

Sun N1 Grid Engine 6.1 Installation Guide



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

Part No: 820-0697
May 2007

Copyright 2007 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

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.

This distribution may include materials developed by third parties.

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, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, N1 Java, 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 SunTM 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.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

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 2007 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, N1 Java et Solaris sont des marques 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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

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 N1 Grid Engine Software Interactively	27
Interactive Installation Overview	28
Performing an Installation	28
▼ How to Install the Master Host	28
▼ How to Install Execution Hosts	43
Registering Administration Hosts	50
Registering Submit Hosts	50
▼ How to Install the Berkeley DB Spooling Server	50
3 Automating the Installation Process	55
Automatic Installation Overview	55
Special Considerations	56
Using the <code>inst_sge</code> Utility and a Configuration Template	57
▼ How to Automate the Master Host Installation	57
Automating Other Installations Through a Configuration File	58
Configuration File Templates	59
Automatic Installation With Increased Security (CSP)	63
Automatic Uninstallation	64

Uninstalling Execution Hosts	64
Uninstalling the Master Host	65
Uninstalling the Shadow Host	65
Uninstalling the Shadow Host	65
Automatic Backup	65
Starting an Automatic Backup	66
Troubleshooting Automatic Installation and Uninstallation	66
4 Installing the Increased Security Features	69
Why Install the Increased Security Features?	69
Additional Setup Required	69
▼ How to Install a CSP-Secured System	70
▼ How to Generate Certificates and Private Keys for Users	73
▼ How to Renew Certificates	75
Checking Certificates	75
5 Upgrading From a Previous Release of N1 Grid Engine Software	77
About Upgrading the Software	77
▼ How to Upgrade to 6.1 Software Using Classic/Berkeley DB Spooling	78
▼ How to Upgrade to 6.1 Software Using Berkeley DB Server	78
▼ How to Upgrade the Software from 5.3 to 6.0 Update 2	79
6 Verifying the Installation	95
Verifying the Installation	95
▼ How to Verify That the Daemons Are Running on the Master Host	95
▼ How to Verify That the Daemons Are Running on the Execution Hosts	96
▼ How to Run Simple Commands	97
▼ How to Submit Test Jobs	98
7 Removing the Grid Engine Software	101
Removing the Software Interactively	101
▼ How to Remove the Software Interactively	101
Removing the Software Using the <code>inst_sge</code> Utility and a Configuration Template	102
▼ How to Remove the Software Interactively	102

8	Installing the Accounting and Reporting Console	105
	Setting Up the Database Software	105
	Using PostgreSQL Database Software with N1 Grid Engine	105
	▼ How to Start the Database Server	105
	▼ How to Set Up a PostgreSQL Database	108
	Using Oracle Database Software With N1 Grid Engine	109
	▼ How to Set Up an Oracle Database	109
	Using MySQL Database Software with N1 Grid Engine	109
	▼ How to Install MySQL Database on UNIX Systems	109
	▼ How to Set Up Accounting and Reporting Database in MySQL	111
	Installing the Accounting and Reporting Software	112
	▼ How to Set Up dbWriter	112
	▼ How to Install Sun Java Web Console	119
	▼ How to Install the Accounting and Reporting Console	120
	▼ How to Add Authorized ARCo Users	124
A	Microsoft Services For UNIX	125
	Overview	125
	Unsupported N1GE Functionality	126
	System Requirements	126
	Services for UNIX Installation	127
	Post SFU Installation Tasks	135
	Troubleshooting SFU	137
B	User Management For N1GE on Windows Hosts	139
	Overview	139
	Managing Users on Windows Hosts	139
	Windows User Example	140
	UNIX User Management	140
	Using N1GE in a Microsoft Windows Environment	141
	Registering Windows User Passwords	141
	Adding Windows Hosts to Existing N1 Grid Engine Systems	142
	▼ How to Add Windows Hosts Later	142

- C Other N1 Grid Engine Installation Issues 145**
 - Checking Linux Motif Libraries 145
 - ▼ Installing Linux Motif Libraries 145
 - Installing N1 Grid Engine on a System with IPMP 146
 - What is IP Multipathing? 146
 - Issues between IPMP and Grid Engine 146
 - Installing the Grid Engine Master Node With IPMP 147
 - Installing a Grid Engine on an Execution Host With IPMP 148
 - Enabling Administrative and Submit Hosts with IPMP 148

- Glossary 149**

- Index 153**

Preface

The *N1 Grid Engine 6 Installation Guide* describes the procedures for installing and upgrading the N1™ N1 Grid Engine 6.1 software.

This version of the manual supports N1 Grid Engine 6.1 software.

