**

- `sge_qmaster`
- `sge_schedd`

```
# ps -ef | grep sge_  
root 429 1 0 Jul 27 3:37 /global/gridmaster/bin/solaris64/sge_qmaster  
root 429 1 0 Jul 27 3:37 /global/gridmaster/bin/solaris64/sge_schedd
```

**6 View the global configuration of the grid.**

- If you are using the command line, type the following command:  

```
$ qconf -sconf
```
- If you are using the QMON graphical user interface (GUI), select Cluster Configuration.

**7 On at minimum one execution host, confirm that these processes are running:**

- `sge_execd`

```
# ps -ef | grep sge_  
root 451 1 0 Jul 27 3:37 /global/gridmaster/bin/solaris64/sge_execd
```

## Installing the HA for Oracle Grid Engine Packages

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

---

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

---

### ▼ How to Install the HA for Oracle Grid Engine Packages

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

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

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

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

- 1 On the cluster node where you are installing the data service packages, become superuser.**
- 2 Load the Oracle Solaris Cluster installation media into the DVD-ROM drive.**

If the Volume Management daemon `vold(1M)` is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `/cdrom` directory.
- 3 Change to the installation wizard directory of the DVD-ROM.**
  - **If you are installing the data service packages on the SPARC platform, type the following command:**

```
# cd /cdrom/cdrom0/Solaris_sparc
```
  - **If you are installing the data service packages on the x86 platform, type the following command:**

```
# cd /cdrom/cdrom0/Solaris_x86
```
- 4 Start the installation wizard.**

```
# ./installer
```
- 5 When you are prompted, accept the license agreement.**
- 6 From the list of Oracle Solaris Cluster agents under Availability Services, select the data service for Oracle Grid Engine.**
- 7 If you require support for languages other than English, select the option to install multilingual packages.**

English language support is always installed.
- 8 When prompted whether to configure the data service now or later, choose Configure Later.**

Choose Configure Later to perform the configuration after the installation.
- 9 Follow the instructions on the screen to install the data service packages on the node.**

The installation wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.
- 10 (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.**

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

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

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

b. Eject the DVD-ROM.

```
# eject cdrom
```

**Next Steps** Refer to the [Oracle Solaris Cluster Data Service for Network File System \(NFS\) Guide](#) on how to also install the Oracle Solaris Cluster HA for NFS packages.

## Configuring the HASStoragePlus Resource Type to Work With HA for Oracle Grid Engine

For maximum availability of the Oracle Grid Engine application, resources that HA for Oracle Grid Engine requires must be available before the Oracle Grid Engine management tier is started. An example of such a resource is the Oracle Grid Engine file system. To ensure that these resources are available, configure the HASStoragePlus resource type to work with HA for Oracle Grid Engine.

For information about the relationship between resource groups and disk device groups, see “[Relationship Between Resource Groups and Device Groups](#)” in [Oracle Solaris Cluster Data Services Planning and Administration Guide](#).

Configuring the HASStoragePlus resource type to work with HA for Oracle Grid Engine involves the following operations:

- Synchronizing the startups between resource groups and disk device groups as explained in “[Synchronizing the Startups Between Resource Groups and Device Groups](#)” in [Oracle Solaris Cluster Data Services Planning and Administration Guide](#)
- Registering and configuring an HASStoragePlus resource

### ▼ How to Register and Configure an HASStoragePlus Resource

**1 Become superuser on a node in the cluster that will host Oracle Grid Engine.**

**2 Register the SUNW.HASStoragePlus resource type.**

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

- 3 Add an `HASStoragePlus` resource for the Oracle Grid Engine file system to the resource group that you created in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#).

```
# clresource create \  
-g sge-rg \  
-t SUNW.HASStoragePlus \  
-p FilesystemMountPoints=sge-root \  
sge-hasp-rs
```

```
-g sge-rg
```

Specifies that the resource is to be added to the resource group that you created in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#)

```
-p FilesystemMountPoints=sge-root
```

Specifies that the mount point for this file system is the root of the Oracle Grid Engine file system

```
sge-hasp-rs
```

Specifies that the resource that you are creating is named `sge-hasp-rs`

## Configuring Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine

You must use NFS to export the Oracle Grid Engine file system to the noncluster nodes. The NFS server that exports this file system must also be protected against failure. To protect the NFS server against failure, use the Oracle Solaris Cluster HA for NFS data service.

