N1 Grid Engine 6 Installation Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817–6118–10
June 2004

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, N1 and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des mardques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

040609@9061

Contents

Preface 7

1 Before You Install the Software 11
Plan the Installation 11
Decisions That You Must Make 11
Gather the Necessary Information 12
Loading the Distribution Files On a Workstation 22
▼ How to Load the Distribution Files On a Workstation 22

2 Installing the Grid Engine Software Interactively 27
Interactive Installation Overview 27
Performing an Installation 28
▼ How to Install the Master Host 28
▼ How to Install Execution Hosts 41
Registering Administration Hosts 45
Registering Submit Hosts 46
▼ How to Install the Berkeley DB Spooling Server 46

3 Automating the Installation Process 51
Using the `inst_sge` Utility and a Configuration Template 51
▼ How to Use the `inst_sge` Utility to Automate the Installation Process 51

4 Installing the Increased Security Features 57
Why Install the Increased Security Features? 57

Additional Setup Required 57
▼ How to Install a CSP-Secured System 58
▼ How to Generate Certificates and Private Keys for Users 61
Checking Certificates 63

5 Upgrading from a Previous Release of Grid Engine Software 65
About Upgrading the Software 65
▼ How to Upgrade the Software 65

6 Verifying the Installation 79
Verifying the Installation 79
▼ How to Verify That the Daemons Are Running on the Master Host 79
▼ How to Verify That the Daemons Are Running on the Execution Hosts 80
▼ How to Run Simple Commands 81
▼ How to Submit Test Jobs 82

7 Removing the Grid Engine Software 85
Removing the Software Interactively 85
▼ How to Remove the Software Interactively 85
Removing the Software Using the `inst_sge` Utility and a Configuration Template 86
▼ How to Remove the Software Interactively 86

8 Installing the Accounting and Reporting Console 89
Setting Up the Database Software 89
Set up the PostgreSQL Database Software 89
▼ How to Set Up a PostgreSQL Database 92
▼ How to Set Up an Oracle Database 93
Install the Accounting and Reporting Software 95
▼ How to Set Up `dbWriter` 95
▼ How to Install Sun Web Console 99
▼ How to Install the Accounting and Reporting Console 100

Glossary 103

Index 107

Preface

The *N1 Grid Engine 6 Installation Guide* describes the procedures for installing and upgrading the N1™ Grid Engine 6 software.

Who Should Use This Book

This manual is for engineers and system administrators, who need to install the N1 Grid Engine 6 software. Installation of the N1 Grid Engine 6 software should be done by someone who is familiar with the administration of the platform on which the software is installed.

Before You Read This Book

While this book contains a complete glossary, if you are new to N1 Grid Engine 6 software, you should familiarize yourself with the concepts that are discussed in Chapter 1, “Introduction to the N1 Grid Engine 6 Software,” in *N1 Grid Engine 6 User’s Guide*.

How This Book Is Organized

Chapter 1 describes each of decisions you must make in the planning process, and how to extract the installation files from CD-ROM.

Chapter 2 describes the procedure for installing the software manually.

Chapter 3 describes the files and commands necessary to automate your installations.

Chapter 4 describes the procedure for installing CSP-based encryption.

Chapter 5 describes the upgrade procedure.

Chapter 6 describes how you can ensure that the software is installed correctly.

Chapter 7 describes the procedure to remove the software from your systems.

Chapter 8 describes the procedure for installing the optional accounting and reporting module.

Glossary is a list of product-specific words and phrases and their definitions.

Related Books

Other books in the N1 Grid Engine 6 software documentation collection include:

- *N1 Grid Engine 6 User's Guide*
- *N1 Grid Engine 6 Administration Guide*
- *N1 Grid Engine 6 Release Notes*

Accessing Sun Documentation Online

The docs.sun.com[SM] Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see "Buy printed documentation" at `http://docs.sun.com`.

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
AaBbCc123	Command-line placeholder: replace with a real name or value	To delete a file, type **`rm`** *filename.*
AaBbCc123	Book titles, new terms, or terms to be emphasized	Read Chapter 6 in *User's Guide.* These are called *class* options. You must be *root* to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name#`
Bourne shell and Korn shell prompt	`$`
Bourne shell and Korn shell superuser prompt	`#`

CHAPTER 1

Before You Install the Software

This chapter describes the steps that need to be taken before you install the N1™ Grid Engine 6 software (*grid engine software*).

This chapter includes the following topics:

- "Plan the Installation" on page 11
- "Loading the Distribution Files On a Workstation" on page 22

Plan the Installation

Whether you have installed previous versions of the grid engine software or this is your first time, you will need to do some planning before you extract and install the software. This section describes the decisions you must make, and, wherever possible, gives you criteria on which you can base your decisions.

Decisions That You Must Make

You must make several decisions before you can plan the installation:

- Decide whether your system of networked computer hosts that run N1 Grid Engine 6 software (*grid engine system*) is to be a single cluster or a collection of sub-clusters, called *cells*. Cells allow you to install separate instances of the grid engine software but share the binary files across those instances.
- Select the machines that are to be grid engine system hosts. Determine the host type of each machine: *master host*, *shadow master host*, *administration host*, *submit host*, *execution host*, or a combination.
- Ensure that all users of the grid engine system have the same user names on all submit and execution hosts.

- Decide how to order grid engine software directories. For example, you could organize directories as a complete tree on each workstation, or you could cross-mount directories, or you could set up a partial directory tree on some workstations. You must also decide where to locate each grid engine software installation directory, *sge-root*.
- Decide on the site's queue structure.
- Determine whether to define network services as an NIS file or as local to each workstation in `/etc/services`.
- Use the information in this chapter to gather the information necessary to complete the installation worksheet.

Gather the Necessary Information

Before you install the grid engine software, you must plan how to achieve the results that fit your environment. This section helps you make the decisions that affect the rest of the procedure. Write down your installation plan in a table similar to the following example.

Parameter	Value
sge-root directory	
Cell name	
Administrative user	
`sge_qmaster` port number	
`sge_execd` port number	
Master host	
Shadow master hosts	
Execution hosts	
Administration hosts	
Submit hosts	
Group ID range for jobs	
Spooling mechanism (Berkeley DB or Classic spooling)	
Berkeley DB server host (the master or another host)	

Parameter	Value
Berkeley DB spooling directory on the database server	
Scheduler tuning profile (Normal, High, Max)	
Installation method (interactive, secure, automated, or upgrade)	

Disk Space Requirements

The grid engine software directory tree has the following fixed disk space requirements:

- 40 Mbytes for the installation files (including documentation) without any binaries
- Between 10 and 15 Mbytes for each set of binaries

The ideal disk space for grid engine system spool directories is as follows:

- 10-200 Mbytes for the master host spool directories
- 10-200 Mbytes for the Berkeley DB spool directories

Note – The spool directories of the master host and of the execution hosts are configurable and need not reside under the default location, *sge-root*.

---

sge-root Installation Directory

Create a directory into which you will load the contents of the distribution media. This directory is called the *root directory*, or *sge-root*. When the grid engine system is running, this directory stores the current cluster configuration and all other data that must be spooled to disk.

Use a valid path name for the directory that is network accessible on all hosts. For example, if the file system is mounted using automounter, set *sge-root* to `/usr/N1GE6`, not to `/tmp_mnt/usr/N1GE6`. Throughout this document, the *sge-root* variable is used to refer to the installation directory.

sge-root is the top level of the grid engine software directory tree. Each grid engine system component in a cell needs read access to the *sge-root*/*cell*/`common` directory, on startup. When grid engine software is installed as a single cluster, the value of *cell* is `default`.

For ease of installation and administration, this directory should be readable on all hosts on which you intend to run the grid engine software installation procedure. For example, you can select a directory available across a network file system, such as

NFS. If you choose to select file systems that are local to the hosts, you must copy the installation directory to each host before you start the installation procedure for the particular machine. See "File Access Permissions" on page 16 for a description of required permissions.

Directory Organization

When determining the directory organization, you must decide the following:

- The directory organization (for example, whether you will install a complete software tree on each workstation, directories cross-mounted, or a partial directory tree on some workstations, and so on)
- Where to locate each root directory, *sge-root*

Note – Because changing the installation directory or the spool directories requires a new installation of the system, use extra care to select a suitable installation directory up front. Note that all important information from a previous installation can be preserved.

By default, the installation procedure installs the grid engine software, manuals, spool areas, and the configuration files in a directory hierarchy under the installation directory (see Figure 1–1). If you accept this default behavior, you should install or select a directory with the access permissions that are described in "File Access Permissions" on page 16.

You can select the spool areas to put in other locations during the primary installation. See "Configuring Queues" in *N1 Grid Engine 6 Administration Guide* for instructions.

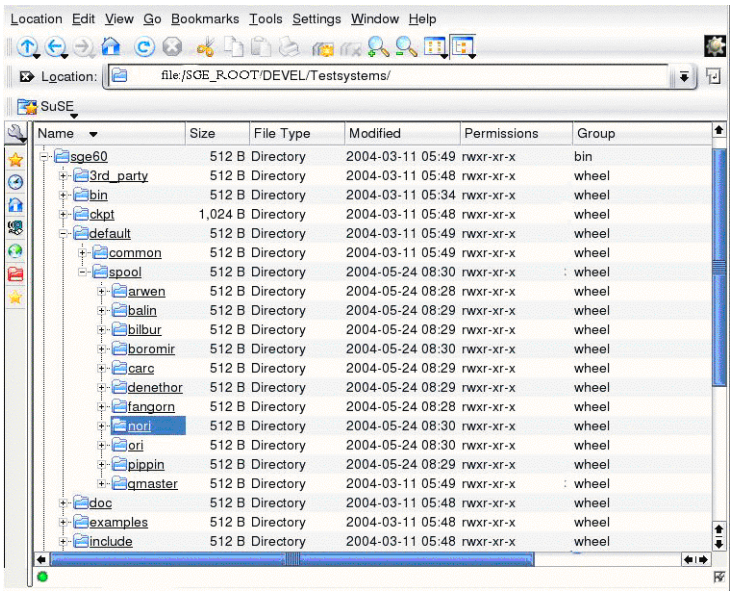

FIGURE 1–1 Sample Directory Hierarchy

Cells

You can set up the grid engine system as a single cluster or as a collection of loosely coupled clusters called *cells*. The `SGE_CELL` environment variable indicates the cluster being referenced. When the grid engine system is installed as a single cluster, `$SGE_CELL` is not set, and the value `default` is assumed for the cell value.

User Names

In order for the grid engine system to verify that users submitting jobs have permission to submit them on the desired execution hosts, users' names must be identical on the submit and execution hosts involved. You may therefore have to change user names on some machines, because grid engine system users map directly to system user accounts.

Note – User names on the master host are not relevant for permission checking. These user names do not have to match or even exist.

Installation Accounts

You can install the grid engine software either as the root user or as an unprivileged user (for example, your own user account). However, if you install the software logged in as an unprivileged user, the installation allows only that user to run grid engine system jobs. Access is denied to all other accounts. Installing the software logged in as the root account resolves this restriction. However, root permission is required for the complete installation procedure. Also, if you install as an unprivileged user, you are not allowed to use the `qrsh`, `qtcsh`, `qmake` commands, nor can you run tightly integrated parallel jobs.

File Access Permissions

If you install the software logged in as root, you might have a problem configuring root read/write access for all hosts on a shared file system. Therefore, you might have problems putting *sge-root* onto a network-wide file system.

You can force grid engine software to run all grid engine system components through a non-root administrative user account (called `sgeadmin`, for example). With this setup, you need only read/write access to the shared *sge-root* file system for this particular user.

The installation procedure asks whether files should be created and owned by an administrative user account. If you answer “Yes” and provide a valid user name, files are created by this user. Otherwise, the user name under which you run the installation procedure is used. It is recommended that you create an administrative user, and answer “ Yes” to this question.

Make sure in all cases that the account used for file handling on all hosts has read/write access to the *sge-root* directory. Also, the installation procedure assumes that the host from which you access the grid engine software distribution media can write to the *sge-root* directory.

Network Services

Determine whether your site’s network services are defined in an NIS database or in an `/etc/services` file that is local to each workstation. If your site uses NIS, find out the host name of your NIS server is so that you can add entries to the NIS “services” map.

The grid engine system services are `sge_execd` and `sge_qmaster`. To add the services to your NIS map, choose reserved, unused port numbers. The following examples show `sge_qmaster` and `sge_execd` entries.

```
sge_qmaster 536/tcp

sge_execd 537/tcp
```

Master Host

The master host controls the grid engine system. This host runs the master daemon `sge_qmaster`, and the scheduling daemon, `sge_schedd`.

The master host must comply with the following requirements:

- The host must be a stable platform.
- The host must *not* be excessively busy with other processing.
- At least 60 – 120 Mbytes of unused main memory must be available to run the grid engine system daemons. For very large clusters that include many hundreds or thousands of hosts and tens of thousands of jobs in the system at any time, 1 GByte or more of unused main memory might be required and two CPUs might be beneficial.
- The master host must be installed before shadow master execution, administration, or submit hosts.
- (Optional) The grid engine software directory, *sge-root*, should be installed locally, to cut down on network traffic.

Shadow Master Hosts

These hosts back up the functionality of `sge_qmaster` in case the master host or the master daemon fails. To be a shadow master host, a machine must have the following characteristics:

- It must run `sge_shadowd`.
- It must share `sge_qmaster` status, job information, and queue configuration information that is logged to disk. In particular, the shadow master hosts need read/write root or administration user access to the `sge_qmaster` spool directory and to the *sge-root*`/`*cell*`/common` directory.
- The *sge-root*`/`*cell*`/common/shadow_masters` file must contain a line defining the host as a shadow master host.

Note – If no cell name is specified during installation, the value of *cell* is `default`.

The shadow master host facility is activated for a host as soon as these conditions are met. You do not need to restart the grid engine system daemons to make a host into a shadow master host.

Spool Directories Under the Root Directory

During the installation of the master host, you must specify the location of a spooling directory. This directory is used to spool jobs from execution hosts that do not have a local spooling directory.

- On the master host, spool directories are maintained under *qmaster-spool-dir*. The location of *qmaster-spool-dir* is defined during the master host installation process. The default value of *qmaster-spool-dir* is *sge-root*/*cell*/`spool/qmaster`.
- On each execution host, a spool directory called *execd-spool-dir* is defined during the execution host installation processes. The default value of *execd-spool-dir* is *sge-root*/*cell*/`spool`/*exec-host*. You will get better performance from execution hosts with local spooling directories than from execution hosts that have NFS mounted the master host's spooling directory.

Note – If no cell name is specified during installation, the value of *cell* is `default`.

You do not need to export these directories to other machines. However, exporting the entire *sge-root* tree and making it write-accessible for the master host and all executable hosts makes administration easier.

Choosing Between Classic Spooling and Database Spooling

During the installation, you are given the option to choose between classic spooling and Berkeley DB spooling server. If you choose Berkeley DB spooling, you are then given the option to spool to a local directory or to a separate host, known as a Berkeley DB spooling server.

While classic spooling is an option, you should see better performance using a Berkeley DB spooling server. Part of this performance increase is because the master host can make non-blocking writes to the database, but has to make blocking writes to the text file used by classic spooling. Other factors that may influence your decision are file format and data integrity. There is a greater level of data integrity when you write to the Berkeley DB, rather than to a text file. However, a text file stores data in a format that humans can read and edit. Normally, you don't need to read these files, but the spooling directory contains the messages from the system daemons, which can be useful during debugging.

Database Server and Spooling Host

The master host can store its configuration and state to a Berkeley DB spooling database. The spooling database can be installed on the master server or on a separate host. When the Berkeley DB spools into a local directory on the master host, the performance is better, but if you want to set up a shadow master host, you need to use a separate Berkeley DB spooling server (host). In this case, you have to choose a host with a configured RPC service. The master host connects through RPC to the Berkeley DB.

Note – Using a shadow master host is more reliable, but using a separate Berkeley DB spooling host results in a potential security hole. RPC communication as used by the Berkeley DB can be easily compromised. Only use this alternative if your site is secure and if users can be trusted to access the Berkeley DB spooling host by means of TCP/IP communication.

If you choose to use Berkeley DB spooling without a shadow master, you don't need to set up a separate spooling server. Likewise, if you choose *not* to use Berkeley DB spooling, you can set up a shadow master host without setting up set up a separate spooling server.

Once you determine whether you need a separate spooling server, you will also need to determine the location for the spooling directory. The spooling directory must be local to the spooling server. A default value for the location of the spooling directory is recommended during installation.

The requirements for the Berkeley DB spooling host are similar to the requirements for the master host:

- The host must be a stable platform.
- The host must *not* be excessively busy with other processing.
- At least 60 – 120 Mbytes of unused main memory must be available to run the grid engine system daemons. For very large clusters that include many hundreds or thousands of hosts and tens of thousands of jobs in the system at any time, 1 GByte or more of unused main memory may be required and two CPUs may be beneficial.
- (Optional) A separate spooling host must be installed before the master host.
- (Optional) The directory, *sge-root*, should be installed locally, to cut down on network traffic.

Execution Hosts

Execution hosts run the jobs that users submit to the grid engine system. An execution host must first be set up as an administration host. You run an installation script on each execution host.

Group IDs

You need to provide a range of IDs that will be assigned dynamically for jobs. The range must be big enough to provide enough numbers for the maximum number of grid engine system jobs running at a single moment on a single host.