Who Should Use This Book

This manual is for engineers and system administrators, who need to install the N1 Grid Engine 6.1 software. Installation of the N1 Grid Engine 6.1 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 6.1 software, you should familiarize yourself with the concepts that are discussed in Chapter 1, “Introduction to the N1 Grid Engine 6.1 Software,” in *Sun N1 Grid Engine 6.1 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.

[Appendix A](#) explains how to prepare a Windows system for installation of N1GE6 by installing Microsoft's Services for UNIX.

[Appendix B](#), explains some differences that apply to using N1 Grid Engine with Windows hosts.

[Appendix C](#) describes some installation corner cases.

[Glossary](#) is a list of product-specific words and phrases and their definitions.

Related Books

Other books in the N1 Grid Engine 6.1 software documentation collection include:

- *Sun N1 Grid Engine 6.1 User's Guide*
- *Sun N1 Grid Engine 6.1 Administration Guide*
- *Sun N1 Grid Engine 6.1 Release Notes*

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation](http://www.sun.com/documentation/) (<http://www.sun.com/documentation/>)
- [Support](http://www.sun.com/support/) (<http://www.sun.com/support/>)
- [Training](http://www.sun.com/training/) (<http://www.sun.com/training/>)

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

Before You Install the Software

This chapter describes the steps that need to be taken before you install the N1™ Grid Engine 6.1 software (sometimes referred to as *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 6.1 software is to be a single cluster or a collection of sub-clusters, called *cells*. Cells enable 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.

Note – Hosts running Windows as operating system cannot have the host type master host or shadow master host.

- 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
<i>sge-root</i> directory	
Cell name	
Administrative user	
sge_qmaster port number	6444 is recommended
sge_execd port number	6445 is recommended
Master host	
Shadow master hosts	
Execution hosts	
Administration hosts	
Submit hosts	
Group ID range for jobs	
Spooling mechanism (Berkeley DB or Classic spooling)	

Parameter	Value
Berkeley DB server host (the master or another host)	
Berkeley DB spooling directory on the database server	
Scheduler tuning profile (Normal, High, Max)	
Installation method (interactive, secure, automated, or upgrade)	
If you are going to install N1 Grid Engine 6.1 on a Windows system, acquire and install Microsoft Services For UNIX. See Appendix A for more information.	
If you are going to install N1 Grid Engine 6.1 on a Windows system, create the required <i>Certificate Security Protocol (CSP)</i> certificates before installing N1GE. See “ How to Install a CSP-Secured System ” on page 70 for information about CSP certificates.	
Check Appendix C for applicability.	

Disk Space Requirements

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

- 40 Mbytes for the installation files 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

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

Note – You must satisfy several Windows-specific prerequisites before you can install N1 Grid Engine on hosts that are running the Windows operating system. You might need to install additional software on your computer which might require additional disk space. See [Appendix A](#)

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.

Note – Spool areas do not have to reside under *sge-root*. Actually this location may be avoided for efficiency reasons.

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
- 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 as shown in [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 *Sun N1 Grid Engine 6.1 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 might 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 `qcrsh`, `qtcsh`, or `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, for example called `sgeadmin`. With this setup, this particular user needs only read/write access to the shared *sge-root* file system.

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. 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.

Note – The name of the root user on Windows hosts depends on the system language of the Windows operating system. You can even change the name of the root user. The default name for many languages is the name Administrator.

If your Windows host is a member of a Windows domain, *only* the local Administrator is the root user. Neither the members of the Administrators group, nor the domain Administrator, nor a member of the Domain Admins group are the root user. See [Appendix B](#) for more information about users on Windows hosts.

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 6444/tcp
```

```
sge_execd 6445/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.

Note – Windows hosts cannot act as master hosts.

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.

Note – Windows hosts cannot act as master hosts.

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 might influence your decision are file format and data integrity. Writing to the Berkeley DB provides a greater level of data integrity than writing to a text file. However, a text file stores data in a format that you can read and edit. Normally, you do not 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. 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.

This configuration does not provide a High-Availability (HA) solution. For example, scripts of pending jobs are not spooled through BDB spool server and thus are not available for a shadow master.

With the introduction of NFS4 software available with the Solaris™ 10 operating system, you can use Berkeley DB spooling on a network file system. You could not use Berkeley DB spooling on previous NFS versions. This circumstance allows a shadow host installation spooled on Berkeley DB without setting up an additional Berkeley DB Spooling Server.

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 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 but this default value is not suitable when the file server is different from the master host.

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.

Consider the following 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 *Sun N1 Grid Engine 6.1 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 runtime 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 *Sun N1 Grid Engine 6.1 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 57.
- 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](#).