The procedure that follows explains only the special requirements for using Oracle Solaris Cluster HA for NFS with HA for Oracle Grid Engine. For complete information about installing and configuring Oracle Solaris Cluster HA for NFS, see [Oracle Solaris Cluster Data Service for Network File System \(NFS\) Guide](#).

### ▼ How to Configure Oracle Solaris Cluster HA for NFS for Use With HA for Oracle Grid Engine

---

**Note** – Commands in this procedure assume that you have set the `$SGE_ROOT` environment variable to specify the root of the Oracle Grid Engine file system.

---

- 1 Register the `SUNW.nfs` resource type.

```
# clresourcetype register SUNW.nfs
```

**2 From any cluster node, create a directory for NFS configuration files.**

Create the directory under root of the Oracle Grid Engine file system. Name the directory `SUNW.nfs`.

```
# mkdir -p $SGE_ROOT/SUNW.nfs
```

**3 In the directory that you created in Step 2, create a file that contains the share command for the root of the Oracle Grid Engine file system.**

Name the file the `dfstab.sge-nfs-rs`, where `sge-nfs-rs` is the name of the NFS resource that you will create in Step 4.

```
# echo "share -F nfs -o rw sge-root" \  
> $SGE_ROOT/SUNW.nfs/dfstab.sge-nfs-rs
```

**4 Add a `SUNW.nfs` resource to the failover resource group that you created in “How to Enable Oracle Grid Engine to Run in a Cluster” on page 22.**

```
# clresource create \  
-g sge-rg \  
-t SUNW.nfs \  
-p Resource_dependencies=sge-hasp-rs \  
sge-nfs-rs
```

**Example 1–2 Creating a `dfstab` File for the Root of the Oracle Grid Engine File System**

This example shows the command for creating a `dfstab` file for the root of the Oracle Grid Engine file system.

- The root of the Oracle Grid Engine file system is `/global/gridmaster`.
- The name of the NFS resource for which this file is created is `sge-nfs-rs`.

```
# echo "share -F nfs -o rw /global/gridmaster" \  
> /global/gridmaster/SUNW.nfs/dfstab.sge-nfs-rs
```

## Registering and Configuring HA for Oracle Grid Engine

Before you perform this procedure, ensure that the HA for Oracle Grid Engine data service packages are installed.

Use the configuration and registration files in the `/opt/SUNWscsge/util` directory to register the HA for Oracle Grid Engine resources. The files define the dependencies that are required between Oracle Grid Engine components. For information about these dependencies, see “Dependencies Between Oracle Grid Engine Components” on page 16. For a listing of these files, see Appendix A, “Files for Configuring and Removing HA for Oracle Grid Engine Resources.”

Registering and configuring HA for Oracle Grid Engine involves the tasks that are explained in the following sections:

1. [“Specifying Configuration Parameters for HA for Oracle Grid Engine Resources” on page 29](#)
2. [“How to Create and Enable HA for Oracle Grid Engine Resources” on page 31](#)

## Specifying Configuration Parameters for HA for Oracle Grid Engine Resources

HA for Oracle Grid Engine provides scripts that automate the process of configuring and removing HA for Oracle Grid Engine resources. These scripts obtain configuration parameters from the `sge_config` file in the `/opt/SUNWscsge/util/` directory. To specify configuration parameters for HA for Oracle Grid Engine resources, edit the `sge_config` file.

Each configuration parameter in the `sge_config` file is defined as a keyword-value pair. The `sge_config` file already contains the required keywords and equals signs. For more information, see [“Listing of `sge\_config`” on page 37](#). When you edit the `sge_config` file, add the required value to each keyword. Use the values that you identified in [“Configuration Planning Questions” on page 17](#).

The keyword-value pairs in the `sge_config` file are as follows:

```
QMASTERS=sge-qmaster-rs
SCHEDDRS=sge-schedd-rs
MASTERRG=sge-rg
MASTERLH=sge-lh-rs
MASTERPORT=portno
MASTERHASP=sge-hasp-rs
SGE_ROOT=sge-root-dir
SGE_CELL=cell-name
SGE_VER=6.0
```

The meaning and permitted values of the keywords in the `sge_config` file are as follows:

`QMASTERS=sge-qmaster-rs`

Specifies the name that you are assigning to the resource for the Oracle Grid Engine queue master daemon `sge_qmaster`. This must be defined.