A group ID is assigned to each grid engine system job to monitor the resource utilization of the job. Each job will be assigned a unique ID during the time it is running. For example, a range of 20000-20100 allows 100 jobs to run concurrently on a single host. You can change the group ID range for your cluster configuration at any time, but the values in the UNIX group ID range must be unused on your system.

Administration Hosts

Operators and managers of the grid engine system use administration hosts to perform administrative tasks such as reconfiguring queues or adding grid engine system users.

The master host installation script automatically makes the master host an administration host. During the master host installation process, you can add other administration hosts. You can also manually add administration hosts on the master host at any time after installation.

Submit Hosts

Jobs can be submitted and controlled from submit hosts. The master host installation script automatically makes the master host a submit host.

Cluster Queues

The installation procedure creates a default cluster queue structure, which is suitable for getting acquainted with the system. The default queue can be removed after installation.

Note – No matter what directory is used for the installation of the software, the administrator can change most settings that were created by the installation procedure. This change can be made while the system is running.

There are several factors you need to consider when determining a queue structure.

- Whether you need cluster queues for sequential, interactive, parallel, and other job types
- Which queue instances to put on which execution hosts
- How many job slots are needed in each queue

For more detailed information on administering cluster queues, see “Configuring Queues” in *N1 Grid Engine 6 Administration Guide*.

Scheduler Profiles

You can choose from three scheduler profiles during the installation process: `normal`, `high`, and `max`. You can use these predefined profiles as a starting point for grid engine tuning.

Using these profiles, you can optimize the scheduler for one or more of the following:

- The amount of information about a scheduling run
- The load adjustment during a scheduling run
- Interval scheduling (the default) or immediate scheduling

You can choose from three scheduler profiles:

- `normal`— This profile uses load adaptation, and interval scheduling, and reports all the information that the scheduler gathers during the dispatch cycle. This profile is the starting point for most grids. Use this profile if your highest priority is gathering and reporting information about a scheduling run.
- `high`— This profile is more appropriate for a large cluster, where throughput is more important than gathering and reporting all the information from the scheduler. This profile also uses interval scheduling. Use this profile if you want to get better performance at the cost of getting less information about your scheduling runs.
- `max`— This profile disables all information gathering and reporting, enables immediate scheduling, and disables load adaptation. Immediate scheduling is very useful for sites with high throughput and very short running jobs. The advantage of immediate scheduling decreases as run time of the jobs increases. This profile can be used in clusters of any size where only throughput is important and everything else is a lower priority.

For more information on how to configure scheduling, see "Administering the Scheduler" in *N1 Grid Engine 6 Administration Guide*.

Installation Method

Several methods are available for installing the grid engine software:

- Interactive
- Interactive, with increased security
- Automated, using the `inst_sge` script and a configuration file
- Upgrade

To decide which installation method you should use, consider the following factors.

- Do you already have the grid engine software installed and running?

  If so, you'll probably want to upgrade. The upgrade process is described in Chapter 5.

  If not, the master host installation is only done once, so the master host is typically installed interactively, as described in Chapter 2.

- Do you need to install just a few execution hosts?

  If so, then you will probably want to install them interactively, as described in Chapter 2.
- Do you need to install a large number of execution hosts?

  If so, then you might want to perform automated installation, using the `inst_sge` script and a configuration file. This process is described in "Using the `inst_sge` Utility and a Configuration Template" on page 51.
- Do you require your grid to use encryption?

  If so, you have to perform an interactive installation with increased security. This process is described in Chapter 4.

Loading the Distribution Files On a Workstation

Grid Engine is distributed on CD-ROM, and through electronic download. For information on how to access CD-ROMs, ask your system administrator or refer to your local system documentation. The CD-ROM distribution contains a directory named `N1_Grid_Engine_6`. The product distribution is in this directory, in both `tar.gz` format and the `pkgadd` format. The `pkgadd` format is provided for the Solaris™ Operating System (*Solaris OS*). For all supported operating systems, the software is distributed in `tar.gz` format.

▼ How to Load the Distribution Files On a Workstation

Before You Begin

Ensure that the file systems and directories that are to contain the grid engine software distribution and the spool and configuration files are set up properly by setting the access permissions as defined in "File Access Permissions" on page 16.

Steps

1. **Provide access to the distribution media.**

   - **If you downloaded the software, rather than getting it on CD-ROM, just unzip the five files into a directory. This directory must be located on a file system that has at least 350 MBytes free space.**

2. **Log in to a system, preferably a system that has a direct connection to a file server.**

3. **Create the installation directory as described in "*sge-root* Installation Directory" on page 13.**

```
# mkdir  /opt/n1ge6
```

In these instructions, the installation directory is abbreviated as *sge-root*.

4. **Install the binaries for all binary architectures that are to be used by any of your master, execution, and submit hosts in your grid engine system cluster.**

   - **The `pkgadd` Method**

     The `pkgadd` format is provided for the Solaris Operating System. To facilitate remote installation, the `pkgadd` directories are also provided in zip files.

     You can install the following packages:

     - `SUNWsgeec` – Architecture independent files
     - `SUNWsgeed` – Documentation
     - `SUNWsgee` – Solaris (SPARC® platform) 32-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
     - `SUNWsgeex` – Solaris (SPARC platform) 64-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
     - `SUNWsgeei`— Solaris (x86 platform) binaries for Solaris 8 and Solaris 9 Operating Systems

     As you type the following commands, you must be prepared to respond to script questions about your base directory, *sge-root* and the administrative user. The script requests the choices you made during the planning steps of this installation. See "Decisions That You Must Make" on page 11.

     At the command prompt, type the following commands, responding to the script questions.

     ```
     # cd cdrom_mount_point/N1_Grid_Engine_6
     # pkgadd -d ./Common/Packages/SUNWsgeec
     # pkgadd -d ./Docs/Packages/SUNWsgeed
     # pkgadd -d ./Solaris_sparc/Packages/SUNWsgee      (This is optional; at least one binary set is required)
     # pkgadd -d ./Solaris_sparc/Packages/SUNWsgeex     (This is optional; at least one binary set is required)
     # pkgadd -d ./Solaris_x86/Packages/SUNWsgeei       (This is optional; at least one binary set is required)
     ```

   - **The `tar` Method**

     For all supported operating systems, the software is distributed in `tar.gz` format.

     The following table contains files that you will need to install, regardless of what platform you are on.

File	*Description*
`Common/tar/n1ge-6_0-common.tar.gz`	Architecture independent files

`Docs/tar/n1ge-6_0-doc.tar.gz`	Release notes, installation guide, user guide and administration guide

The tar files that contain platform-specific binaries use the naming convention of `n1ge-6_0-bin-`*architecture*`.tar.gz`. The following table lists the platform-specific binaries. You need to install the file for each platform that you need to support. Note that each platform has its own directory under `N1_Grid_Engine_6`.

Platform-Specific File	*Platform*
`Solaris_sparc/tar/n1ge-6_0-bin-solaris.tar.gz`	Solaris (SPARC platform) 32-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
`Solaris_sparc/tar/n1ge-6_0-bin-solaris-sparcv9.tar.gz`	Solaris (SPARC platform) 64-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
`Solaris_x86/tar/n1ge-6_0-bin-solaris-i586.tar.gz`	Solaris (x86 platform) binaries for Solaris 8, and Solaris 9 Operating Systems
`Linux24_i586/tar/n1ge-6_0-bin-linux24-i586.tar.gz`	Linux (x86 platform) binaries for the 2.4 kernel
`Linux24_amd64/tar/n1ge-6_0-bin-linux24-amd64.tar.gz`	Linux (AMD platform) binaries for the 2.4 kernel
`MacOSX/tar/n1ge-6_0-bin-darwin.tar.gz`	Apple Mac OS/X
`HPUX11/tar/n1ge-6_0-bin-hp11.tar.gz`	Hewlitt Packard HP-UX 11
`Aix51/tar/n1ge-6_0-bin-irix65.tar.gz`	SGI Irix 6.5
`Aix43/tar/n1ge-6_0-bin-aix51.tar.gz`	IBM AIX 5.1
`Irix65/tar/n1ge-6_0-bin-aix43.tar.gz`	IBM AIX 4.3

Type the following commands at the command prompt. In the example, *basedir* is the abbreviation for the full directory, *cdrom_mount_point*`/N1_Grid_Engine_6`.

```
% su
# cd sge-root
# gzip -dc basedir/Common/tar/n1ge-6_0-common.tar.gz | tar xvpf -
# gzip -dc basedir/Docs/tar/n1ge-6_0-doc.tar.gz | tar xvpf -
# gzip -dc basedir/Solaris_sparc/tar/n1ge-6_0-bin-solsparc32.tar.gz | tar xvpf -
# gzip -dc basedir/Solaris_sparc/tar/n1ge-6_0-bin-solsparc64.tar.gz | tar xvpf -
# SGE_ROOT=sge-root; export SGE_ROOT
```

```
# util/setfileperm.sh $SGE_ROOT
```

CHAPTER 2

Installing the Grid Engine Software Interactively

This chapter discusses the steps that you must take to manually install the grid engine software. Please note that this chapter does not cover the upgrade process or the installation of the Accounting and Reporting Module, ARCo. Those processes are covered in subsequent chapters.

Note – The instructions in this chapter assume that you are installing the software on a computer running the Solaris™ Operating System. Any difference in functionality created by other operating system architecture that grid engine software runs on is documented in files starting with the string `arc_depend_` in the *sge-root*`/doc` directory. The remainder of the file name indicates the operating system architectures to which the comments in the files apply, as in the `arc_depend_irix.asc` file.

This chapter discusses the following topics:

- "Interactive Installation Overview" on page 27
- "Performing an Installation" on page 28

Interactive Installation Overview

Note – The instructions in this section are for a *fresh* grid engine system only. For instructions on how to install a new system with additional security protection, see Chapter 4. For instructions on how to *upgrade* an existing installation of an earlier version of grid engine software, see Chapter 5.

Full installation includes the following tasks:

- Running an installation script on the master host and on every execution host in the grid engine system
- Registering information about administration hosts and submit hosts

Performing an Installation

This section includes instructions for performing the following tasks:

- "How to Install the Master Host" on page 28
- "How to Install Execution Hosts" on page 41
- "Registering Administration Hosts" on page 45
- "Registering Submit Hosts" on page 46
- "How to Install the Berkeley DB Spooling Server" on page 46

The following sections describe how to install all the components of the grid engine system, including the master, execution, administration, and submit hosts. If you need to install the system with enhanced security, see Chapter 4 before you continue installation.

▼ How to Install the Master Host

The master installation procedure creates the appropriate directory hierarchy required by `sge_qmaster` and `sge_schedd`. The procedure starts up the grid engine system daemons, `sge_qmaster` and `sge_schedd`, on the master host. The master host is also registered as a host with administrative and submit permission. The installation procedure creates a default configuration for the system on which it is run. The installation script queries the system for the type of operating system. The script then makes meaningful settings based on this information.

If, at any time during the installation, you think something went wrong, you can quit the installation procedure and restart it.

Before You Begin

If you have decided to use an administrative user, as described in "User Names" on page 15, you should create that user now. This procedure assumes that you have already extracted the grid engine software, as described in "Loading the Distribution Files On a Workstation" on page 22.

Steps

1. **Log in to the master host as `root`.**
2. **Ensure that you have set the `$SGE_ROOT` environment variable by typing:**

```
# echo $SGE_ROOT
```

- **If the `$SGE_ROOT` environment variable is not set, set it now, by typing:**

  # **`SGE_ROOT=`***sge-root*`; export SGE_ROOT`**

3. **Change to the installation directory.**

   - **If the directory where the installation files reside is visible from the master host, change directories (`cd`) to the installation directory** *sge-root***, and then proceed to Step 4.**

   - **If the directory is not visible and cannot be made visible, do the following:**

     a. **Create a local installation directory,** *sge-root***, on the master host.**

     b. **Copy the installation files to the local installation directory** *sge-root* **across the network (for example, by using `ftp` or `rcp`).**

     c. **Change directories (`cd`) to the local** *sge-root* **directory.**

4. **Type the `install_qmaster` command, adding the `-csp` flag if you are installing using the Certificate Security Protocol method described in Chapter 4.**

   This command starts the master host installation procedure. You are asked several questions, and you might be required to run some administrative actions.

```
% ./install_qmaster
Welcome to the Grid Engine installation
---------------------------------------

Grid Engine qmaster host installation
-------------------------------------

Before you continue with the installation please read these hints:

   - Your terminal window should have a size of at least
     80x24 characters

   - The INTR character is often bound to the key Ctrl-C.
     The term >Ctrl-C< is used during the installation if you
     have the possibility to abort the installation

The qmaster installation procedure will take approximately 5-10 minutes.

Hit <RETURN> to continue >>
```

5. **Choose an administrative account owner.**

```
Choosing Grid Engine admin user account
---------------------------------------

You may install Grid Engine that all files are created with the user id of an
unprivileged user.

This will make it possible to install and run Grid Engine in directories
where user >root< has no permissions to create and write files and directories.
```

```
   - Grid Engine still has to be started by user >root<

   - this directory should be owned by the Grid Engine administrator

Do you want to install Grid Engine
under an user id other than >root< (y/n) [y] >> y

Choosing a Grid Engine admin user name
--------------------------------------

Please enter a valid user name >> sgeadmin

Installing Grid Engine as admin user >sgeadmin<

Hit <RETURN> to continue >>
```

6. **Verify the *sge-root* directory setting.**

   In the following example, the value of *sge-root* is `/opt/n1ge6`.

```
Checking $SGE_ROOT directory
----------------------------

The Grid Engine root directory is:

   $SGE_ROOT = /opt/n1ge6

If this directory is not correct (e.g. it may contain an automounter
prefix) enter the correct path to this directory or hit <RETURN>
to use default [/opt/n1ge6] >>

Your $SGE_ROOT directory: /opt/n1ge6

Hit <RETURN> to continue >>
```

7. **Set up the TCP/IP services for the grid engine software.**

   a. **You will be notified if the TCP/IP services have not been configured.**

```
Grid Engine TCP/IP service >sge_qmaster<
----------------------------------------

There is no service >sge_qmaster< available in your >/etc/services< file
or in your NIS/NIS+ database.

You may add this service now to your services database or choose a port number.
It is recommended to add the service now. If you are using NIS/NIS+ you should
add the service at your NIS/NIS+ server and not to the local >/etc/services<
file.

Please add an entry in the form

   sge_qmaster <port_number>/tcp

to your services database and make sure to use an unused port number.
```

```
Please add the service now or press <RETURN> to go to entering a port number >>
```

b. Start a new terminal session or window to add the information `/etc/services` file or your NIS maps.

c. Add the correct ports to the `/etc/services` file or your NIS services map, as described in "Network Services" on page 16.

The following example shows how you might edit your `/etc/services` file.

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

Note – In this example, the entries for both `sge_qmaster` and `sge_execd` are added to `/etc/services`. Subsequent steps in this example assume that both entries have been made.

Save your changes.

d. Return to the window where the installation script is running.

```
Please add the service now or press <RETURN> to go to entering a port number >>
```

Press `<RETURN>`. You will see the following output:

```
sge_qmaster 536

Service >sge_qmaster< is now available.

Hit <RETURN> to continue >>

Grid Engine TCP/IP service >sge_execd<
-------------------------------------

Using the service

   sge_execd

for communication with Grid Engine.

Hit <RETURN> to continue >>
```

8. Enter the name of your cell.

The use of grid engine system cells is described in "Cells" on page 15.

```
Grid Engine cells
-----------------

Grid Engine supports multiple cells.

If you are not planning to run multiple Grid Engine clusters or if you don't
```

```
know yet what is a Grid Engine cell it is safe to keep the default cell name

   default

If you want to install multiple cells you can enter a cell name now.

The environment variable

   $SGE_CELL=<your_cell_name>

will be set for all further Grid Engine commands.

Enter cell name [default] >>
```

- **If you have decided to use cells, then enter the cell name now.**
- **If you have decided not to use cells, then press `<RETURN>` to continue.**

```
Using cell >default<.
Hit <RETURN> to continue >>
```

Press `<RETURN>` to continue.