Check the Other Installation Issues Appendix

If you are installing N1 Grid Engine on a Linux system or on a system with IPMP, see [Appendix C](#) for important information.

Loading the Distribution Files on a Workstation

The N1 Grid Engine 6.1 software 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_1`. 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.

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.

You can use either the pkgadd method or the tar method.

■ 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
- SUNWsgeeax - Solaris (x64 platform) binaries for Solaris 10 Operating System
- SUNWsgeea - Accounting and Reporting Console (ARCo) packages for the Solaris and Linux 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_6u4
# pkgadd -d ./Common/Packages/SUNWsgeec
# pkgadd -d ./Docs/Packages/SUNWsgeed
```

Depending on the Solaris binary that you need, type one of the following commands:

```
# pkgadd -d ./Solaris_sparc/Packages/SUNWsgee
# pkgadd -d ./Solaris_sparc/Packages/SUNWsgeex
# pkgadd -d ./Solaris_x86/Packages/SUNWsgeei
# pkgadd -d ./Solaris_x64/Packages/SUNWsgeeax
```

■ **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/nlge-6_1-common.tar.gz	Architecture independent files

The tar files that contain platform-specific binaries use the naming convention of `nlge-6_1-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_1`.

Platform-Specific File	Platform
Solaris_sparc/tar/nlge-6_1-bin-solaris.tar.gz	Solaris (SPARC platform) 32-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
Solaris_sparc/tar/nlge-6_1-bin-solaris-sparcv9.tar.gz	Solaris (SPARC platform) 64-bit binaries for Solaris 7, Solaris 8, and Solaris 9 Operating Systems
Solaris_x86/tar/nlge-6_1-bin-solaris-i586.tar.gz	Solaris (x86 platform) binaries for Solaris 8, and Solaris 9 Operating Systems
Solaris_x64/tar/nlge-6_1-bin-solaris-x64.tar.gz	Solaris (x64 platform) 64-bit binaries for Solaris 10
Gemm/tar/nlge-6_1-gemm.tar.gz	Grid Engine Management Module, to be installed with SCS 2.2.1.
Windows/tar/nlge-6_1-bin-win32-x86.tar.gz	Microsoft Windows (x86 platform 32-bit binaries for Windows 2000, XP and Windows Server 2003
Linux24_i586/tar/nlge-6_1-bin-linux24-i586.tar.gz	Linux (x86 platform) binaries for the 2.4 kernel
Linux24_amd64/tar/nlge-6_1-bin-linux24-amd64.tar.gz	Linux (AMD platform) binaries for the 2.4 kernel
MacOSX/tar/nlge-6_1-bin-darwin.tar.gz	Apple Mac OS/X
HPUX11/tar/nlge-6_1-bin-hp11.tar.gz	Hewlett Packard HP-UX 11
Aix51/tar/nlge-6_1-bin-irix65.tar.gz	SGI Irix 6.5
Aix43/tar/nlge-6_1-bin-aix51.tar.gz	IBM AIX 5.1
Irix65/tar/nlge-6_1-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_1*.

```
% su
# cd sge-root
# gzip -dc basedir/Common/tar/nlge-6_1-common.tar.gz | tar xvpf -
# gzip -dc basedir/Solaris_sparc/tar/nlge-6_1-bin-solsparc32.tar.gz | tar xvpf -
# gzip -dc basedir/Solaris_sparc/tar/nlge-6_1-bin-solsparc64.tar.gz | tar xvpf -
# SGE_ROOT=sge-root; export SGE_ROOT
# util/setfileperm.sh $SGE_ROOT
```


Installing the N1 Grid Engine Software Interactively

This chapter discusses the steps that you must take to manually install the grid engine software.

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_iris.asc` file.

Also note that there are several prerequisites that you must satisfy for Windows systems before you can install N1 Grid Engine 6.1. See [Appendix A](#).

This chapter discusses the following topics:

- “Interactive Installation Overview” on page 28
- “Performing an Installation” on page 28

Note – This chapter does not cover the upgrade process or the installation of the Accounting and Reporting Module, ARCo. For information about upgrading, see [Chapter 5](#). For information about installing ARCo, see [Chapter 8](#).

Interactive Installation Overview

Note – The instructions in this section are for a *new* 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 43](#)
- [“Registering Administration Hosts” on page 50](#)
- [“Registering Submit Hosts” on page 50](#)
- [“How to Install the Berkeley DB Spooling Server” on page 50](#)

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](#).

Note – Windows hosts cannot act as master hosts.

1 Log in to the master host as root.

2 If the `$SGE_ROOT` environment variable is not set, set it by typing:

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

To confirm that you have set the `$SGE_ROOT` environment variable, type:

```
# echo $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/nlge6.