`SCHEDDRS=sge-schedd-rs`

Specifies the name that you are assigning to the resource for the Oracle Grid Engine scheduling daemon `sge_schedd`. This must be defined.

`MASTERRG=sge-rg`

Specifies the name of the resource group that contains the HA for Oracle Grid Engine resources. This name must be the name that you assigned when you created the resource group as explained in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#). This must be defined.

`MASTERLH=sge-lh-rs`

Specifies the name of the logical host name resource for Oracle Grid Engine. This name must be the name that you assigned when you created the resource in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#). This must be defined.

`MASTERPORT=portno`

Specifies the port number that is configured for `sge_qmaster`, the default is set to 536. It must be an integer and must be defined.

`MASTERHASP=sge-hasp-rs`

Specifies the name of the `SUNW.HASStoragePlus` resource for Oracle Grid Engine. This name must be the name that you assigned when you created the resource in “[Configuring the HASStoragePlus Resource Type to Work With HA for Oracle Grid Engine](#)” on page 26. If this resource is used it must be defined.

`SGE_ROOT=sge-root-dir`

Specifies the root directory of the Oracle Grid Engine file system. This directory must be the directory that you created for root of the Oracle Grid Engine file system in “[Preparing the Nodes and Disks](#)” on page 18. This must be defined.

`SGE_CELL=cell-name`

Specifies the cell that Oracle Grid Engine references. This must be defined.

`SGE_VER=6.0`

Specifies the version of the installed Sun Grid Engine configuration. This keyword must be defined and set the value to "6.0".

---

**Note** – You must set the `SGE_VER` keyword to “6.0”, even if you are using Oracle Grid Engine version 6.1.

---

#### EXAMPLE 1-3 Sample `sge_config` File

This example shows an `sge_config` file in which configuration parameters are set as follows:

- The name of the resource for the Oracle Grid Engine queue master daemon `sge_qmaster` is `sge_qmaster-rs`.
- The name of the resource for the Oracle Grid Engine scheduling daemon `sge_schedd` is `sge_schedd-rs`.
- The name of the resource group that contains the HA for Oracle Grid Engine resources is `sge-rg`.
- The name of the logical host name resource for Oracle Grid Engine is `sge-lh-rs`.
- The port number for `sge_qmaster` is set to 536.
- The name of the `SUNW.HASStoragePlus` resource for Oracle Grid Engine is `sge-hasp-rs`.
- The root directory of the Oracle Grid Engine file system is `/global/gridmaster`.
- Oracle Grid Engine references the default cell.
- The version for Oracle Grid Engine is set to 6.0.

```
QMASTERRS=sge_qmaster-rs
SCHEDDRS=sge_schedd-rs
MASTERRG=sge-rg
```

EXAMPLE 1-3 Sample sge\_config File (Continued)

```
MASTERLH=sge-lh-rs
MASTERPORT=536
MASTERHASP=sge-hasp-rs
SGE_ROOT=/global/gridmaster
SGE_CELL=default
SGE_VER=6.0
```

## ▼ How to Create and Enable HA for Oracle Grid Engine Resources

Before you begin, ensure that you have edited the sge\_config file or a copy of it to specify configuration parameters for HA for Oracle Grid Engine resources. For more information, see [“Specifying Configuration Parameters for HA for Oracle Grid Engine Resources” on page 29](#).

### 1 Register the SUNW.gds resource type.

```
# clresourcetype register SUNW.gds
```

### 2 Go to the directory that contains the script for creating the Oracle Grid Engine resources.

```
# cd /opt/SUNWscsge/util/
```

### 3 Run the script that creates the Oracle Grid Engine resources.

```
# ./sge_register -f /mypath/sge_config
```

### 4 Bring online the failover resource group that you created in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#).

This resource group contains the following resources:

- Logical host name resource
- HAStoragePlus resource
- NFS resource
- Oracle Grid Engine application resources

```
# clresourcegroup online -M sge-rg
```

*sge-rg* Specifies the resource group that you created in [“How to Enable Oracle Grid Engine to Run in a Cluster” on page 22](#) is to be brought online



---

**Caution** – Make sure that the Oracle Grid Engine daemons (`sgc_qmaster` and `sgc_schedd`) are not running before bringing the failover resource group online. They may be running because the `install_qmaster` installation script started them or they are still running after performing the verification described in “[How to Verify the HA for Oracle Grid Engine Installation and Configuration](#)” on page 33.