9. **Specify a spool directory.**

   For guidelines on disk space requirements for the spool directory, see "Disk Space Requirements" on page 13. For information on where spool directory is installed, see "Spool Directories Under the Root Directory" on page 17.

```
Grid Engine qmaster spool directory
-----------------------------------

The qmaster spool directory is the place where the qmaster daemon stores
the configuration and the state of the queuing system.

The admin user >sgeadmin< must have read/write access
to the qmaster spool directory.

If you will install shadow master hosts or if you want to be able to start
the qmaster daemon on other hosts (see the corresponding sectionin the
Grid Engine Installation and Administration Manual for details) the account
on the shadow master hosts also needs read/write access to this directory.

The following directory

[/opt/n1ge6/default/spool/qmaster]

will be used as qmaster spool directory by default!

Do you want to select another qmaster spool directory (y/n) [n] >>
```

   - **If you want to accept the default spool directory, press `<RETURN>` to continue.**
   - **If you do *not* want to accept the default spool directory, then answer `y`.**

     In the following example the `/my/spool` directory is specified as the master host spool directory.

```
   Do you want to select another qmaster spool directory (y/n) [n] >> y

   Please enter a qmaster spool directory now! >>/my/spool
```

10. Set the correct file permissions.

```
Verifying and setting file permissions
--------------------------------------

Did you install this version with >pkgadd< or did you already
verify and set the file permissions of your distribution (y/n) [y] >> n

Verifying and setting file permissions
--------------------------------------

We may now verify and set the file permissions of your Grid Engine
distribution.

This may be useful since due to unpacking and copying of your distribution
your files may be unaccessible to other users.

We will set the permissions of directories and binaries to

   755 - that means executable are accessible for the world

and for ordinary files to

   644 - that means readable for the world

Do you want to verify and set your file permissions (y/n) [y] >> y

Verifying and setting file permissions and owner in >3rd_party<
Verifying and setting file permissions and owner in >bin<
Verifying and setting file permissions and owner in >ckpt<
Verifying and setting file permissions and owner in >examples<
Verifying and setting file permissions and owner in >install_execd<
Verifying and setting file permissions and owner in >install_qmaster<
Verifying and setting file permissions and owner in >mpi<
Verifying and setting file permissions and owner in >pvm<
Verifying and setting file permissions and owner in >qmon<
Verifying and setting file permissions and owner in >util<
Verifying and setting file permissions and owner in >utilbin<
Verifying and setting file permissions and owner in >catman<
Verifying and setting file permissions and owner in >doc<
Verifying and setting file permissions and owner in >man<
Verifying and setting file permissions and owner in >inst_sge<
Verifying and setting file permissions and owner in >bin<
Verifying and setting file permissions and owner in >lib<
Verifying and setting file permissions and owner in >utilbin<

Your file permissions were set

Hit <RETURN> to continue >>
```

11. Specify whether all of your grid engine system hosts are located in a single DNS domain.

```
Select default Grid Engine hostname resolving method
----------------------------------------------------

Are all hosts of your cluster in one DNS domain? If this is
the case the hostnames

   >hostA< and >hostA.foo.com<

would be treated as eqal, because the DNS domain name >foo.com<
is ignored when comparing hostnames.

Are all hosts of your cluster in a single DNS domain (y/n) [y] >>
```

- **If all of your grid engine system hosts are located in a single DNS domain, then answer `y`.**

  ```
  Are all hosts of your cluster in a single DNS domain (y/n) [y] >> y

  Ignoring domainname when comparing hostnames.

  Hit <RETURN> to continue >>
  ```

- **If all of your grid engine system hosts are *not* located in a single DNS domain, then answer `n`.**

  ```
  Are all hosts of your cluster in a single DNS domain (y/n) [y] >> n

  The domainname is not ignored when comparing hostnames.

  Hit <RETURN> to continue >>

  Default domain for hostnames
  ----------------------------

  Sometimes the primary hostname of machines returns the short hostname
  without a domain suffix like >foo.com<.

  This can cause problems with getting load values of your execution hosts.
  If you are using DNS or you are using domains in your >/etc/hosts< file or
  your NIS configuration it is usually safe to define a default domain
  because it is only used if your execution hosts return the short hostname
  as their primary name.

  If your execution hosts reside in more than one domain, the default domain
  parameter must be set on all execution hosts individually.

  Do you want to configure a default domain (y/n) [y] >>
  ```

  Press `<RETURN>` to continue.

  a. **If you want to specify a default domain, then answer `y`.**

     In the following example, `sun.com` is specified as the default domain.

     ```
     Do you want to configure a default domain (y/n) [y] >> y
     ```

```
Please enter your default domain >> sun.com

Using >sun.com< as default domain. Hit <RETURN> to continue >>
```

b. If you do *not* want to specify a default domain, then answer `n`.

In the following example, `sun.com` is specified as the default domain.

```
Do you want to configure a default domain (y/n) [y] >> n
```

12. Press `<RETURN>` to continue.

```
Making directories
------------------

creating directory: default/common
creating directory: /opt/n1ge6/default/spool/qmaster
creating directory: /opt/n1ge6/default/spool/qmaster/job_scripts
Hit <RETURN> to continue >>
```

13. Specify whether you want to use classic spooling or Berkeley DB.

For more information on how to determine the type of spooling mechanism you want, please see "Choosing Between Classic Spooling and Database Spooling" on page 18.

```
Setup spooling
--------------
Your SGE binaries are compiled to link the spooling libraries
during runtime (dynamically). So you can choose between Berkeley DB
spooling and Classic spooling method.
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >>
```

- **If you want to specify Berkeley DB spooling, press `<RETURN>` to continue.**

```
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >>

The Berkeley DB spooling method provides two configurations!

1) Local spooling:
The Berkeley DB spools into a local directory on this host (qmaster host)
This setup is faster, but you can't setup a shadow master host

2) Berkeley DB Spooling Server:
If you want to setup a shadow master host, you need to use
Berkeley DB Spooling Server!
In this case you have to choose a host with a configured RPC service.
The qmaster host connects via RPC to the Berkeley DB. This setup is more
failsafe, but results in a clear potential security hole. RPC communication
(as used by Berkeley DB) can be easily compromised. Please only use this
alternative if your site is secure or if you are not concerned about
security. Check the installation guide for further advice on how to achieve
failsafety without compromising security.

Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >>
```

- **If you want to use a use a Berkeley DB spooling server, enter `y`.**

```
Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >> y

Berkeley DB Setup

-----------------
Please, log in to your Berkeley DB spooling host and execute "inst_sge -db"
Please do not continue, before the Berkeley DB installation with
"inst_sge -db" is completed, continue with <RETURN>
```

Note – *Do not* press `<RETURN>` until you have completed the Berkeley DB installation on the spooling server.

---

a. **Start a new terminal session or window.**

b. **Log in to the spooling server.**

c. **Install the software, as described in "How to Install the Berkeley DB Spooling Server" on page 46.**

d. **After you have installed the software on the spooling server, return to the master installation window, and press `<RETURN>` to continue.**

e. **Enter the name of the spooling server.**

   In the following example, `vector` is the host name of the spooling server.

```
Berkeley Database spooling parameters
-------------------------------------

Please enter the name of your Berkeley DB Spooling Server! >> vector
```

f. **Enter the name of the spooling directory.**

   In the following example, `/opt/n1ge6/default/spooldb` is the spooling directory.

```
Please enter the Database Directory now!

Default: [/opt/n1ge6/default/spooldb] >>
Dumping bootstrapping information
Initializing spooling database

Hit <RETURN> to continue >>
```

- **If you do *not* want to use a use a Berkeley DB spooling server, enter `n`.**

```
Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >> n


Hit <RETURN> to continue >>
```

```
Berkeley Database spooling parameters
-------------------------------------

Please enter the Database Directory now, even if you want to spool locally
it is necessary to enter this Database Directory.

Default: [/opt/n1ge6/default/spool/spooldb] >>
```

Specify an alternate directory, or press <RETURN> to continue.

```
creating directory: /opt/n1ge6/default/spool/spooldb
Dumping bootstrapping information
Initializing spooling database

Hit <RETURN> to continue >>
```

- **If you want to specify classic spooling, then enter `classic`.**

```
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >> classic

Dumping bootstrapping information
Initializing spooling database

Hit <RETURN> to continue >>
```

14. Enter a group ID range

For more information, see "Group IDs" on page 19.

```
Grid Engine group id range
--------------------------

When jobs are started under the control of Grid Engine an additional group id
is set on platforms which do not support jobs. This is done to provide maximum
control for Grid Engine jobs.

This additional UNIX group id range must be unused group id's in your system.
Each job will be assigned a unique id during the time it is running.
Therefore you need to provide a range of id's which will be assigned
dynamically for jobs.

The range must be big enough to provide enough numbers for the maximum number
of Grid Engine jobs running at a single moment on a single host. E.g. a range
like >20000-20100< means, that Grid Engine will use the group ids from
20000-20100 and provides a range for 100 Grid Engine jobs at the same time
on a single host.

You can change at any time the group id range in your cluster configuration.

Please enter a range >> 20000-20100

Using >20000-20100< as gid range. Hit <RETURN> to continue >>
```

15. Verify the spooling directory for the execution daemon.

For information on spooling, see "Spool Directories Under the Root Directory" on page 17.

```
Grid Engine cluster configuration
---------------------------------

Please give the basic configuration parameters of your Grid Engine
installation:

   <execd_spool_dir>

The pathname of the spool directory of the execution hosts. User >sgeadmin<
must have the right to create this directory and to write into it.

Default: [/opt/n1ge6/default/spool] >>
```

16. **Enter the email address of the user who should receive problem reports.**

    In this example, the user who will receive problem report is `me@my.domain`.

```
Grid Engine cluster configuration (continued)
---------------------------------------------

<administator_mail>

The email address of the administrator to whom problem reports are sent.

It's is recommended to configure this parameter. You may use >none<
if you do not wish to receive administrator mail.

Please enter an email address in the form >user@foo.com<.

Default: [none] >> me@my.domain
```

17. **Verify the configuration parameters.**

```
The following parameters for the cluster configuration were configured:

   execd_spool_dir        /opt/n1ge6/default/spool
   administrator_mail     me@my.domain

Do you want to change the configuration parameters (y/n) [n] >> n

Creating local configuration
----------------------------
Creating >act_qmaster< file
Adding default complex attributes
Reading in complex attributes.
Adding default parallel environments (PE)
Reading in parallel environments:
        PE "make".
Adding SGE default usersets
Reading in usersets:
        Userset "deadlineusers".
        Userset "defaultdepartment".
Adding >sge_aliases< path aliases file
Adding >qtask< qtcsh sample default request file
Adding >sge_request< default submit options file
Creating >sgemaster< script
```

```
Creating >sgeexecd< script
Creating settings files for >.profile/.cshrc<

Hit <RETURN> to continue >>
```

18. Specify whether you want the daemons to start when the system is booted.

```
qmaster/scheduler startup script
--------------------------------

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

Installing startup script /etc/rc2.d/S95sgemaster

Hit <RETURN> to continue >>
...
```

19. Identify the hosts that you will later install as execution hosts.

```
Adding Grid Engine hosts
------------------------

Please now add the list of hosts, where you will later install your execution
daemons. These hosts will be also added as valid submit hosts.

Please enter a blank separated list of your execution hosts. You may
press <RETURN> if the line is getting too long. Once you are finished
simply press <RETURN> without entering a name.

You also may prepare a file with the hostnames of the machines where you plan
to install Grid Engine. This may be convenient if you are installing Grid
Engine on many hosts.

Do you want to use a file which contains the list of hosts (y/n) [n] >> n

Adding admin and submit hosts
-----------------------------

Please enter a blank seperated list of hosts.

Stop by entering <RETURN>. You may repeat this step until you are
entering an empty list. You will see messages from Grid Engine
when the hosts are added.

Host(s): host1 host2 host3 host4

host1 added to administrative host list
host1 added to submit host list
host2 added to administrative host list
host2 added to submit host list
host3 added to administrative host list
host3 added to submit host list
host4 added to administrative host list
host4 added to submit host list
```

```
Hit <RETURN> to continue >>

Creating the default <all.q> queue and <allhosts> hostgroup
-----------------------------------------------------------

root@vector added "@allhosts" to host group list
root@vector added "all.q" to cluster queue list

Hit <RETURN> to continue >>
```

20. Select a scheduler profile.

For information on how to determine which profile you should use, see "Scheduler Profiles" on page 21.

```
Scheduler Tuning
----------------

The details on the different options are described in the manual.

Configurations
--------------
1) Normal
          Fixed interval scheduling, report scheduling information,
          actual + assumed load

2) High
          Fixed interval scheduling, report limited scheduling information,
          actual load

3) Max
          Scheduling on demand, report no scheduling information,
          actual load

Enter the number of your prefered configuration and hit <RETURN>!
Default configuration is [1] >>
```

Once you answer this question, the installation process is complete. Several screens of information will be displayed before the script exits. The commands that are noted in those screens are also documented in this chapter.

21. Create the environment variables for use with the grid engine software.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  `% source` *sge-root*`/`*cell*`/common/settings.csh`

- **If you are using a Bourne shell or Korn shell, type the following command:**

  `$ .` *sge-root*`/`*cell*`/common/settings.sh`

See Also For details on how you can verify that the execution host has been set up correctly, see "How to Verify That the Daemons Are Running on the Master Host" on page 79.

▼ How to Install Execution Hosts

The execution host installation procedure creates the appropriate directory hierarchy required by `sge_execd`, and starts the `sge_execd` daemon on the execution host. This section describes how to install execution hosts interactively from the command line. You can automate the installation of execution of multiple hosts by using the procedure described in Chapter 3.

Before You Begin Before installing an execution host, you first need to have installed the master server as described in "How to Install the Master Host" on page 28.

Steps

1. **Log in to the execution host as `root`.**
2. **As you did for the master installation, either copy the installation files to a local installation directory *sge-root* or use a network installation directory.**
3. **Ensure that you have set the `$SGE_ROOT` environment variable by typing:**

   ```
   # echo $SGE_ROOT
   ```

   - **If the `$SGE_ROOT` environment variable is not set, set it now, by typing:**

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

4. **Change directory (`cd`) to the installation directory, *sge-root*.**
5. **Verify that the execution host has been declared on the administration host.**

   ```
   # qconf -sh
   ```

   - **If you do not see the name of this execution host in the output of the `qconf` command, you will need to declare it as an administration host.**
     a. **Start a new terminal session or window.**
     b. **In that window, log into the master host.**
     c. **Declare the execution host as an administration host, using the `qconf` command.**

        ```
        # qconf -ah quark
        quark added to administrative host list
        ```

     d. **Log back out of the master host, and continue with the installation of the execution host.**
6. **Run the `install_execd` command.**

If you are installing using the Certificate Security Protocol method described in Chapter 4, add the `-csp` option to the `install_execd` command.

```
% ./install_execd
```

This command starts the execution host installation procedure.

```
Welcome to the Grid Engine execution host installation
------------------------------------------------------

If you haven't installed the Grid Engine qmaster host yet, you must execute
this step (with >install_qmaster<) prior the execution host installation.

For a sucessfull installation you need a running Grid Engine qmaster. It is
also neccesary that this host is an administrative host.

You can verify your current list of administrative hosts with
the command:

   # qconf -sh

You can add an administrative host with the command:

   # qconf -ah <hostname>

The execution host installation will take approximately 5 minutes.

Hit <RETURN> to continue >>
```

7. **Verify the *sge-root* directory setting.**

   In the following example, the value of *sge-root* is `/opt/n1ge6`.

```
Checking $SGE_ROOT directory
----------------------------

The Grid Engine root directory is:

   $SGE_ROOT = /opt/n1ge6

If this directory is not correct (e.g. it may contain an automounter
prefix) enter the correct path to this directory or hit <RETURN>
to use default [/opt/n1ge6] >>

Your $SGE_ROOT directory: /opt/n1ge6

Hit <RETURN> to continue >>
```

8. **Enter the name of your cell(s).**

   The use of grid engine system cells is described in "Cells" on page 15.

```
Grid Engine cells
-----------------

Grid Engine supports multiple cells.
```

```
If you are not planning to run multiple Grid Engine clusters or if you don't
know yet what is a Grid Engine cell it is safe to keep the default cell name

   default

If you want to install multiple cells you can enter a cell name now.

The environment variable

   $SGE_CELL=<your_cell_name>

will be set for all further Grid Engine commands.

Enter cell name [default] >>
```

- **If you have decided to use cells, then enter the cell names now.**
- **If you have decided not to use cells, then press `<RETURN>` to continue.**

```
Using cell >default<.
Hit <RETURN> to continue >>
```

Press `<RETURN>` to continue.

9. **Press `<RETURN>` to continue.**

   The script verifies that the execution host has been declared as an administration host.

```
Checking hostname resolving
---------------------------

This hostname is known at qmaster as an administrative host.

Hit <RETURN> to continue >>
```

10. **Specify whether you want to use a local spool directory.**

    For information on spooling, see "Spool Directories Under the Root Directory" on page 17.

```
Local execd spool directory configuration
-----------------------------------------

During the qmaster installation you've already entered a global
execd spool directory. This is used, if no local spool directory is configured.

Now you can enter a local spool directory for this host.

Do you want to configure a local spool directory
for this host (y/n) [n] >>
```

- **If you do *not* want a local spool directory, answer `n`.**

```
Do you want to configure a local spool directory
for this host (y/n) [n] >> n
```

```
Creating local configuration
----------------------------
sgeadmin@host1 modified "host1" in configuration list

Local configuration for host >host1< created.

Hit <RETURN> to continue >>
```

- **If you do want a local spool directory, answer `y`.**

  In the following example, `/var/tmp/spool` is used as the local spool directory on `host1`. Choose any directory that meets the disk space requirements described in "Disk Space Requirements" on page 13.

```
Do you want to configure a local spool directory
for this host (y/n) [n] >> y

Please enter the local spool directory now! >> /var/tmp/spool
Using local execd spool directory [/var/tmp/spool]
Hit <RETURN> to continue >>

Creating local configuration
----------------------------
sgeadmin@host1 modified "host1" in configuration list

Local configuration for host >host1n< created.

Hit <RETURN> to continue >>
```

11. **Specify whether you want `execd` to start automatically at boot time.**

    You may not wish to install the startup script if you are installing a test cluster or you decide you'd rather start the daemon manually on reboot.

```
execd startup script
--------------------

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

Installing startup script /etc/rc2.d/S95sgeexecd

Hit <RETURN> to continue >>

Grid Engine execution daemon startup
------------------------------------

Starting execution daemon. Please wait ...
   starting sge_execd

Hit <RETURN> to continue >>
```

12. **Specify a queue for this host.**

```
Adding a queue for this host
----------------------------
```

```
We can now add a queue instance for this host:

   - it is added to the >allhosts< hostgroup
   - the queue provides 1 slot(s) for jobs in all queues
     referencing the >allhosts< hostgroup

You do not need to add this host now, but before running jobs on this host
it must be added to at least one queue.

Do you want to add a default queue instance for this host (y/n) [y] >>
```

Once you answer this question, the installation process is complete. Several screens of information will be displayed before the script exits. The commands that are noted in those screens are also documented in this chapter.