---

## Setting HA for Oracle Grid Engine Extension Properties

Extension properties for HA for Oracle Grid Engine resources are set when you run the script that creates these resources. You need to set these properties only if you require values other than the values that are set by the script. For information about HA for Oracle Grid Engine extension properties, see the [SUNW.gds\(5\)](#) man page. You can update some extension properties dynamically. You can update other properties, however, only when you create or disable a resource. The Tunable entry indicates when you can update a property.

To update an extension property of a resource, run the `clresource(1CL)` command with the following option to modify the resource:

`-p property=value`  
`-p property`      Identifies the extension property that you are setting  
`value`              Specifies the value to which you are setting the extension property

You can also use the procedures in [Chapter 2, “Administering Data Service Resources,”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure resources after the resources are created.

## Verifying the HA for Oracle Grid Engine Installation and Configuration

After you install, register, and configure HA for Oracle Grid Engine, verify the HA for Oracle Grid Engine installation and configuration. Verifying the HA for Oracle Grid Engine installation and configuration determines if the HA for Oracle Grid Engine data service makes the Oracle Grid Engine application highly available.

## ▼ How to Verify the HA for Oracle Grid Engine Installation and Configuration

- 1 Become superuser a node that will host Oracle Grid Engine.
- 2 Verify that all Oracle Grid Engine resources are online.  
# `cluster status -t rg,rs`
- 3 If an Oracle Grid Engine resource is not online, enable the resource.  
# `clresource enable sge-rs`
- 4 Switch the Oracle Grid Engine resource group to another cluster node.  
# `clresourcegroup switch -n node sge-rg`

## Tuning the HA for Oracle Grid Engine Fault Monitors

The HA for Oracle Grid Engine fault monitors verify that the following daemons are running correctly:

- Queue master daemon `sge_qmaster`
- Scheduling daemon `sge_schedd`

Each HA for Oracle Grid Engine fault monitor is contained in the resource that represents Oracle Grid Engine component. You create these resources when you register and configure HA for Oracle Grid Engine. For more information, see [“Registering and Configuring HA for Oracle Grid Engine” on page 28](#).

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

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

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

For more information, see [“Tuning Fault Monitors for Oracle Solaris Cluster Data Services” in \*Oracle Solaris Cluster Data Services Planning and Administration Guide\*](#).

# Debugging HA for Oracle Grid Engine

The config file in the `/opt/SUNWscsge/etc` directory enables you to activate debugging for Oracle Grid Engine resources. This file enables you to activate debugging for all Oracle Grid Engine resources or for a specific Oracle Grid Engine resource on a particular node. If you require debugging for HA for Oracle Grid Engine to be enabled throughout the cluster, repeat this procedure on all nodes.

## ▼ How to Activate Debugging for HA for Oracle Grid Engine

### 1 Determine whether debugging for HA for Oracle Grid Engine is active.

If debugging is inactive, `daemon.notice` is set in the file `/etc/syslog.conf`.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                  operator
#
```

### 2 If debugging is inactive, edit the `/etc/syslog.conf` file to change `daemon.notice` to `daemon.debug`.

### 3 Confirm that debugging for HA for Oracle Grid Engine is active.

If debugging is active, `daemon.debug` is set in the file `/etc/syslog.conf`.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                  operator
#
```

### 4 Restart the `syslogd` daemon.

```
# svcadm restart system-log
```

### 5 Edit the `/opt/SUNWscsge/etc/config` file to change `DEBUG=` to `DEBUG=ALL` or `DEBUG=sge-rs`.

```
# cat /opt/SUNWscsge/etc/config
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident "@(#)config 1.1 06/02/18 SMI"
#
# Usage:
#     DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
#
```

---

**Note** – To deactivate debugging, reverse the preceding steps.

---



## Files for Configuring and Removing HA for Oracle Grid Engine Resources

---

The /opt/SUNWscsge/util directory contains files that automate the process of configuring and removing HA for Oracle Grid Engine resources. Listings of these files are provided in the following sections:

- “Listing of sge\_config” on page 37
- “Listing of sge\_register” on page 38
- “Listing of sge\_remove” on page 39

### Listing of sge\_config

```
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident "@(#)sge_config 1.3 06/07/14 SMI"
#
# This file will be sourced in by sge_register and the parameters
# listed below will be used.
#
# These parameters can be customized in (key=value) form and are always needed:
#
# QMASTERRS - name of the resource for the sge_qmaster daemon.
# SCHEDDRS - name of the resource for the sge_schedd daemon.
# MASTERRG - name of the resource group containing the resources
#           for sge_qmaster and sge_schedd.
# MASTERPORT - tcp port number that sge_qmaster will listen to
#              (normally 536).
# MASTERLH - name of the LogicalHostname SC resource within MASTERRG.
# MASTERHASP - name of the HAStoragePlus SC resource within MASTERRG.
# SGE_ROOT - SGE_ROOT of this Sun GridEngine Installation.
# SGE_CELL - SGE_CELL of this Sun GridEngine Installation.
# SGE_VER - Version of this Sun GridEngine Installation.
#           Currently this can only be set to 6.0 .
#
#
```

```
QMASTERRS=
```

```

SCHEDDRS=
MASTERRG=
MASTERPORT=536
MASTERLH=
MASTERHASP=
SGE_ROOT=
SGE_CELL=
SGE_VER=6.0

```

## Listing of sge\_register

```

#!/bin/ksh
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident      "@(#)sge_register.ksh      1.6      06/07/14 SMI"
#

MYCONFIG=""
MYNAME='basename $0'
GDSTYPE=SUNW.gds

typeset opt

while getopts 'f:' opt
do
    case "${opt}" in
        f)  MYCONFIG=${OPTARG};;
        *)  echo "ERROR: ${MYNAME} Option ${OPTARG} unknown - early End."
            echo "Only the option -f <filename> is valid."
            exit 1;;
    esac
done

if [ -n "${MYCONFIG}" ] && [ -f "${MYCONFIG}" ]
then
    echo "sourcing ${MYCONFIG}"
    . ${MYCONFIG}
else
    PKGCONF='dirname $0'/sge_config
    echo "sourcing ${PKGCONF}"
    . ${PKGCONF}
fi

if [ "${SGE_VER}" = "6.0" ]; then
    # Register resource for sge_qmaster
    /usr/cluster/bin/scrgadm -a -j ${QMASTERS} -g ${MASTERRG} -t ${GDSTYPE} \
-x Start_command="/opt/SUNWscsge/bin/sge_qmaster6/start_sge_qmaster \
-R ${QMASTERS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-x Stop_command="/opt/SUNWscsge/bin/sge_qmaster6/stop_sge_qmaster \
-R ${QMASTERS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-x Probe_command="/opt/SUNWscsge/bin/sge_qmaster6/probe_sge_qmaster \
-R ${QMASTERS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-y Port_list=${MASTERPORT}/tcp -y Network_resources_used=${MASTERLH} \
-x Stop_signal=9 \

```

```

-x probe_timeout=90 -y Thorough_probe_interval=120 \
-y retry_count=2 -y retry_interval=900 \
-y Resource_dependencies=${MASTERHASP}
  St=$?
  if [ "${St}" -ne 0 ]; then
    echo "Error: Registration of resource ${QMASTERRS} failed."
    echo "Please correct the wrong parameters."
    exit 1
  else
    echo "Registration of resource ${QMASTERRS} succeeded"
  fi

  # Register resource for sge_schedd
  /usr/cluster/bin/scrgadm -a -j ${SCHEDDRS} -g ${MASTERRG} -t ${GDSTYPE} \
-x Start_command="/opt/SUNWscsge/bin/sge_schedd6/start_sge_schedd \
-R ${SCHEDDRS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-x Stop_command="/opt/SUNWscsge/bin/sge_schedd6/stop_sge_schedd \
-R ${SCHEDDRS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-x Probe_command="/opt/SUNWscsge/bin/sge_schedd6/probe_sge_schedd \
-R ${SCHEDDRS} -G ${MASTERRG} -S ${SGE_ROOT} -C ${SGE_CELL}" \
-y Port_list=${MASTERPORT}/tcp -y Network_resources_used=${MASTERLH} \
-x Stop_signal=9 \
-x probe_timeout=90 -y Thorough_probe_interval=120 \
-y retry_count=2 -y retry_interval=900 \
-y Resource_dependencies=${QMASTERRS}
  St=$?
  if [ "${St}" -ne 0 ]; then
    echo "Error: Registration of resource ${SCHEDDRS} failed."
    echo "Please correct the wrong parameters and try again."
    echo "Removing resource ${QMASTERRS}."
    /usr/cluster/bin/scrgadm -r -j ${QMASTERRS}
    exit 1
  else
    echo "Registration of resource ${SCHEDDRS} succeeded"
  fi
fi
else
  echo "Fatal: Please set variable SGE_VER properly in 'dirname $0'/sge_config!"
  exit 1
fi