13. Create the environment variables for use with the grid engine software.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  `% source` *sge-root*`/`*cell*`/common/settings.csh`

- **If you are using a Bourne shell or Korn shell, type the following command:**

  `$ .` *sge-root*`/`*cell*`/common/settings.sh`

See Also For details on how you can verify that the execution host has been set up correctly, see "How to Verify That the Daemons Are Running on the Execution Hosts" on page 80.

Registering Administration Hosts

The master host is implicitly allowed to run administrative tasks and to submit, monitor, and delete jobs. The master host does not require any kind of additional installation as administration. By contrast, *pure* administration hosts do require registration.

Note – You can also install execution hosts by using the `QMON` graphical user interface. For information on how to complete this task using `QMON`, see "Configuring Administration Hosts With QMON" in *N1 Grid Engine 6 Administration Guide*.

From the master host, using the grid engine system administrative account (for example, the `sgeadmin` account), type the following command.

```
% qconf -ah admin_host_name[,...]
```

Registering Submit Hosts

Note – You can also install execution hosts by using the `QMON` graphical user interface. For information on how to complete this task using `QMON`, see "Configuring Submit Hosts With QMON" in *N1 Grid Engine 6 Administration Guide*.

From the master host, using the grid engine system administrative account (for example, the `sgeadmin` account), type the following command.

```
% qconf -as submit_host_name[,...]
```

Refer to "About Hosts and Daemons" in *N1 Grid Engine 6 Administration Guide* for more details and other means to configure the different host types.

▼ How to Install the Berkeley DB Spooling Server

The installation procedure installs the grid engine software necessary for Berkeley DB spooling.

Before You Begin

The grid engine software must be loaded onto a local file system. For details on how to extract the files, see "How to Load the Distribution Files On a Workstation" on page 22.

Steps

1. **Log in to the spooling server host as `root`.**
2. **Ensure that you have set the `$SGE_ROOT` environment variable by typing:**

   ```
   # echo $SGE_ROOT
   ```

   - **If the `$SGE_ROOT` environment variable is not set, set it now, by typing:**

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

3. **Change to the installation directory.**

   ```
   # cd $SGE_ROOT
   ```

4. **Type the `inst_sge` command with the `-db` option.**

   ```
   # sge-root/inst_sge -db
   ```

   This command starts the spooling server installation procedure. You are asked several questions. If you think something went wrong, you can quit the installation procedure and restart it at any time.

5. **Choose an administrative account owner.**

```
Choosing Grid Engine admin user account
---------------------------------------

You may install Grid Engine that all files are created with the user id of an
unprivileged user.

This will make it possible to install and run Grid Engine in directories
where user >root< has no permissions to create and write files and directories.

   - Grid Engine still has to be started by user >root<

   - this directory should be owned by the Grid Engine administrator

Do you want to install Grid Engine
under an user id other than >root< (y/n) [y] >> y

Choosing a Grid Engine admin user name
--------------------------------------

Please enter a valid user name >> sgeadmin
Installing Grid Engine as admin user >sgeadmin<

Hit <RETURN> to continue >>
```

6. **Verify the *sge-root* directory setting.**

   In the following example, the value of *sge-root* is `/opt/n1ge6`.

```
Checking $SGE_ROOT directory
----------------------------

The Grid Engine root directory is:

   $SGE_ROOT = /opt/n1ge6

If this directory is not correct (e.g. it may contain an automounter
prefix) enter the correct path to this directory or hit <RETURN>
to use default [/opt/n1ge6] >>

Your $SGE_ROOT directory: /opt/n1ge6

Hit <RETURN> to continue >>
```

7. **Enter the name of your cell.**

   The use of grid engine system cells is described in "Cells" on page 15.

```
Grid Engine cells
-----------------

Grid Engine supports multiple cells.

If you are not planning to run multiple Grid Engine clusters or if you don't
know yet what is a Grid Engine cell it is safe to keep the default cell name
```

```
   default

If you want to install multiple cells you can enter a cell name now.

The environment variable

   $SGE_CELL=<your_cell_name>

will be set for all further Grid Engine commands.

Enter cell name [default] >>
```

8. **Select Berkeley DB spooling.**

```
Setup spooling
--------------
Your SGE binaries are compiled to link the spooling libraries
during runtime (dynamically). So you can choose between Berkeley DB
spooling and Classic spooling method.
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >>
```

9. **Verify your host name.**

   In this example, the installation script is being run on `host2`.

```
Berkeley Database spooling parameters
-------------------------------------

You are going to install a RPC Client/Server machanism!
In this case, qmaster will
contact a RPC server running on a separate server machine.
If you want to use the SGE shadowd, you have to use the
RPC Client/Server mechanism.


Enter database server name or
hit <RETURN> to use default [host2] >>
```

10. **Enter the directory path of your spooling directory.**

    You may need to change this path if this directory is NFS mounted, or if you don't have write permissions to this directory.

```
Enter the database directory
or hit <RETURN> to use default [/opt/n1ge6/default//spooldb] >>

creating directory: /opt/n1ge6/default//spooldb
```

11. **Start the RPC server.**

```
Now we have to startup the rc script
>/opt/n1ge6/default/common/sgebdb<
on the RPC server machine

If you already have a configured Berkeley DB Spooling Server,
you have to restart the Database with the rc script now and continue with >NO<
```

```
Shall the installation script try to start the RPC server? (y/n) [y] >> y
Starting rpc server on host host2!
The Berkeley DB has been started with these parameters:

Spooling Server Name: host2
DB Spooling Directory: /opt/n1ge6/default//spooldb

Please remember these values, during Qmaster installation
you will be asked for! Hit <RETURN> to continue!
```

12. Specify whether you want Berkeley DB service to start automatically at boot time.

```
Berkeley DB startup script
--------------------------

We can install the startup script that
Grid Engine is started at machine boot (y/n) [y] >> y
```

Once you answer this question, the installation process is complete.

13. Create the environment variables for use with the grid engine software.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  `% source` *sge-root*`/`*cell*`/common/settings.csh`

- **If you are using a Bourne shell or Korn shell, type the following command:**

  `$ .` *sge-root*`/`*cell*`/common/settings.sh`

CHAPTER 3

Automating the Installation Process

This chapter describes how the software installation process for grid engine software can be automated. Currently, you may only perform this procedure as the administrative user. If you attempt to use this procedure running as the `root` user, the procedure will fail.

Using the `inst_sge` Utility and a Configuration Template

The `inst_sge` utility is a general-purpose script that can be used on any host to install locally as a master, administration, execution, submit or shadow master host. You can use this utility in interactive mode to supplant any of the commands that were described in Chapter 2. In addition, the `inst_sge` utility can be used in combination with a configuration file to automate system installation.

Note – When the `inst_sge` utility is used in combination with a configuration file to automate system installation, the user will receive no feedback from the system. However, a log file is created in the *sge-root*/*cell*/`spool` directory. If the installation script fails to complete, the log file can be found under the `/tmp` directory.

▼ How to Use the `inst_sge` Utility to Automate the Installation Process

Before You Begin

You need to complete the planning process as outlined in "Plan the Installation" on page 11.

In addition, you will need to be able to connect to each of the remote hosts using the `rsh` or `ssh` commands, without supplying a password. If this type of access is not allowed on your network, you will not be able to use this method of installation.

Steps

1. **Change the ownership of the *sge-root* directory to the grid engine system administrative user.**

   Currently, the owner of the *sge-root* directory must be the grid engine system administrative user. You must change the ownership of the *sge-root* directory to belong to your existing administrative user. If you attempt to use this procedure running as the `root` user, the procedure will fail.

   In the following example, your administrative user is `sgeadmin`, and *sge-root* is `/opt/n1ge6`.

   ```
   # chown -R sgeadmin /opt/n1ge6
   ```

2. **Create a copy of the configuration template, *sge-root*`/util/install_modules/inst_template.conf`.**

   ```
   # cd sge-root/util/install_modules
   # cp inst_template.conf  my_configuration.conf
   ```

3. **Edit your configuration template, using the values from the worksheet you completed in "Plan the Installation" on page 11.**

4. **Start the installation.**

   In this example, the local machine is installed as a master host and an execution host. In addition, the `-auto` option sets up any remote hosts, as specified in the configuration file.

   ```
   # cd sge-root
   # ./inst_sge -m -x -auto ./util/install_modules/my_configuration.conf
   ```

   A log file of the master installation is created in the *sge-root*`/default/spool/qmaster` directory. The file name is created using the format `install_`*hostname_date_time*`.log`.

Example 3–1 Example Configuration File

```
#-------------------------------------------------
# SGE default configuration file
#-------------------------------------------------

# Use always fully qualified pathnames, please

# SGE_ROOT Path, this is basic information
#(mandatory for qmaster and execd installation)
SGE_ROOT="/opt/n1ge6"

# SGE_QMASTER_PORT is used by qmaster for communication
# Please enter the port in this way: 1300
# Please do not this: 1300/tcp
```

```
#(mandatory for qmaster installation)
SGE_QMASTER_PORT="536"

# SGE_EXECD_PORT is used by execd for communication
# Please enter the port in this way: 1300
# Please do not this: 1300/tcp
#(mandatory for qmaster installation)
SGE_EXECD_PORT="537"

# CELL_NAME, will be a dir in SGE_ROOT, contains the common dir
# Please enter only the name of the cell. No path, please
#(mandatory for qmaster and execd installation)
CELL_NAME="default"

# The dir, where qmaster spools this parts, which are not spooled by DB
#(mandatory for qmaster installation)
QMASTER_SPOOL_DIR="/opt/n1ge6/default/spool/qmaster"

# The dir, where the execd spools (active jobs)
# This entry is needed, even if your are going to use
# berkeley db spooling. Only cluster configuration and jobs will
# be spooled in the database. The execution daemon still needs a spool
# directory
#(mandatory for qmaster installation)
EXECD_SPOOL_DIR="/opt/n1ge6/default/spool"

# For monitoring and accounting of jobs, every job will get
# unique GID. So you have to enter a free GID Range, which
# is assigned to each job running on a machine.
# If you want to run 100 Jobs at the same time on one host you
# have to enter a GID-Range like that: 16000-16100
#(mandatory for qmaster installation)
GID_RANGE="20000-20100"

# If SGE is compiled with -spool-dynamic, you have to enter here, which
# spooling method should be used. (classic or berkeleydb)
#(mandatory for qmaster installation)
SPOOLING_METHOD="berkeleydb"

# Name of the Server, where the Spooling DB is running on
# if spooling methode is berkeleydb, it must be "none", when
# using no spooling server and it must containe the servername
# if a server should be used. In case of "classic" spooling,
# can be left out
DB_SPOOLING_SERVER="none"

# The dir, where the DB spools
# If berkeley db spooling is used, it must contain the path to
# the spooling db. Please enter the full path. (eg. /tmp/data/spooldb)
# Remember, this directory must be local on the qmaster host or on the
# Berkeley DB Server host. No NSF mount, please
DB_SPOOLING_DIR="/opt/n1ge6/default/spooldb"

# A List of Host which should become admin hosts
# If you do not enter any host here, you have to add all of your hosts
```

```
# by hand, after the installation. The autoinstallation works without
# any entry
ADMIN_HOST_LIST="host1"

# A List of Host which should become submit hosts
# If you do not enter any host here, you have to add all of your hosts
# by hand, after the installation. The autoinstallation works without
# any entry
SUBMIT_HOST_LIST="host1"

# A List of Host which should become exec hosts
# If you do not enter any host here, you have to add all of your hosts
# by hand, after the installation. The autoinstallation works without
# any entry
# (mandatory for execution host installation)
EXEC_HOST_LIST="host1"

# The dir, where the execd spools (local configuration)
# If you want configure your execution daemons to spool in
# a local directory, you have to enter this directory here.
# If you do not want to configure a local execution host spool directory
# please leave this empty
EXECD_SPOOL_DIR_LOCAL=""

# If true, the domainnames will be ignored, during the hostname resolving
# if false, the fully qualified domain name will be used for name resolving
HOSTNAME_RESOLVING="true"

# Shell, which should be used for remote installation (rsh/ssh)
# This is only supported, if your hosts and rshd/sshd is configured,
# not to ask for a password, or promting any message.
SHELL_NAME="rsh"

# Enter your default domain, if you are using /etc/hosts or NIS configuration
DEFAULT_DOMAIN="none"

# If a job stops, fails, finnish, you can send a mail to this adress
ADMIN_MAIL="my.name@sun.com"

# If true, the rc scripts (sgemaster, sgeexecd, sgebdb) will be added,
# to start automatically during boottime
ADD_TO_RC="true"

#If this is "true" the file permissions of executables will be set to 755
#and of ordenary file to 644.
SET_FILE_PERMS="true"

# This option is not implemented, yet.
# When a exechost should be uninstalled, the running jobs will be rescheduled
RESCHEDULE_JOBS="wait"

# Enter a one of the three distributed scheduler tuning configuration sets
# (1=normal, 2=high, 3=max)
SCHEDD_CONF="1"
```

```
# The name of the shadow host. This host must have read/write permission
# to the qmaster spool directory
# If you want to setup a shadow host, you must enter the servername
# (mandatory for shadowhost installation)
SHADOW_HOST="hostname"

# Remove this execution hosts in automatic mode
# (mandatory for unistallation of executions hosts)
EXEC_HOST_LIST_RM="host2 host3 host4"
```

Troubleshooting If you do not want your execution hosts to spool locally, be sure to set `EXECD_SPOOL_DIR_LOCAL=""`, with no space between the double quotes ("").

CHAPTER 4

Installing the Increased Security Features

Use the instructions in this section to set up your system more securely. These instructions will help you set up your system with *Certificate Security Protocol* (CSP)-based encryption.

Why Install the Increased Security Features?

Instead of transferring messages in clear text, the messages in this secure system are encrypted with a secret key. The secret key is exchanged using a public/private key protocol. Users present their certificates through the grid engine system to prove identity. Users receive the certificate to ensure that they are communicating with the correct systems. After this initial announcement phase, communication continues transparently in encrypted form. The session is valid only for a certain period, after which the session must be re-announced.

Additional Setup Required

The steps required to set up the Certificate Security Protocol enhanced version of the grid engine system are similar to the standard setup. You generally follow the instructions in “Plan the Installation” on page 11, “How to Load the Distribution Files On a Workstation” on page 22, “How to Install the Master Host” on page 28, “How to Install Execution Hosts” on page 41, and “Registering Administration Hosts” on page 45.

However, the following additional tasks are required:

- Generating the Certificate Authority (CA) system keys and certificates on the master host by calling the installation script with the `-csp` flag
- Distributing the system keys and certificates to the execution and submit hosts using a secure method such as `ssh`
- Generating user keys and certificates automatically, after master installation
- Adding new users

▼ How to Install a CSP-Secured System

Install the grid engine software as outlined in "Performing an Installation" on page 28, with the following exception: use the additional flag `-csp` when invoking the various installation scripts.

Steps

1. **To install a CSP-secured system, change the master host installation procedure by typing the following command and responding to the prompts from the installation script.**

   ```
   # ./install_qmaster -csp
   ```

2. **To generate the CSP certificates and keys, you must supply the following information:**

   - Two-letter country code, for example, US for the United States
   - State
   - Location, such as a city
   - Organization
   - Organizational unit
   - CA email address

   As the installation proceeds, the Certificate Authority is created. A CA specific to the grid engine system is created on the master host. The directories that contain information relevant to security are as follows:

   - The publicly accessible CA and daemon certificate are stored in *sge-root*/*cell*/`common/sgeCA`.
   - The corresponding private keys are stored in `/var/sgeCA/{`*sge_service* | `port`*SGE_QMASTER_PORT*`}/`*cell*`/private`.
   - User keys and certificates are stored in `/var/sgeCA/{`*sge_service* | `port`*SGE_QMASTER_PORT*`}/`*cell*`/userkeys/$USER`.

3. **The script prompts you for site information.**

4. **Confirm whether the information you supplied is correct.**

5. **After the security-related setup of the master host `sge_qmaster` is finished, the script prompts you to continue with the rest of the installation procedure, as in the following example:**

```
SGE startup script
--------------------

Your system wide SGEEE startup script is installed as:

     "/scratch2/eddy/sge_sec/default/common/sgemaster"

Hit Return to continue >>
```

6. **Determine whether the shared file system is secure enough to hold the CSP security information.**

   - **If you think that the shared file system *is* secure enough, continue with the basic installation procedure as outlined in "How to Install Execution Hosts" on page 41.**

     Note – Remember to add the `-csp` flag when you call the `./install_execd` script for the execution host installation.

     If the `root` user does not have write permissions in the *sge-root* directory on *all* of the machines where grid engine software will be installed, you will be asked whether to install the software as the user to whom the directory belongs. If you answer `yes`, you must install the security-related files into that user's `$HOME/.sge` directory, as shown in the following example. In the following example, *sgeadmin* is the name of the user who owns the installation directory.

     ```
     % su - sgeadmin
     % source sge-root/default/common/settings.csh
     % sge-root/util/sgeCA/sge_ca -copy
     % logout
     ```

     After completing all remaining installation steps, refer to the instructions in "How to Generate Certificates and Private Keys for Users" on page 61.

   - **If you think that the shared file system is *not* secure enough to hold the CSP security information in a place that can be accessed by the execution daemons, you must transfer the directory containing the daemon's private key and the random file to the *execution* host.**

     a. **As `root` on the *master* host, type the following commands to prepare to copy the private keys to the machines you set up as execution hosts.**

        ```
        # umask 077
        # cd /
        # tar cvpf /var/sgeCA/port536.tar /var/sgeCA/port536/default
        ```

b. As `root` on each *execution* host, use the following commands to securely copy the files.

```
# umask 077
# cd /
# scp masterhost:/var/sgeCA/port536.tar .
# umask 022
# tar xvpf /port536.tar
# rm /port536.tar
```

c. Type the following command to verify the file permissions.

```
# ls -lR /var/sgeCA/port536/
```

The output should look like the following example:

```
/var/sgeCA/port536/:
total 2
drwxr-xr-x   4 eddy     other        512 Mar  6 10:52 default
/var/sgeCA/port536/default:
total 4
drwx------   2 eddy     staff        512 Mar  6 10:53 private
drwxr-xr-x   4 eddy     staff        512 Mar  6 10:54 userkeys
/var/sgeCA/port536/default/private:
total 8
-rw-------   1 eddy     staff        887 Mar  6 10:53 cakey.pem
-rw-------   1 eddy     staff        887 Mar  6 10:53 key.pem
-rw-------   1 eddy     staff       1024 Mar  6 10:54 rand.seed
-rw-------   1 eddy     staff        761 Mar  6 10:53 req.pem
/var/sgeCA/port536/default/userkeys:
total 4
dr-x------   2 eddy     staff        512 Mar  6 10:54 eddy
dr-x------   2 root     staff        512 Mar  6 10:54 root
/var/sgeCA/port536/default/userkeys/eddy:
total 16
-r--------   1 eddy     staff       3811 Mar  6 10:54 cert.pem
-r--------   1 eddy     staff        887 Mar  6 10:54 key.pem
-r--------   1 eddy     staff       2048 Mar  6 10:54 rand.seed
-r--------   1 eddy     staff        769 Mar  6 10:54 req.pem
/var/sgeCA/port536/default/userkeys/root:
total 16
-r--------   1 root     staff       3805 Mar  6 10:54 cert.pem
-r--------   1 root     staff        887 Mar  6 10:54 key.pem
-r--------   1 root     staff       2048 Mar  6 10:53 rand.seed
-r--------   1 root     staff        769 Mar  6 10:54 req.pem
```

d. Install the security-related files into the admin user's `$HOME/.sge` directory.

If the `root` user does not have write permissions in the *sge-root* directory on *all* of the machines where grid engine software will be installed, you will be asked whether to install the software as the user to whom the directory belongs. If you answer `yes`, you must install the security-related files into that user's `$HOME/.sge` directory, as shown in the following example. In the following example, *sgeadmin* is the name of the user who owns the installation directory.

```
% su - sgeadmin
% source sge-root/default/common/settings.csh
% sge-root/util/sgeCA/sge_ca -copy
% logout
```

e. Continue with the grid engine software installation on the execution host by typing the following commands.

Note – You can also use the master host for executing jobs. To do so, you must carry out the execution host installation for the master machine. However, if you use a very slow machine as master host, or if your cluster is significantly large, do not use the master host as an execution host.

```
# cd sge-root
# ./install_execd -csp
```

f. Respond to the prompts from the installation script.

The execution host installation procedure creates the appropriate directory hierarchy required by `sge_execd`, and starts the `sge_execd` daemon on the execution host.

g. Create the environment variables for use with the grid engine software.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  ```
  % source sge_root/cell/common/settings.csh
  ```

- **If you are using a Bourne shell or Korn shell, type the following command:**

  ```
  $ . sge_root/cell/common/settings.sh
  ```

7. **Proceed to the next task,"How to Generate Certificates and Private Keys for Users" on page 61.**

▼ How to Generate Certificates and Private Keys for Users

To use the CSP-secured system, the user must have access to a user-specific certificate and private key. The most convenient method of gaining access is to create a text file identifying the users.

Steps

1. **On the master host, create and save a text file that identifies users.**

   Use the format of the file `myusers.txt` shown in the following example. The fields of the file are *UNIX_username:Gecos_field:email_address*.

   ```
   eddy:Eddy Smith:eddy@my.org
   sarah:Sarah Miller:sarah@my.org
   leo:Leo Lion:leo@my.org
   ```

2. **As `root` on the master host, type the following command:**

   ```
   # sge-root/util/sgeCA/sge_ca -usercert myusers.txt
   ```

3. **Confirm by typing the following command:**

   ```
   # ls -l /var/sgeCA/port536/default/userkeys
   ```

   This directory listing produces output similar to the following example.

   ```
   dr-x------  2 eddy  staff       512 Mar  5 16:13 eddy
   dr-x------  2 sarah staff        512 Mar  5 16:13 sarah
   dr-x------  2 leo   staff        512 Mar 5 16:13 leo
   ```

4. **Tell each user listed in the file (`myusers.txt` in the example) to install the security-related files in their `$HOME/.sge` directories by typing the following commands.**

   ```
   % source sge-root/default/common/settings.csh
   % sge-root/util/sgeCA/sge_ca -copy
   ```

   Users should see the following confirmation (user `eddy` in the example).

   ```
   Certificate and private key for user
   eddy have been installed
   ```

   For every grid engine software installation, a subdirectory for the corresponding `SGE_QMASTER_PORT` number is installed. The following example, based on the `myusers.txt` file, results from issuing the command preceding the output.

   ```
   % ls -lR $HOME/.sge

   /home/eddy/.sge:
   total 2
   drwxr-xr-x  3 eddy staff        512 Mar  5 16:20 port536

   /home/eddy/.sge/port536:
   total 2
   drwxr-xr-x  4 eddy staff        512 Mar  5 16:20 default

   /home/eddy/.sge/port536/default:
   total 4
   drwxr-xr-x  2 eddy staff        512 Mar  5 16:20 certs
   drwx------  2 eddy staff        512 Mar 5 16:20 private

   /home/eddy/.sge/port536/default/certs:
   total 8
   -r--r--r--  1 eddy staff       3859 Mar  5 16:20 cert.pem
   ```

```
/home/eddy/.sge/port536/default/private:
total 6
-r--------  1 eddy staff        887 Mar  5 16:20 key.pem
-r--------  1 eddy staff       2048 Mar 5 16:20 rand.seed
```

Checking Certificates

The following sections provide examples of commands related to certificates, where *arch* is your system architecture, as in `sol-sparc64`. Depending on what you want to do, type one or more of the following commands.

Display a Certificate

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -in ~/.sge/port536/default/certs/cert.pem -text
```

Check Issuer

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -issuer -in ~/.sge/port536/default/certs/cert.pem -noout
```

Check Subject

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -subject -in ~/.sge/port536/default/certs/cert.pem
-noout
```

Show Email of Certificate

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -email -in ~/.sge/default/port536/certs/cert.pem -noout
```

Show Validity

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -dates -in ~/.sge/default/port536/certs/cert.pem -noout
```

Show Fingerprint

Type the following as one string (the command is too long to fit on one line in this guide), with a space between the `-in` and the `~/.sge` components.

```
% sge-root/utilbin/arch/openssl
x509 -fingerprint -in ~/.sge/port536/default/certs/cert.pem -noout
```

CHAPTER 5

Upgrading from a Previous Release of Grid Engine Software

This chapter describes the steps necessary to upgrade your existing software to this release. You can upgrade from either Sun™ ONE Grid Engine 5.3 or Sun ONE Grid Engine, Enterprise Edition, 5.3.

About Upgrading the Software

The upgrade procedure is non destructive. The upgrade procedure installs the N1 Grid Engine 6 software on the master host by using the cluster configuration information from the older version of the software. The older version of the software will not be removed or modified in any way.

▼ How to Upgrade the Software

Before You Begin

Please review "Plan the Installation" on page 11 for the information that you will need during the upgrade process. If you have decided to use an administrative user, as described in "User Names" on page 15, you should create that user now. This procedure assumes that you have already extracted the grid engine software, as described in "Loading the Distribution Files On a Workstation" on page 22.

Note – While you can run N1 Grid Engine 6 software concurrently with your older version of grid engine software, it is better to run the upgrade procedure when there are no running jobs.

Steps

1. **Log in to the master host as `root`.**
2. **Load the distribution files.**

For details, see "Loading the Distribution Files On a Workstation" on page 22.

3. **Ensure that you have set the `$SGE_ROOT` environment variable by typing:**

   ```
   # echo $SGE_ROOT
   ```

   - **If the `$SGE_ROOT` environment variable is not set, set it now, by typing:**

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

4. **Change to the installation directory, *sge-root*.**

   - **If the directory where the installation files reside is visible from the master host, change directories (`cd`) to the installation directory *sge-root*, and then proceed to Step 4.**
   - **If the directory is not visible and cannot be made visible, do the following:**

     a. **Create a local installation directory, *sge-root*, on the master host.**

     b. **Copy the installation files to the local installation directory *sge-root* across the network (for example, by using `ftp` or `rcp`).**

     c. **Change directories (`cd`) to the local *sge-root* directory.**

5. **Run the upgrade command on the master host, and respond to the prompts.**

   This command starts the master host installation procedure. You are asked several questions, and you might be required to run some administrative actions.

   The syntax of the upgrade command is:

   `inst_sge -upd` *5.3-sge-root-directory 5.3-cell-name*

   In the following example, the 5.3 *sge-root* directory is `/sge/gridware` and the cell name is `default`.

```
# ./inst_sge -upd /sge/gridware default
Welcome to the Grid Engine Upgrade
----------------------------------

Before you continue with the installation please read these hints:

   - Your terminal window should have a size of at least
     80x24 characters

   - The INTR character is often bound to the key Ctrl-C.
     The term >Ctrl-C< is used during the upgrade if you
     have the possibility to abort the upgrade

The upgrade procedure will take approximately 5-10 minutes.
After this upgrade you will get a running qmaster and schedd with
the configuration of your old installation. If the upgrade was
successfully completed it is necessary to install your execution hosts
with the install_execd script.

Hit <RETURN> to continue >>
```

6. **Choose an administrative account owner.**

   In the following example, the value of *sge-root* is `/opt/n1ge6`, and the administrative user is `sgeadmin`.

```
Grid Engine admin user account
------------------------------

The current directory

   /opt/n1ge6

is owned by user

   sgeadmin

If user >root< does not have write permissions in this directory on *all*
of the machines where Grid Engine will be installed (NFS partitions not
exported for user >root< with read/write permissions) it is recommended to
install Grid Engine that all spool files will be created under the user id
of user >sgeadmin<.

IMPORTANT NOTE: The daemons still have to be started by user >root<.

Do you want to install Grid Engine as admin user >sgeadmin< (y/n) [y] >>
```

7. **Verify the *sge-root* directory setting.**

   In the following example, the value of *sge-root* is `/opt/n1ge6`.

```
Checking $SGE_ROOT directory
----------------------------

The Grid Engine root directory is:

   $SGE_ROOT = /opt/n1ge6

If this directory is not correct (e.g. it may contain an automounter
prefix) enter the correct path to this directory or hit <RETURN>
to use default [/opt/n1ge6] >>
```

8. **Set up the TCP/IP services for the grid engine software.**

   a. **You will be notified if the TCP/IP services have not been configured.**

```
Grid Engine TCP/IP service >sge_qmaster<
----------------------------------------

There is no service >sge_qmaster< available in your >/etc/services< file
or in your NIS/NIS+ database.

You may add this service now to your services database or choose a port number.
It is recommended to add the service now. If you are using NIS/NIS+ you should
add the service at your NIS/NIS+ server and not to the local >/etc/services<
file.

Please add an entry in the form
```

```
   sge_qmaster <port_number>/tcp

to your services database and make sure to use an unused port number.

Please add the service now or press <RETURN> to go to entering a port number >>
```

b. Start a new terminal session or window to add the information `/etc/services` file or your NIS maps.

c. Add the correct ports to the `/etc/services` file or your NIS services map, as described in "Network Services" on page 16.

The following example shows how you might edit your `/etc/services` file.

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

Note – In this example, the entries for both `sge_qmaster` and `sge_execd` are added to `/etc/services`. Subsequent steps in this example assume that both entries have been made.

Save your changes.

d. Return to the window where the installation script is running.

```
Please add the service now or press <RETURN> to go to entering a port number >>
```

Press `<RETURN>`. You will see the following output:

```
sge_qmaster 536

Service >sge_qmaster< is now available.

Hit <RETURN> to continue >>

Grid Engine TCP/IP service >sge_execd<
-------------------------------------

Using the service

   sge_execd

for communication with Grid Engine.

Hit <RETURN> to continue >>
```

9. Enter the name of your cell.

The use of grid engine system cells is described in "Cells" on page 15.

```
Grid Engine cells
-----------------
```

```
Grid Engine supports multiple cells.

If you are not planning to run multiple Grid Engine clusters or if you don't
know yet what is a Grid Engine cell it is safe to keep the default cell name

   default

If you want to install multiple cells you can enter a cell name now.

The environment variable

   $SGE_CELL=<your_cell_name>

will be set for all further Grid Engine commands.

Enter cell name [default] >>
```

- **If you have decided to use cells, then enter the cell name now.**
- **If you have decided not to use cells, then press <RETURN> to continue.**

```
Using cell >default<.
Hit <RETURN> to continue >>
```

Press <RETURN> to continue.

10. **Specify a spool directory.**

    For guidelines on disk space requirements for the spool directory, see "Disk Space Requirements" on page 13. For information on where spool directory is installed, see "Spool Directories Under the Root Directory" on page 17.

```
Grid Engine qmaster spool directory
-----------------------------------

The qmaster spool directory is the place where the qmaster daemon stores
the configuration and the state of the queuing system.

The admin user >sgeadmin< must have read/write access
to the qmaster spool directory.

If you will install shadow master hosts or if you want to be able to start
the qmaster daemon on other hosts (see the corresponding sectionin the
Grid Engine Installation and Administration Manual for details) the account
on the shadow master hosts also needs read/write access to this directory.

The following directory

[/opt/n1ge6/default/spool/qmaster]

will be used as qmaster spool directory by default!

Do you want to select another qmaster spool directory (y/n) [n] >>
```

- **If you want to accept the default spool directory, press <RETURN> to continue.**

- **If you do *not* want to accept the default spool directory, then answer `y`.**

  In the following example the `/my/spool` directory is specified as the master host spool directory.

```
Do you want to select another qmaster spool directory (y/n) [n] >> y

Please enter a qmaster spool directory now! >>/my/spool
```

11. Set the correct file permissions.

```
Verifying and setting file permissions
--------------------------------------

Did you install this version with >pkgadd< or did you already
verify and set the file permissions of your distribution (y/n) [y] >> n

Verifying and setting file permissions
--------------------------------------

We may now verify and set the file permissions of your Grid Engine
distribution.

This may be useful since due to unpacking and copying of your distribution
your files may be unaccessible to other users.

We will set the permissions of directories and binaries to

   755 - that means executable are accessible for the world

and for ordinary files to

   644 - that means readable for the world

Do you want to verify and set your file permissions (y/n) [y] >> y

Verifying and setting file permissions and owner in >3rd_party<
Verifying and setting file permissions and owner in >bin<
Verifying and setting file permissions and owner in >ckpt<
Verifying and setting file permissions and owner in >examples<
Verifying and setting file permissions and owner in >install_execd<
Verifying and setting file permissions and owner in >install_qmaster<
Verifying and setting file permissions and owner in >mpi<
Verifying and setting file permissions and owner in >pvm<
Verifying and setting file permissions and owner in >qmon<
Verifying and setting file permissions and owner in >util<
Verifying and setting file permissions and owner in >utilbin<
Verifying and setting file permissions and owner in >catman<
Verifying and setting file permissions and owner in >doc<
Verifying and setting file permissions and owner in >man<
Verifying and setting file permissions and owner in >inst_sge<
Verifying and setting file permissions and owner in >bin<
Verifying and setting file permissions and owner in >lib<
Verifying and setting file permissions and owner in >utilbin<

Your file permissions were set
```

```
Hit <RETURN> to continue >>
```

12. **Specify whether all of your grid engine system hosts are located in a single DNS domain.**

```
Select default Grid Engine hostname resolving method
----------------------------------------------------

Are all hosts of your cluster in one DNS domain? If this is
the case the hostnames

   >hostA< and >hostA.foo.com<

would be treated as eqal, because the DNS domain name >foo.com<
is ignored when comparing hostnames.

Are all hosts of your cluster in a single DNS domain (y/n) [y] >>
```

- **If all of your grid engine system hosts are located in a single DNS domain, then answer `y`.**

```
Are all hosts of your cluster in a single DNS domain (y/n) [y] >> y

Ignoring domainname when comparing hostnames.

Hit <RETURN> to continue >>
```

- **If all of your grid engine system hosts are *not* located in a single DNS domain, then answer `n`.**

```
Are all hosts of your cluster in a single DNS domain (y/n) [y] >> n

The domainname is not ignored when comparing hostnames.

Hit <RETURN> to continue >>

Default domain for hostnames
----------------------------

Sometimes the primary hostname of machines returns the short hostname
without a domain suffix like >foo.com<.

This can cause problems with getting load values of your execution hosts.
If you are using DNS or you are using domains in your >/etc/hosts< file or
your NIS configuration it is usually safe to define a default domain
because it is only used if your execution hosts return the short hostname
as their primary name.

If your execution hosts reside in more than one domain, the default domain
parameter must be set on all execution hosts individually.

Do you want to configure a default domain (y/n) [y] >>
```

Press `<RETURN>` to continue.

a. If you want to specify a default domain, then answer `y`.