```

## Listing of sge\_remove

```

#!/bin/ksh
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident    "@(#)sge_remove.ksh    1.3    06/04/24 SMI"
#

MYCONFIG=""
MYNAME='basename $0'

typeset opt

while getopts 'f:' opt

```

```
do
  case "${opt}" in
    f) MYCONFIG=${OPTARG};;
    *) echo "ERROR: ${MYNAME} Option ${OPTARG} unknown - early End."
       echo "Only the option -f <filename> is valid."
       exit 1;;
  esac
done

if [ -n "${MYCONFIG}" ] && [ -f "${MYCONFIG}" ]
then
  echo "sourcing ${MYCONFIG}"
  . ${MYCONFIG}
else
  PKGCONF='dirname $0'/sge_config
  echo "sourcing ${PKGCONF}"
  . ${PKGCONF}
fi

/usr/cluster/bin/scswitch -n -j ${SCHEDDRS}
/usr/cluster/bin/scswitch -n -j ${QMASTERRS}

/usr/cluster/bin/scrgadm -r -j ${SCHEDDRS}
/usr/cluster/bin/scrgadm -r -j ${QMASTERRS}
```

# Index

---

## A

administrative user account, 19

## B

Berkley DB spooling server, Oracle Grid Engine  
application, 14

binary files

- determining location, 17
- disk space requirements, 15
- distribution among file systems, 17

boot-time options

- restrictions, 14
- setting, 22

## C

`clnode` command, 9

commands, node information, 8

`config` file, 34

configuration parameters, Oracle Grid  
Engineresources, 29–31

configuring

- file systems, 26–27
- HA for Oracle Grid Engine
  - overview, 13
  - performing, 28–32
  - planning, 14–18

HASStoragePlus resource, 26–27

Oracle Grid Engine application, 20–23

configuring (*Continued*)

Oracle Solaris Cluster HA for NFS, 27–28

## D

daemons, shadow, 14

debugging, HA for Oracle Grid Engine, 34–35

dependencies, Oracle Grid Engine components, 16

device groups, 26

`dfstab` file, 28

directories

`/opt/SUNWscsge/util`, 28

spool

- disk space requirements, 15
- distribution among file systems, 17

disk space requirements, 15

disks, preparing for Oracle Grid Engine

installation, 18–20

## E

environment variables

Oracle Grid Engine root directory, 21, 23

`/etc/services` file, 19

extension properties

- effect on fault monitors, 33
- HA for Oracle Grid Engine, 32

**F**

fault monitors, tuning, 33

file systems

configuring, 26–27

distribution of files on, 17

Oracle Grid Engine

configuration requirements, 16

creating root directory, 19

environment variables, 21, 23

installing software on, 21

performance, 17

requirements, 16

files

config, 34

determining location, 17

dfstab, 28

/etc/services, 19

/opt/SUNWscsge/etc/config, 34

services, 19

settings.sh, 23

sgc\_config

listing, 37–38

syntax, 29–31

sgc\_register

listing, 38–39

running, 31

sgc\_remove, 39–40

system configuration, 18

**G**

global zone, 16

grid configuration, effect on file systems, 17

group IDs, 21

**H**

HA for Oracle Grid Engine

overview, 12

configuration

performing, 28–32

planning, 14–18

with HASStoragePlus resource, 26–27

HA for Oracle Grid Engine (*Continued*)

debugging, 34–35

fault monitors, 33

installing, 24–26

overview, 13

verifying installation, 32–33

service configuration, 16

software packages, installing, 24–26

HASStoragePlus resource

configuring, 26–27

creating, 27

enabling, 31

help, 8–9

hosts, file system requirements, 17

**I**

installing

HA for Oracle Grid Engine, 24–26

overview, 13

verifying installation, 32–33

Oracle Grid Engine application, 20–23

**J**

jobs, memory requirements, 15

**L**

location, Oracle Grid Engine files, 17

logical host names

resource group for

creating, 22

planning, 18

resources

creating, 22

enabling, 31

planning, 18

**M**

management tier, operating system for, 15  
 memory requirements, 15  
 messages file, 9

**N**

network addresses, *See* logical host names  
 NFS file system  
   configuration requirements, 16  
   distribution of files on, 17  
   performance, 17  
   resources  
     configuring, 27–28  
     enabling, 31  
 nodes, preparing for Oracle Grid Engine  
   installation, 18–20  
 non-global zones, 16

**O**

operating systems, requirements on, 15  
 /opt/SUNWscsge/util directory, 28  
 /opt/SUNWscsge/etc/config file, 34  
 Oracle Grid Engine, Berkley DB spooling server, 14  
 Oracle Grid Engine application  
   administrative user account, 19  
   configuring, 20–23  
   dependencies between components, 16  
   disk space requirements, 15  
   disks, 18–20  
   fault monitor, 33  
   file system, 16  
     creating root directory, 19  
     environment variables, 21, 23  
     installing software on, 21  
   group IDs, 21  
   installing, 20–23  
   location of binary files, 17  
   memory requirements, 15  
   nodes, 18–20  
   resource group for  
     creating, 22

Oracle Grid Engine application, resource group for  
 (*Continued*)

  enabling, 31  
   planning, 18  
   resources  
     configuration parameters, 29–31  
     enabling, 31  
     extension properties, 32  
     planning, 18  
     shadow daemon, 14  
     software versions, 15  
     start at boot option, 14, 22  
     verifying installation and configuration, 23–24  
 Oracle Solaris Cluster HA for NFS  
   configuration requirements, 16  
   configuring, 27–28

**P**

packages, 24–26  
 performance, NFS file system, 17  
 port numbers  
   sge\_execd service, 19  
   sge\_qmaster service, 19  
 protocols  
   sge\_execd service, 19  
   sge\_qmaster service, 19  
 prtconf -v command, 9  
 prtdiag -v command, 9  
 psrinfo -v command, 9

**R**

registering  
   SUNW.gds resource type, 31  
   SUNW.HAStoragePlus resource type, 26  
   SUNW.nfs resource type, 27  
 resource groups  
   creating, 22  
   enabling, 31  
   planning, 18  
 resource types  
   fault monitors, 33

resource types (*Continued*)

- SUNW.gds
    - extension properties, 32
    - registering, 31
  - SUNW.HAStoragePlus
    - dependencies on, 16
    - registering, 26
  - SUNW.nfs
    - registering, 27
- resources
- dependencies between, 16
  - HAStoragePlus
    - creating, 27
    - enabling, 31
  - logical host name
    - creating, 22
    - enabling, 31
    - planning, 18
  - NFS file system
    - configuring, 27–28
    - enabling, 31
  - Oracle Grid Engine application
    - debugging, 34–35
    - enabling, 31
    - extension properties, 32
    - planning, 18
  - Oracle Grid Engineapplication
    - configuration parameters, 29–31
- restrictions
- HA for Oracle Grid Engine configuration, 14
  - zones, 16
- root directory
- creating, 19
  - environment variables, 21, 23
  - installing software on, 21

**S**

- services file, 19
- settings.sh file, 23
- sgc\_config file
  - listing, 37–38
  - syntax, 29–31
- sgc\_execd service, 19

- sgc\_qmaster service, 19
- sgc\_register file
  - listing, 38–39
  - running, 31
- sgc\_remove file, 39–40
- SGE\_ROOT environment variable, 21, 23
- sgcadmin user account, 19
- shadow daemon, Oracle Grid Engine application, 14
- share command, 28
- show-rev subcommand, 9
- showrev -p command, 9
- software packages, 24–26
- spool directories
  - disk space requirements, 15
  - distribution among file systems, 17
- spooling server, Berkeley DB, 14
- starting
  - Oracle Grid Engine application, 14, 22
- SUNW.gds resource type
  - extension properties, 32
  - registering, 31
- SUNW.HAStoragePlus resource type
  - dependencies on, 16
  - registering, 26
- SUNW.nfs resource type, registering, 27
- system configuration files, location, 18
- system properties, effect on fault monitors, 33

**T**

- TCP (Transmission Control Protocol), 19
- technical support, 8–9
- Transmission Control Protocol (TCP), 19
- tuning, fault monitors, 33

**U**

- user accounts, 19

**V**

- /var/adm/messages file, 9

verifying, Oracle Grid Engine application, 23–24  
version requirements, Oracle Grid Engine  
application, 15

## **Z**

zones, 16