In the following example, `sun.com` is specified as the default domain.

```
Do you want to configure a default domain (y/n) [y] >> y


Please enter your default domain >> sun.com

Using >sun.com< as default domain. Hit <RETURN> to continue >>
```

b. If you do *not* want to specify a default domain, then answer `n`.

In the following example, `sun.com` is specified as the default domain.

```
Do you want to configure a default domain (y/n) [y] >> n
```

13. Press `<RETURN>` to continue.

```
Making directories
------------------

creating directory: default/common
creating directory: /opt/n1ge6/default/spool/qmaster
creating directory: /opt/n1ge6/default/spool/qmaster/job_scripts
Hit <RETURN> to continue >>
```

14. Specify whether you want to use classic spooling or Berkeley DB.

For more information on how to determine the type of spooling mechanism you want, please see "Database Server and Spooling Host" on page 18.

```
Setup spooling
--------------
Your SGE binaries are compiled to link the spooling libraries
during runtime (dynamically). So you can choose between Berkeley DB
spooling and Classic spooling method.
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >>
```

- **If you want to specify Berkeley DB spooling, press `<RETURN>` to continue.**

```
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >>

The Berkeley DB spooling method provides two configurations!

1) Local spooling:
The Berkeley DB spools into a local directory on this host (qmaster host)
This setup is faster, but you can't setup a shadow master host

2) Berkeley DB Spooling Server:
If you want to setup a shadow master host, you need to use
Berkeley DB Spooling Server!
In this case you have to choose a host with a configured RPC service.
The qmaster host connects via RPC to the Berkeley DB. This setup is more
failsafe, but results in a clear potential security hole. RPC communication
(as used by Berkeley DB) can be easily compromised. Please only use this
alternative if your site is secure or if you are not concerned about
security. Check the installation guide for further advice on how to achieve
```

```
failsafety without compromising security.

Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >>
```

- **If you want to use a use a Berkeley DB spooling server, enter `y`.**

  ```
  Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >> y

  Berkeley DB Setup

  -----------------
  Please, log in to your Berkeley DB spooling host and execute "inst_sge -db"
  Please do not continue, before the Berkeley DB installation with
  "inst_sge -db" is completed, continue with <RETURN>
  ```

  Note – *Do not* press `<RETURN>` until you have completed the Berkeley DB installation on the spooling server.

  ---

  a. **Start a new terminal session or window.**

  b. **Log in to the spooling server.**

  c. **Install the software, as described in "How to Install the Berkeley DB Spooling Server" on page 46.**

  d. **After you have installed the software on the spooling server, return to the master installation window, and press `<RETURN>` to continue.**

  e. **Enter the name of the spooling server.**

     In the following example, `vector` is the host name of the spooling server.

     ```
     Berkeley Database spooling parameters
     -------------------------------------

     Please enter the name of your Berkeley DB Spooling Server! >> vector
     ```

  f. **Enter the name of the spooling directory.**

     In the following example, `/opt/n1ge6/default/spooldb` is the spooling directory.

     ```
     Please enter the Database Directory now!

     Default: [/opt/n1ge6/default/spooldb] >>
     Dumping bootstrapping information
     Initializing spooling database

     Hit <RETURN> to continue >>
     ```

- **If you do *not* want to use a use a Berkeley DB spooling server, enter `n`.**

```
Do you want to use a Berkeley DB Spooling Server? (y/n) [n] >> n


Hit <RETURN> to continue >>

Berkeley Database spooling parameters
-------------------------------------

Please enter the Database Directory now, even if you want to spool locally
it is necessary to enter this Database Directory.

Default: [/opt/n1ge6/default/spool/spooldb] >>
```

Specify an alternate directory, or press <RETURN> to continue.

```
creating directory: /opt/n1ge6/default/spool/spooldb
Dumping bootstrapping information
Initializing spooling database

Hit <RETURN> to continue >>
```

- **If you want to specify classic spooling, then enter `classic`.**

```
Please choose a spooling method (berkeleydb|classic) [berkeleydb] >> classic

Dumping bootstrapping information
Initializing spooling database

Hit <RETURN> to continue >>
```

15. Enter a group ID range

For more information, see "Group IDs" on page 19.

```
Grid Engine group id range
--------------------------

When jobs are started under the control of Grid Engine an additional group id
is set on platforms which do not support jobs. This is done to provide maximum
control for Grid Engine jobs.

This additional UNIX group id range must be unused group id's in your system.
Each job will be assigned a unique id during the time it is running.
Therefore you need to provide a range of id's which will be assigned
dynamically for jobs.

The range must be big enough to provide enough numbers for the maximum number
of Grid Engine jobs running at a single moment on a single host. E.g. a range
like >20000-20100< means, that Grid Engine will use the group ids from
20000-20100 and provides a range for 100 Grid Engine jobs at the same time
on a single host.

You can change at any time the group id range in your cluster configuration.

Please enter a range >> 20000-20100

Using >20000-20100< as gid range. Hit <RETURN> to continue >>
```

16. **Verify the spooling directory for the execution daemon.**

    For information on spooling, see "Spool Directories Under the Root Directory" on page 17.

```
Grid Engine cluster configuration
---------------------------------

Please give the basic configuration parameters of your Grid Engine
installation:

   <execd_spool_dir>

The pathname of the spool directory of the execution hosts. User >sgeadmin<
must have the right to create this directory and to write into it.

Default: [/opt/n1ge6/default/spool] >>
```

17. **Enter the email address of the user who should receive problem reports.**

    In this example, the user who will receive problem report is `me@my.domain`.

```
Grid Engine cluster configuration (continued)
---------------------------------------------

<administator_mail>

The email address of the administrator to whom problem reports are sent.

It's is recommended to configure this parameter. You may use >none<
if you do not wish to receive administrator mail.

Please enter an email address in the form >user@foo.com<.

Default: [none] >> me@my.domain
```

    Once you answer this question, the installation process is complete. Several screens of information will be displayed before the script exits. The commands that are noted in those screens are also documented in this chapter.

    The upgrade process uses your existing configuration to customize the installation. You will see output similar to the following:

```
Creating >act_qmaster< file
Creating >sgemaster< script
Creating >sgeexecd< script
creating directory: /tmp/centry
Reading in complex attributes.
Reading in administrative hosts.
Reading in execution hosts.
Reading in submit hosts.
Reading in users:
    User "as114086".
    User "md121042".
Reading in usersets:
    Userset "defaultdepartment".
    Userset "deadlineusers".
```

```
    Userset "admin".
    Userset "bchem1".
    Userset "bchem2".
    Userset "bchem3".
    Userset "bchem4".
    Userset "damtp7".
    Userset "damtp8".
    Userset "damtp9".
    Userset "econ1".
    Userset "staff".
Reading in calendars:
    Calendar "always_disabled".
    Calendar "always_suspend".
    Calendar "test".
Reading in projects:
    Project "ap1".
    Project "ap2".
    Project "high".
    Project "low".
    Project "p1".
    Project "p2".
    Project "staff".
Reading in parallel environments:
    PE "bench_tight".
    PE "make".
Creating settings files for >.profile/.cshrc<
```

Caution – Do not rename any of the binaries of the distribution. If you use any scripts or tools in your cluster that monitor the daemons, make sure to check for the new names.

18. Create the environment variables for use with the grid engine software.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  `% source` *sge-root*`/`*cell*`/common/settings.csh`

- **If you are using a Bourne shell or Korn shell, type the following command:**

  `$ .` *sge-root*`/`*cell*`/common/settings.sh`

19. Install or upgrade the execution hosts.

There are two ways that you can install the N1 Grid Engine 6 software on your execution hosts: installation or upgrade. If you *install* the execution hosts, the local spool directory configuration, and some `execd` parameters will be overwritten. If

you *upgrade* the execution hosts, those files will remain untouched.

- **Upgrade the software on the execution host.**
  You need to log into each execution host, and run the following command:

  ```
  # sge-root/inst_sge -x -upd
  ```

- **Install the software on the execution host.**

  a. **If you only have a few execution hosts, you can install them interactively.**
     You need to log into each execution host, and run the following command:

     ```
     # sge-root/inst_sge -x
     ```

     Complete instructions for installing execution hosts interactively are in "How to Install Execution Hosts" on page 41.

  b. **If you have a large number of execution hosts, you should consider installing them non-interactively.**
     Instructions for installing execution hosts in an automated way are in "Using the `inst_sge` Utility and a Configuration Template" on page 51.

20. **If you have configured load sensors on your execution hosts, you will need to copy these load sensors to the new directory location.**

21. **Check your complexes.**
    Both the structure of complexes and the rules for configuring complexes have changed. You can use `qconf -sc` to list your complexes. Review the log file that was generated during the master host upgrade, `update.`*pid*. the `update.`*pid* file will be placed in the master host spool directory, which is *sge-root*`/`*cell*`/spool/`by default.

    sge-root`/`*cell*`/default/spool/update.`*pid*. Please note that If you chose a non-default directory, the

    If necessary, you can use `qconf -mc` to reconfigure your complexes. For details, see Chapter 3, "Configuring Complex Resource Attributes," in *N1 Grid Engine 6 Administration Guide*.

22. **Reconfigure your queues.**
    During the upgrade process, a single default cluster queue is created. Within this queue you will find all of your installed execution hosts. It is recommended that you reconfigure your queues. For details, see "Configuring Queues" in *N1 Grid Engine 6 Administration Guide*.

CHAPTER 6

Verifying the Installation

The verification phase includes the following tasks:

- Ensuring that the daemons are running on the master host
- Ensuring that the daemons are running on all execution hosts
- Ensuring that you can run simple commands
- Submitting test jobs

Verifying the Installation

To ensure that the grid engine system daemons are running, look for the `sge_qmaster` and `sge_schedd` daemons on the master host, and then the `sge_execd` daemon on the execution hosts. Once you have verified that the daemons are running, you can try to use commands, and prepare to submit jobs.

Note – If no cell name was specified during installation, the value of *cell* is `default`.

▼ How to Verify That the Daemons Are Running on the Master Host

Steps

1. **Log in to the master host.**
2. **Verify that the daemons are running by typing one of the following commands, depending on the operating system you are running.**
   - **On BSD-based UNIX systems, type the following command.**

```
% ps -ax | grep sge
```

- **On systems running a UNIX System 5-based operating system (such as the Solaris Operating System), type the following command.**

```
% ps -ef | grep sge
```

3. **Verify that the daemons are running by looking through the output for `sge` strings that are similar to the following examples.**

   Specifically, you should see that the `sge_qmaster` daemon and the `sge_schedd` daemon are running.

   - **On a BSD-based UNIX system, you should see output such as the following example.**

```
14676 p1 S <  4:47 /gridware/sge/bin/solaris/sge_qmaster

14678 p1 S <  9:22 /gridware/sge/bin/solaris/sge_schedd
```

   - **On a UNIX System 5-based system, you should see output such as the following example.**

```
root 439 1 0 Jun 2 ? 3:37 /gridware/sge/bin/solaris/sge_qmaster

root 446 1 0 Jun 2 ? 3:37 /gridware/sge/bin/solaris/sge_schedd
```

4. **If you *do not see* the appropriate strings, one or more daemons that are required on the master host are not running.**

   a. **Look in the file *sge-root*`/`*cell*`/common/act_qmaster` to see if you really are on the master host.**

   b. **Restart the daemons by hand.**

      To start the master host daemons, `sge_qmaster` and `sge_schedd`:

```
# sge-root/cell/common/sgemaster  start
```

5. **Continue the verification process.**

   After you have verified that the master host and the execution host daemons are running, continue the verification process. See "How to Run Simple Commands" on page 81.

▼ How to Verify That the Daemons Are Running on the Execution Hosts

Steps

1. **Log in to the execution hosts on which you ran the execution host installation procedure.**

2. **Verify that the daemons are running by typing one of the following commands, depending on the operating system you are running.**

   - **On BSD-based UNIX systems, type the following command.**

     ```
     % ps -ax | grep sge
     ```

   - **On systems running a UNIX System 5--based operating system (such as the Solaris Operating System), type the following command.**

     ```
     % ps -ef | grep sge
     ```

3. **Verify the daemons are running by looking for the `sge_execd` string in the output.**

   Specifically, you should see that the `sge_execd` daemon is running.

   - **On a BSD-based UNIX system, you should see output such as the following example.**

     ```
     14688 p1 S <    4:27  /gridware/sge/bin/solaris/sge_execd
     ```

   - **On a UNIX System 5-based system, such as the Solaris Operating System, you should see output such as the following example.**

     ```
     root 171 1 0 Jun 22 ? 7:11 /gridware/sge/bin/solaris/sge_execd
     ```

4. **If you** *do not see* **similar output, the daemon required on the execution host is not running. Restart the daemon by hand.**

   # *sge-root*/*cell*/**common/sgeexecd start**

5. **Continue the verification process.**

   After you have verified that the master host and the execution host daemons are running, continue the verification process. See "How to Run Simple Commands" on page 81.

▼ How to Run Simple Commands

If both the necessary daemons are running on the master and execution hosts, the grid engine software should be operational. Check by issuing a trial command.

Steps

1. **Log in to either the master host or another administrative host.**

   In your standard search path, make sure to include *sge-root*/`bin`.

2. **From the command line, type the following command.**

   ```
   % qconf -sconf
   ```

   This `qconf` command displays the current global cluster configuration (see "Basic Cluster Configuration" in *N1 Grid Engine 6 Administration Guide*).

If this command fails, your SGE_ROOT environment variable is not set correctly.

a. **Check whether the environment variables `SGE_EXECD_PORT` and `SGE_QMASTER_PORT` are set in the script files,** *sge-root*`/`*cell*`/common/settings.csh` **or** *sge-root*`/`*cell*`/common/settings.sh`**.**

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If so, make sure that the environment variables `SGE_EXECD_PORT` and `SGE_QMASTER_PORT` are set to the correct value before you try the command again.**
- **If not, verify whether your NIS services map contains entries for `sge_qmaster` and `sge_execd`.**

  If the SGE_EXECD_PORT and SGE_QMASTER_PORT variables are not used in these files, then the services database (for example, /etc/services or the NIS services map) on the machine from which you run the command must provide entries for both sge_qmaster and sge_execd. If it doesn't, add such an entry to the machine's services database, giving it the same value as is configured on the master host.

d. **Retry the `qconf` command.**

3. **Try to submit test jobs.**

▼ How to Submit Test Jobs

Before you start submitting batch scripts to the grid engine system, check to see whether your site's standard shell resource files (.cshrc, .profile or ..kshrc) as well as your personal resource files contain commands such as stty. Batch jobs do not have a terminal connection by default, and therefore calls to stty result in an error.

Steps

1. **Log in to the master host.**
2. **Type the following command.**

   % **rsh** *exec-host-name* **date**

   exec-host-name refers to one of the already installed execution hosts. You should try this test on all execution hosts if your login or home directories differ from host to host. The rsh command should give you output similar to the date command run locally on the master host. If there are any additional lines containing error messages, you must fix the cause of the errors before you can run a batch job successfully.

For all command interpreters you can check on an actual terminal connection before you run a command such as `stty`.

The following is an example of a Bourne shell script to test the terminal connection.

```
tty -s
if [ $? = 0 ]; then
   stty erase ^H
fi
```

The following example shows C Shell syntax.

```
tty -s
if ( $status = 0 ) then
   stty erase ^H
endif
```

3. **Submit one of the sample scripts contained in the** *sge-root*`/examples/jobs` **directory.**

   ```
   % qsub sge-root/examples/jobs/simple.sh
   ```

4. **Use the `qstat` command to monitor the job's behavior.**

   For more information about submitting and monitoring batch jobs, see "Submitting Batch Jobs" in *N1 Grid Engine 6 User's Guide*.

5. **After the job finishes executing, check your home directory for the redirected stdout/stderr files** *script-name*`.e`*job-id* **and** *script-name*`.o`*job-id*.

   job-id is a consecutive unique integer number assigned to each job.

   In case of problems, see Chapter 8, "Fine Tuning, Error Messages, and Troubleshooting," in *N1 Grid Engine 6 Administration Guide*.

CHAPTER 7

Removing the Grid Engine Software

This chapter describes the two ways that the grid engine software can be removed:

- "Removing the Software Interactively" on page 85
- "Removing the Software Using the `inst_sge` Utility and a Configuration Template" on page 86

Removing the Software Interactively

Note – Remove the software from the execution hosts before removing it from the master host. If you remove the software from the master host first, you cannot automate the removal of the software from the execution hosts.

▼ How to Remove the Software Interactively

Steps 1. **Ensure that your environment variables are set up properly.**

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  # **`source`** *sge-root*`/`*cell*`/common/settings.csh`

- **If you are using a Bourne shell or Korn shell, type the following command:**

```
# . sge-root/cell/common/settings.sh
```

2. **On the master host, issue the *sge-root*`/inst_sge` command.**

   This example uninstalls the execution hosts, `host1`, `host2` and `host3`.

   ```
   # sge-root/inst_sge -ux -host "host1 host2 host3"
   ```

Note – You are not prompted for any information during this process. However, the output from this process will be displayed to the terminal window where you run the command.

---

3. **(Optional) If you have any shadow master hosts, uninstall them now by typing:**

   ```
   # sge-root/inst_sge -usm -host "host4"
   ```

4. **Uninstall the master host.**

   ```
   # sge-root/inst_sge -um
   ```

Removing the Software Using the `inst_sge` Utility and a Configuration Template

One difference between interactive and automated uninstallation is that the automated method suppresses output during the process. The other difference is that the automate method requires a properly formatted configuration file.

Note – Remove the software from the execution hosts before removing it from the master host. If you remove the software from the master host first, you cannot automate the removal of the software from the execution hosts.

▼ How to Remove the Software Interactively

Steps

1. **Ensure that your environment variables are set up properly.**

Note – If no cell name was specified during installation, the value of *cell* is `default`.

- **If you are using a C shell, type the following command:**

  # **source** *sge-root*/*cell*/**common/settings.csh**

- **If you are using a Bourne-shell or Korn shell, type the following command:**

  # **.** *sge-root*/*cell*/**common/settings.sh**

2. **Create a copy of the configuration template,** *sge-root*`/util/inst_sge_modules/inst_sge_template.conf`**.**

```
# cd sge-root/util/inst_sge_modules/
# cp inst_sge_template.conf  my_configuration.conf
```

3. **Edit your configuration template.**

   Every host that is in the `EXEC_HOST_LIST_RM` list will be automatically removed.

```
# Remove this execution hosts in automatic mode
EXEC_HOST_LIST_RM="host1 host2 host3 host4"
```

4. **On the master host type the** *sge-root*`/inst_sge` **command.**

   This example uninstalls the execution hosts, `host1`, `host2` and `host3`.

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `-auto` and the `sge-root/util/inst_sge_modules/my_configuration.conf` components.

```
# sge-root/inst_sge -ux -auto
sge-root/util/inst_sge_modules/my_configuration.conf
```

Note – You are not prompted for any information during this process. However, the output from this process will be displayed to the terminal window where you run the command.

5. **(Optional) If you have any shadow master hosts, uninstall them.**

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `-auto` and the `sge-root/util/inst_sge_modules/my_configuration.conf` components.

```
# sge-root/inst_sge -usm -auto
sge-root/util/inst_sge_modules/my_configuration.conf
```

6. **Uninstall the master host.**

Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `-auto` and the `sge-root/util/inst_sge_modules/my_configuration.conf` components.

```
# sge-root/inst_sge -um -auto
sge-root/util/inst_sge_modules/my_configuration.conf
```

CHAPTER 8

Installing the Accounting and Reporting Console

This chapter describes the steps to install the accounting and reporting console software:

- "Setting Up the Database Software" on page 89
- "Install the Accounting and Reporting Software" on page 95

Setting Up the Database Software

Database software must be installed and configured properly before you can install and use the Accounting and Reporting Console. This section describes how to set up the database software using either a PostgreSQL database or an Oracle database.

Set up the PostgreSQL Database Software

Detailed information on the PostgreSQL database software can be found in the Postgres documentation (`http://www.postgresql.org/docs/7.4/static/index.html`).

▼ How to Start the Database Server

Once the PostgreSQL software is installed, you can start the database server.

Before You Begin Download, compile, and install the PostgreSQL database software. Create a user account to own the database processes. Usually, this user is `postgres`. Add the PostgreSQL `bin` directory and necessary `LD_LIBRARY_PATH` settings to your environment.

Steps

1. **Create a home directory for the `postgres` user.**

   In this example, the home directory is `/space/postgres/data`.

   ```
   % mkdir -p /space/postgres/data
   % useradd -d /space/postgres  postgres
   % chown postgres /space/postgres/data
   % su - postgres
   ```

2. **Continue as described in the PostgreSQL documentation to set up a database.**

   ```
   > initdb -D /space/postgres/data

   creating directory /space/postgres/data... ok
   creating directory /space/postgres/data/base... ok
   creating directory /space/postgres/data/global... ok
   creating directory /space/postgres/data/pg_xlog... ok
   creating directory /space/postgres/data/pg_clog... ok
   creating template1 database in /space/postgres/data/base/1... ok
   creating configuration files... ok
   initializing pg_shadow... ok
   enabling unlimited row size for system tables... ok
   initializing pg_depend... ok
   creating system views... ok
   loading pg_description... ok
   creating conversions... ok
   setting privileges on built-in objects... ok
   vacuuming database template1... ok
   copying template1 to template0... ok

   Success. You can now start the database server using:
      postmaster -D /space/postgres/data
      or
      pg_ctl -D /space/postgres/data -l logfile start
   ```

3. **Make the following changes to the `pg_hba.conf` file.**

   This change permits unrestricted and password free access to the database superuser `postgres` but requires md5 encrypted passwords for all other database users. Replace `nnn.nnn.nnn` with your subnet address without the trailing `0`. You also can add access rules on a per-host basis by adding similar lines with host IP addresses.

   ```
   # TYPE  DATABASE  USER  IP-ADDRESS     IP-MASK         METHOD
   local   all       postgres                             trust
   local   all       all                                  md5
   # IPv4-style local connections:
   #host    all       all   nnn.nnn.nnn.0  255.255.255.0  md5
   ```

4. **Make the following changes to the `postgresql.conf` file, to enable TCP/IP access from other hosts.**

Note – Ensure that the value of `shared_buffers` is at least twice the value of `max_connections`.

```
tcpip_socket = true
max_connections = 40   (increase if necessary)
```

5. **Start the database.**

   In this example, `-i` enables TCP/IP communication, and `-S` is for "silent" mode.

```
> postmaster -S -i
```

6. **Verify the installation.**

   As the `postgres` user, try the following commands:

```
% su - postgres

> createuser -P test_user
Enter password for new user:
Enter it again:
Shall the new user be allowed to create databases? (y/n) y
Shall the new user be allowed to create more new users? (y/n) n
CREATE USER

> createdb -O test_user -E UNICODE test
CREATE DATABASE
```

7. **Execute commands as the database super user.**

```
 > psql test
Welcome to psql 7.3, the PostgreSQL interactive terminal.
Type:  \copyright for distribution terms
       \h for help with SQL commands
       \? for help on internal slash commands
       \g or terminate with semicolon to execute query
       \q to quit


test=# create table test (x int, y text);
CREATE TABLE
test=# insert into test values (1, 'one');
INSERT 16982 1
test=# insert into test values (2, 'two');
INSERT 16983 1
test=# select * from test;
x |  y
---+------
1 | one
2 | two
(2 rows)

test=# \q
```

```
 > psql -U test_user test
Password:
Welcome to psql 7.4.1, the PostgreSQL interactive terminal.

Type:  \copyright for distribution terms
       \h for help with SQL commands
       \? for help on internal slash commands
       \g or terminate with semicolon to execute query
       \q to quit
test=>
```

8. **After you have successfully tested your database, you should proceed to the next task, "How to Set Up a PostgreSQL Database" on page 92.**

▼ How to Set Up a PostgreSQL Database

Steps

1. **Log in as the database superuser, for example, `postgres`.**

```
# su - postgres
```

2. **Create the owner of the database.**

```
> createuser -P arco_write
Enter password for new user:
Enter it again:
Shall the new user be allowed to create databases? (y/n) y
Shall the new user be allowed to create more new users? (y/n) n
CREATE USER
```

3. **Create the database for accounting and reporting.**

```
> createdb -O arco_write arco
CREATE DATABASE
```

4. **Create the tables and views of `arco_data` as user `arco_write`.**

```
> psql -f sge-root/dbwriter/database/postgres/setup.sql arco arco_write
> psql -f sge-root/dbwriter/database/postgres/view.sql arco arco_write
```

5. **Crate a database user for reading the database.**

   If you do not use the `arco_read` user, you will need to modify the privileges script that is used in the next step.

```
> createuser -P arco_read
Enter password for new user:
Enter it again:
Shall the new user be allowed to create databases? (y/n) n
Shall the new user be allowed to create more new users? (y/n) n
CREATE USER
```

6. **Grant `arco_user` the necessary permissions.**

```
> psql -f sge-root/dbwriter/database/postgres/privileges.sql arco arco_write
```

▼ How to Set Up an Oracle Database

Steps

1. **Ask your database administrator for an instance of an Oracle database.**

   You need two database users for this instance, `arco_write` and `arco_read`. The `arco_write` user must be able to create or alter tables, views, and indexes. Privileges for the `arco_read` user are set later by an SQL script.

2. **Ask your database administrator for the connection string to the database.**

   To set up the database, you also need the `sqlplus` utility in your `PATH`.

3. **Set your environment variables.**

   - **If you are using a Bourne shell or Korn shell, type the following command:**

     ```
     > . /usr/local/bin/oraenv ORACLE_SID = [] ? SID of your database
     ORACLE_HOME = [/space/oracle] ? path to your oracle installation
     ```

   - **If you are using a C shell, type the following command:**

     ```
     > source  /usr/local/bin/coraenv ORACLE_SID = [] ? SID of your database
     ORACLE_HOME = [/space/oracle] ? path to your oracle installation
     ```

4. **Create the database tables.**

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `$ORACLE_HOME/bin/sqlplus` and the `arco_write@my-host.my-domain` components.

   ```
   > $ORACLE_HOME/bin/sqlplus
        arco_write@my-host.my-domain @sge-root/dbwriter/database/oracle/setup.sql

   SQL*Plus: Release 9.0.1.0.0 - Production on Thu May 13 11:14:06 2004

   (c) Copyright 2001 Oracle Corporation.  All rights reserved.

   Enter password:

   Connected to:
   Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
   With the Partitioning option
   JServer Release 9.0.1.0.0 - Production


   Table created.
   ...
   Disconnected from Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
   With the Partitioning option
   JServer Release 9.0.1.0.0 - Production
   ```

5. **Create the views.**

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `$ORACLE_HOME/bin/sqlplus` and the `arco_write@my-host.my-domain` components.

```
> $ORACLE_HOME/bin/sqlplus
     arco_write@my-host.my-domain @sge-root/dbwriter/database/oracle/view.sql

SQL*Plus: Release 9.0.1.0.0 - Production on Thu May 13 11:20:05 2004

(c) Copyright 2001 Oracle Corporation.  All rights reserved.

Enter password:

Connected to:
Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
With the Partitioning option
JServer Release 9.0.1.0.0 - Production


View created.

Commit complete.

Disconnected from Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
With the Partitioning option
JServer Release 9.0.1.0.0 - Production
```

6. **Grant `arco_user` the necessary permissions.**

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `arco_write@my-host.my-domain` and the `@sge-root/dbwriter/database/oracle/privileges.sql` components.

```
> $ORACLE_HOME/bin/sqlplus arco_write@my-host.my-domain
     @sge-root/dbwriter/database/oracle/privileges.sql
SQL*Plus: Release 9.0.1.0.0 - Production on Thu May 13 11:20:05 2004

(c) Copyright 2001 Oracle Corporation.  All rights reserved.

Enter password:

Connected to:
Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
With the Partitioning option
JServer Release 9.0.1.0.0 - Production

Grant succeeded.
...
Synonym created.
...
Commit complete.

Disconnected from Oracle9i Enterprise Edition Release 9.0.1.0.0 - Production
```

```
With the Partitioning option
JServer Release 9.0.1.0.0 - Production
```

Install the Accounting and Reporting Software

▼ How to Set Up `dbWriter`

Steps

1. **Extract the accounting and reporting software.**

   - **The `tar` Method**

     # **cd** *sge-root*

     Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `-dc` and the path to the tar file.

     # **gunzip -dc**
     cdrom_mount_point/**N1_Grid_Engine_6/ARCo/tar/n1ge-6_0-arco-1_0.tar.gz** | **tar xvpf -**

   - **The `pkgadd` Method**

     At the command prompt, type the following command and respond to the script questions.

     # **cd** *cdrom_mount_point*/**N1_Grid_Engine_6/ARCo/Packages**
     # **pkgadd -d . SUNWsgeea**

2. **As the administrative user, set your environment variables.**

   - **If you are using a Bourne shell or Korn shell, type the following command:**

     $ **.** *sge-root*/**default/common/settings.sh**

   - **If you are using a C shell, type the following command:**

     % **source** *sge-root*/**default/common/settings.csh**

3. **Change the global configuration to enable reporting.**

   For details on how to enable reporting, see "Report Statistics (ARCo)" in *N1 Grid Engine 6 Administration Guide*.

   ```
   % qconf -mconf
   <......>
   reporting_params    accounting=true \
   ```

```
 reporting=true flush_time=00:00:15 joblog=true \
 sharelog=00:00:00<......>
```

By default, no report variables are activated. You can use the `qconf` command to enable statistics gathering on specific variables, as shown in the following example.

```
% qconf -me global
hostname              global
<......>
report_variables      cpu,np_load_avg,mem_free,virtual_free
<......>
```

4. **Install the `dbWriter` software.**

```
# cd sge-root/dbwriter
# ./inst_dbwriter
Welcome to the Grid Engine installation
---------------------------------------

Grid Engine dbWriter installation
---------------------------------
The dbWriter installation will take approximately 5 minutes

Hit <RETURN> to continue >>
```

5. **Specify the location of *sge-root*, and the name of your cell.**

```
Generic Parameters
------------------
Please enter your SGE_ROOT [<PATH_TO_SGE>] >> /opt/n1ge6

Now the name of the cell is needed:
Please enter your SGE_CELL [default] >> default
```

6. **Specify the location of your Java™ Software Development Kit.**

   Version 1.4.1 or higher is required. If your `JAVA_HOME` environment variable is set, the script will use that as a default value.

```
Please enter the path to your java 1.4 installation [/usr/java] >>/opt/j2sdk1.4.1_01
```

7. **Specify the connection parameters for the reporting database.**

```
Setup your database connection parameters
-----------------------------------------

Enter your database type ( o = Oracle, p = PostgreSQL ) [p] >> p

Please enter the name of your postgres db host []>> my-host.my-domain

Please enter the port of your postgres db [5432] >>

Please enter the name of your postgres database [arco] >>
```

8. **Enter the user name and password of the database user.**

```
Please enter the name of the database user [arco_write] >>

Please enter the password of the database user >>
Please retype the password >>
```

9. **Enter the name of the database schema.**

   - **If the JDBC driver is found, you will be asked for the name of the schema.**

     For PostgreSQL this value is normally `public`. For oracle this should be the username of the database owner (`arco_write`).

     ```
     Please enter the name of the database schema [public] >>

     Search for the jdbc driver org.postgresql.Driver
     in directory /opt/n1ge6/dbwriter/lib ..........
     found in /opt/n1ge6/dbwriter/lib/pg73jdbc2.jar
     ```

   - **If no corresponding JDBC driver is found, you will be asked to copy the JAR file of the database driver into the** *sge-root*`/dbwriter/lib` **directory.**

     After you have copied the JAR file to the correct location, press `RETURN`, and the search will be repeated.

     ```
     Please copy your driver jar into the
     directory /opt/n1ge6//dbwriter/lib
     Press enter to continue >>
     ```

10. **Verify that the database has been set up correctly.**

    ```
    Should the connection to the database be tested? (y/n) [y] >>
    ```

    - **If the test succeeds, you will be notified.**

      ```
      Test db connection to 'jdbc:postgresql://my-host.my-domain:5432/arco' ... OK
      ```

    - **If the test fails, you are given the option to repeat the set up procedure.**

      ```
      Test db connection to 'jdbc:postgresql://my-host.my-domain:5432/arco' ... Failed
      Do you want to repeat database connection setup? (y/n) [y] >>
      ```

11. **Specify how often the `dbWriter` program should check the grid engine system log files for new data.**

    ```
    Please enter the interval between two dbwriter runs in seconds [60] >>
    ```

12. **Enter the location of the file containing the rules for calculating derived values.**

    Based on the reporting and accounting data, the `dbWriter` program calculates some derived values. The *derived values* are reporting information that are calculated from one or more raw values. The `dbWriter` program also allows you to delete data from the reporting database, when that data is not longer needed. The rules governing the derived value calculations and data deletion are stored in a file. The `dbWriter` program distribution contains an example file for a PostgreSQL database and an example file for an Oracle database. The example files are `$SGE_ROOT/dbwriter/database/postgres/dbwriter.xml` and

```
$SGE_ROOT/dbwriter/database/oracle/dbwriter.xml
```

You can accept the default file location, or you can specify a alternative path.

```
Please enter the file with the derived value rules
[$SGE_ROOT/dbwriter/database/postgres/dbwriter.xml] >>
```

13. Enter the level of logging that the `dbWriter` program should use.

```
The dbWriter can run with different debug levels
Possible values: WARNING INFO CONFIG FINE FINER FINEST
Please enter the debug level of the dbwriter [INFO] >>
```

14. Verify the settings.

If you answer `n`, you are given the option to repeat the set up.

```
All parameters are now collected
--------------------------------

      SGE_ROOT=/opt/n1ge6
      SGE_CELL=default
     JAVA_HOME=/usr/java (java version "1.4.1")
        DB_URL=jdbc:postgresql://my-host.my-domain:5432/arco
       DB_USER=arco_write
      INTERVAL=60
REPORTING_FILE=/opt/n1ge6/default/common/reporting
  DERIVED_FILE=/opt/n1ge6/dbwriter/database/postgres/dbwriter.xml
   DEBUG_LEVEL=INFO

Are this settings correct? (y/n) [y] >> y
```

The installation scripts creates a start script, *sge-root*`/dbwriter/bin/sgedbwriter`. Your configuration parameters are stored in *sge-root*/*cell*`/common/dbwriter.conf`. If any configuration changes are necessary, you can edit the `dbwriter.conf` file or repeat the installation script.

```
Create configuration file for dbWriter in
/opt/n1ge6/default/common
Installation of dbWriter completed
Start the dbWriter with
/opt/n1ge6/dbwriter/bin/sgedbwriter start
>
```

15. If you want the `dbWriter` program to start at boot time, copy the start script to the `/etc/init.d` directory.

```
# cp sge-root/cell/dbwriter/bin/sgedbwriter /etc/init.d
# ln -s /etc/init.d/sgedbwriter /etc/rc2.d/S98sgedbwriter
```

16. Start the `dbWriter` program.

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

▼ How to Install Sun Web Console

Steps

1. **Extract the web console package under a temporary directory.**

   ```
   # cd /tmp
   # umask 022
   # mkdir swc
   # cd swc
   ```

   Type the following command as one string (the command is too long to fit on one line in this guide), with a space between the `xvf` and the path to the tar file.

   ```
   # tar xvf
   cdrom_mount_point/N1_Grid_Engine_6/SunWebConsole/tar/swc_sparc_2.0.3.tar
   ```

2. **If the user `noaccess` and the group `noaccess` do not exist in your password file or NIS passwd map, add them by typing:**

   ```
   # groupadd -g 60002 noaccess
   # useradd -u 60002 -g 60002 -d /tmp -s /bin/csh -c "No Access User" noaccess
   ```

3. **If you are running SuSE 9.0, create symbolic links for each of the `/etc/rc#.d` directories.**

   ```
   # ln -s /etc/rc.d/rc0.d /etc/rc0.d
   # ln -s /etc/rc.d/rc1.d /etc/rc1.d
   # ln -s /etc/rc.d/rc2.d /etc/rc2.d
   ```

4. **Run the Sun Web Console `setup` script.**

   ```
   # ./setup
   <....>
   Installation complete.

   Starting Sun(TM) Web Console Version 2.0.3...
   See /var/log/webconsole/console_debug_log for server logging information
   ```

   After running the setup script, the Sun Web Console is started. You can stop, start, or restart the console at any time, using the following commands:

   ```
   # /usr/sadm/bin/smcwebserver start
   # /usr/sadm/bin/smcwebserver stop
   # /usr/sadm/bin/smcwebserver restart
   ```

5. **Connect to the web console, by accessing the following URL in your browser. Replace *hostname* with the name of your master host.**

   `https://`*hostname*`:6789`

   The web console login page is displayed.

6. **Login with a valid UNIX account.**

   Once you have logged in, you should see the examples that are shipped with the web console. Now you can continue to install the accounting and reporting console.

▼ How to Install the Accounting and Reporting Console

Steps

1. **Change directory to *sge-root*`/reporting`.**

   ```
   # cd sge-root/reporting
   ```

2. **Use the `inst_reporting` script to install the software.**

   ```
   # ./inst_reporting

   Welcome to the  N1 SGE reporting module installation
   ----------------------------------------------------
   The installation will take approximately 5 minutes

   Hit <RETURN> to continue >>
   ```

3. **Enter the path to your Java Software Development Kit.**

   Version 1.4.1 or higher is required. If your `JAVA_HOME` environment variable is set, the script will use that as a default value.

   ```
   Please enter the path to your java 1.4 installation [] >> /opt/j2sdk1.4.1_01
   ```

4. **Specify the directory where the accounting and reporting software can store its results.**

   If this directory does not exist, it will be created for you.

   ```
   Spool directory
   ---------------

   In the spool directory the N1 SGE reporting module will
   store all queries and results

   Please enter the path to the spool directory [/var/spool/arco] >>
   ```

5. **Specify the parameters for the database connection.**

   ```
   Database Setup
   --------------

   Enter your database type ( o = Oracle, p = PostgreSQL ) [p] >> o


   Please enter the name of your oracle db host [] >> my-host

   Please enter the port of your oracle db [1521] >>

   Please enter the name of your oracle database [arco] >>
   ```

6. **Specify an accounting and reporting database user.**

   Do not use the same database user as you did for the `dbwriter` application. For security reasons, it is recommended that the database user for accounting and

reporting only have read permission for the database tables.

```
Please enter the name of the database user [arco_read] >>

Please enter the password of the database user >>
Please retype the password >>

Please enter the name of the database schema [arco_write] >>

Search for the jdbc driver oracle.jdbc.driver.OracleDriver
in directory /opt/n1ge/reporting/WEB-INF/lib ...
found in /opt/n1ge/reporting/WEB-INF/lib/classes12.jar
Should the connection to the database be tested? (y/n) [y] >> y

Test db connection to 'jdbc:oracle:thin:@my-system:1521:arco' ... OK
h150
```

7. **Enter the login names of users who are allowed to store the queries and results.**

```
Configure users with write access
---------------------------------
Enter a login name of a user (Press enter to finish) >> user1
Users: user1
Enter a login name of a user (Press enter to finish) >> user2
Users: user1 user2
Enter a login name of a user (Press enter to finish) >>
```

8. **Verify the information.**

```
All parameters are now collected
--------------------------------
       SPOOL_DIR=/var/spool/arco
          DB_URL=jdbc:oracle:thin://my-system:1521/arco
         DB_USER=arco_read
ARCO_WRITE_USERS=user1 user2

Are this settings correct? (y/n) [y] >> y

Shutting down Sun(TM) Web Console Version 2.0.3...
```

9. **Create the query directory.**

   If the query directory already exists, the predefined queries will not be installed. To install this queries manually, you can copy them from *sge-root*/reporting/database/example_queries/queries.

```
Directory /var/spool/arco does not exist, create it? (y/n) [y] >> y

Create query directory /var/spool/arco
Create query directory /var/spool/arco/queries
Copy predefined queries into /var/spool/arco/queries
Create query directory /var/spool/arco/results
Create query directory /opt/n1ge6/reporting/charts
Register the N1 SGE reporting module in the webconsole

    Registering com.sun.grid.arco_6.
```

```
Starting Sun(TM) Web Console Version 2.0.3...
See /var/log/webconsole/console_debug_log for server logging information
```

10. **Check the log file for error or warning messages.**

    ```
    # more /var/log/webconsole/console_debug_log
    ```

    The accounting and reporting logs are written to the `/var/log/webconsole/console_debug_log` file. The default log level is `INFO`, but you can modify the log level from the command line:

    ```
    # smreg add -p -e arco_logging_level=FINE
    ```

    The new log take effect the next time the console is started or restarted. The possible log levels are `WARNING`, `INFO`, `FINE`, `FINER`, and `FINEST`.

11. **Connect to the Sun Web Console by accessing the following URL in your browser:**

    `https://`*hostname*`:6789`

12. **Login with your UNIX account.**

13. **Select N1 Grid Engine 6 Accounting and Reporting Console.**

Glossary

administration host — Administration hosts are hosts that have permission to carry out administrative activity for the grid engine system.

access list — A list of users and UNIX groups who are permitted or denied access to a resource such as a queue or a host. Users and groups can belong to multiple access lists, and the same access lists can be used in various contexts.

array job — A job made up of a range of independent identical tasks. Each task is similar to a separate job. Array job tasks differ among themselves only by having unique task identifiers, which are integer numbers.

campus grid — A grid that enables multiple projects or departments within an organization to share computing resources.

cell — A separate cluster with a separate configuration and a separate master machine. Cells can be used to loosely couple separate administrative units.

checkpointing — A procedure that saves the execution status of a job into a *checkpoint*, thereby allowing the job to be aborted and resumed later without loss of information and already completed work. The process is called *migration* if the checkpoint is moved to another host before execution resumes.

checkpointing environment — A grid engine system configuration entity that defines events, interfaces, and actions that are associated with a certain method of checkpointing.

cluster — A collection of machines, called hosts, on which grid engine system functions occur.

cluster grid — The simplest form of a grid, consisting of computer *hosts* working together to provide a single point of access to users in a single project or department.

cluster queue — A container for a class of jobs that are allowed to run concurrently. A queue determines certain job attributes, for example, whether it can be migrated. Throughout its lifetime, a running job is associated with its queue. Association with a queue affects some of the things that can happen to a job. For example, if a queue is suspended, all jobs associated with that queue are also suspended.

complex — A set of resource attribute definitions that can be associated with a queue, a host, or the entire cluster.

department — A list of users and groups who are treated alike in the functional and override scheduling policies of the grid engine system. Users and groups can belong to only one department.

entitlement — The same as *share*. The amount of resources that are planned to be consumed by a certain job, user, user group, or project.

execution host — Systems that have permission to run grid engine system jobs. These systems host queue instances, and run the execution daemon `sge_execd`.

functional policy — A policy that assigns specific levels of importance to jobs, users, user groups, and projects. For instance, through the functional policy, a high-priority project and all its jobs can receive a higher resource share than a low-priority project.

global grid — A collection of campus grids that cross organizational boundaries to create very large virtual systems.

grid — A collection of computing resources that perform tasks. Users treat the grid as a *single* computational resource.

group — A UNIX group.

hard resource requirements — The resources that must be allocated before a job can be started. Contrast with *soft resource requirements*.

host — A system on which grid engine system functions occur.

job — A request from a user for computational resources from the grid.

batch job — A batch job is a UNIX shell script that can be run without user intervention and does not require access to a terminal.

interactive job — An interactive job is a session started with the commands `qrsh`, `qsh`, or `qlogin`, which open an *xterm* window for user interaction or provide the equivalent of a remote login session.

job class — A set of jobs that are equivalent in some sense and treated similarly. A job class is defined by the identical requirements of the corresponding jobs and by the characteristics of the queues that are suitable for those jobs.

manager	A user who can manipulate all aspects of the grid engine software. The superusers of the master host and of any other machine that is declared to be an administration host have manager privileges. Manager privileges can be assigned to nonroot user accounts as well.
migration	The process of moving a checkpointing job from one host to another before execution of the job resumes.
operator	Users who can perform the same commands as managers except that they cannot change the configuration. Operators are supposed to maintain operation.
override policy	A policy commonly used to override the automated resource entitlement management of the functional and share-based policies. The cluster administrator can modify the automated policy implementation to assign override to jobs, users, user groups, and projects.
owner	Users who can suspend or resume, and disable or enable, the queues they own. Typically, users are owners of the queue instances that reside on their workstations.
parallel environment	A grid engine system configuration that defines the necessary interfaces for the grid engine software to correctly handle parallel jobs.
parallel job	A job that is made up of more than one closely correlated task. Tasks can be distributed across multiple hosts. Parallel jobs usually use communication tools such as shared memory or message passing (MPI, PVM) to synchronize and correlate tasks.
policy	A set of rules and configurations that the administrator can use to define the behavior of the grid engine system. Policies are implemented automatically by the system.
priority	The relative level of importance of a job compared to others.
project	A grid engine system project.
resource	A computational device consumed or occupied by running jobs. Typical examples are memory, CPU, I/O bandwidth, file space, software licenses, and so forth.
master host	The master host is central to the overall cluster activity. It runs the master daemon `sge_qmaster` and the scheduler daemon `sge_schedd`. By default, the master host is also an administration host and a submit host.
share	The same as *entitlement*. The amount of resources that are planned to be consumed by a certain job, user, or project.
share-based policy	A policy that allows definition of the entitlements of user and projects and arbitrary groups thereof in a hierarchical fashion. An enterprise, for instance, can be subdivided into divisions, departments, projects

	active in the departments, user groups working on those projects, and users in those user groups. The share-based hierarchy is called a share-tree, and once a share-tree is defined, its entitlement distribution is automatically implemented by the grid engine software.
share-tree	The hierarchical definition of a share-based policy.
soft resource requirements	Resources that a job needs but that do not have to be allocated before a job can be started. Allocated to a job on an as-available basis. Contrast with *hard resource requirements*.
submit host	Submit hosts allow for submitting and controlling *batch jobs only*. In particular, a user who is logged in to a submit host can submit jobs using `qsub`, can control the job status using `qstat`, and can use the grid engine system OSF/1 Motif graphical user interface `QMON`.
suspension	The process of holding a running job but keeping it on the execution host (in contrast to checkpointing, where the job is aborted). A suspended job still consumes some resources, such as swap memory or file space.
ticket	A generic unit for resource share definition. The more ticket shares that a job, user, project, or other component has, the more important it is. If a job has twice as many tickets as another job, for example, that job is entitled to twice the resource consumption.
usage	Another term for "resources consumed." Usage is determined by an administrator-configurable weighted sum of CPU time consumed, memory occupied over time, and amount of I/O performed.
users	People who can submit jobs to the grid and run them if they have a valid login ID on at least one submit host and one execution host.
userset	Either an *access list* or a *department*.

Index

A

abort installation procedure, 28, 46
access list, 103
account for installation, 16
accounting and reporting
  console
    install, 100
  `dbWriter`
    enable reporting, 95
accounting and reporting console, install, 100
add administrative host, 45, 59, 60
add submit host, 46
administration host, 103
administrative host, 20
  add, 45, 59, 60
administrative user, file handling, 16
array, job, 103
automounter, 13

B

batch job, monitoring, 83
Berkeley DB host, 18
Berkeley DB installation procedure, 46

C

campus grid, 103
cell, 11, 13, 15, 103
checkpointing, 103
checkpointing environment, 103
classic spooling, 18
cluster, 103
cluster grid, 103
complex, 104

D

database
  for accounting and reporting, 89
    Oracle, 93
    PostgreSQL, 89, 92
database host, 18
database spooling host, 18
`date`, 82
`dbWriter`
  for accounting and reporting
    enable reporting, 95
    install, 95
    set environment variables, 95
department, 104
directory organization, 14
disk space requirements, 13

E

entitlement, 104
environment
  checkpointing, 103
  parallel, 105
/etc/services, 12, 16
execution host, 19, 104

execution host (Continued)
installation procedure, 41
execution host spool directory, 18

F

file access permission, 16
file handling, administrative user, 16
functional policy, 104

G

global grid, 104
grid, 104
campus, 103
cluster, 103
global, 104
group, 104

H

hard resource requirements, 104
host, 104
administration, 103
Berkeley DB, 18
execution, 104
master, 105
shadow master, 17
submit, 106

I

id, equivalent user, 11
installation
with root account, 16
with unprivileged account, 16
installation account, 16
installation directory, 13
installation procedure
abort, 28, 46
Berkeley DB, 46
Berkeley DB spooling server, 46
execution host, 41
master, 28
installation procedure (Continued)
spooling server, 46
interactive job, 104

J

job, 104
batch, 104
class, 104
monitoring, 83
parallel, 105
job array, 103
job class, 104
job_id, 83

M

manager, 105
master, as execution host, 61
master host, 17, 105
master installation procedure, 28
migration, 105

N

network services, 16
NIS, 12, 16, 82

O

operator, 105
Oracle database
accounting and reporting, 93
create tables, 93
create views, 94
environment variables, 93
grant permissions, 94
override policy, 105
owner, 105

P

parallel environment, 105

parallel job, 105
permission, file access, 16
policy, 105
  functional, 104
  override, 105
  share-based, 105
PostgreSQL database
  accounting and reporting, 89, 92
    create database, 92
    create database owner, 92
    create database tables, 92
    create tables, 92
    create user, 92
    create views, 92
priority, 105
project, 105

Q

`qconf -ah`, 45
qconf -ah, 59, 60
qconf -as, 46
qmaster spool directory, 18
`qstat`, monitor batch jobs, 83
queue, 20, 104

R

redirection
  stderr, 83
  stdout, 83
requirements
  hard, 104
  soft, 106
resource, 105
resource requirements
  hard, 104
  soft, 106
root account, 16
root directory, 13
`rsh`, 82

S

schedd spool directory, 18
scheduler, 21
  profile, 21
    high, 21
    max, 21
    normal, 21
services, 12, 82
services database, 16
`settings.csh`, 82
`settings.sh`, 82
`sge_execd`, 41
`SGE_EXECD_PORT`, 82
`sge_qmaster`, 28
`SGE_QMASTER_PORT`, 82
sge_schedd, 28
shadow master host, 17
shadow_masters file, 17
share, 105
share-based policy, 105
share-tree, 106
soft resource requirements, 106
spool directories, 17
spooling, classic, 18
spooling server, installation procedure, 46
stderr redirection, 83
stdout redirection, 83
submit host, 20, 106
  add, 46
Sun Web Console
  for accounting and reporting
    install, 99
suspension, 106
system requirements
  disk space, 13

T

ticket, 106

U

unprivileged account, 16
usage, 106
user, 106
user id, equivalent, 11
user names, 15

userset, 106