Checking \$SGE_ROOT directory

The Grid Engine root directory is:

```
$SGE_ROOT = /opt/nlge6
```

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/nlge6] >>

Your \$SGE_ROOT directory: /opt/nlge6

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      6444/tcp
sge_execd        6445/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 the Return key. You will see the following output:

```
sge_qmaster 6444
```

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 Type 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, type the cell name now.**
- **If you have decided not to use cells, press the Return key to continue.**

Using cell >default<.

Hit <RETURN> to continue >>

Press the Return key 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 18](#).

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 >sgadmin< 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 section in 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/nlge6/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 the Return key 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 The next question concerns Windows-based execution hosts.

If you do not plan to use Windows support, answer **No**. If you want Windows support, answer **Yes**.

If you answer yes, you will be asked some Windows-specific questions further on in the installation process. These questions will be marked as WINDOWS ONLY.

Windows Execution Host Support

Are you going to install Windows Execution Hosts? (y/n) [n]

11 Verify or set the correct file permissions.

If you used pkgadd or you know that the file permissions are correct, you should answer **Yes**. Answering **No** will direct the script to set the permissions for you as shown in the next step.

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] >> **y**

12 Set the correct file permissions.

- **WINDOWS ONLY – If you specified that you wanted Windows Execution Host support in the previous question, you should let the script set the file permissions for you. Answer No to the following question.**

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] >>

In some cases, eg: the binaries are stored on a NTFS or on any other
filesystem, which provides additional file permissions, the UNIX file
permissions can be wrong. In this case we would advise to verify and
to set the file permissions (enter: n) (y/n) [n] >>**n**

- **Verify and set file permissions.**

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 inaccessible 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 >>

13 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 equal, 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 the Return key 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

14 Press the Return key to continue.

Making directories

creating directory: default/common
creating directory: /opt/nlge6/default/spool/qmaster
creating directory: /opt/nlge6/default/spool/qmaster/job_scripts
Hit <RETURN> to continue >>

15 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] >>

- **To specify Berkeley DB spooling, press the Return key 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] >>

- **To 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 the Return key 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 50.](#)****d. After you have installed the software on the spooling server, return to the master installation window, and press the Return key to continue.****e. Type 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. Type the name of the spooling directory.

In the following example, /opt/nlge6/default/spooldb is the spooling directory.

Please enter the Database Directory now!

Default: [/opt/nlge6/default/spooldb] >>

Dumping bootstrapping information

Initializing spooling database

Hit <RETURN> to continue >>

■ If you do *not* want to use a Berkeley DB spooling server, type 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/nlge6/default/spool/spooldb] >>

Specify an alternate directory, or press the Return key to continue.

creating directory: /opt/nlge6/default/spool/spooldb

Dumping bootstrapping information

Initializing spooling database

Hit <RETURN> to continue >>

- **To specify classic spooling, type classic.**

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

Dumping bootstrapping information

Initializing spooling database

Hit <RETURN> to continue >>

16 Type a group ID range

For more information, see [“Group IDs” on page 20](#).

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 >>

17 Verify the spooling directory for the execution daemon.

For information on spooling, see [“Spool Directories Under the Root Directory” on page 18](#).

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/nlge6/default/spool] >>

18 Type the email address of the user who should receive problem reports.

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

Grid Engine cluster configuration (continued)

<administrator_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**

19 Verify the configuration parameters.

The following parameters for the cluster configuration were configured:

execd_spool_dir	/opt/nlge6/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< qtask 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 >>

20 WINDOWS-ONLY – If you specified that you want Windows support, you are asked to create Certificate Security Protocol (CSP) certificates.

Read “[How to Install a CSP-Secured System](#)” on page 70 for information about CSP certificates before you continue.

21 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 >>

...

22 WINDOWS-ONLY – Add the Windows Administrator name to the SGE manager list.

Windows Administrator Name

For a later execution host installation it is recommended to add the
Windows Administrator name to the SGE manager list

Please, enter the Windows Administrator name [Default: Administrator] >>

23 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 separated 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 >>

24 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 preferred 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.

- 25 WINDOWS-ONLY – If you are using CSP mode, copy the certificate files to each execution host.**
You can use a script to perform this function.

Tip – To use this functionality without being asked for a password, the root user should use `rsh` or `ssh` to access the execution hosts.

Should the script try to copy the cert files, for you, to each
execution host? (y/n) [y] >>

- 26 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 about 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 95](#).

▼ 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](#) and shared the common directory.

You must satisfy several prerequisites before you can install N1 Grid Engine execution hosts with Windows operating systems. You might have to install additional software on your computer. See [Appendix A](#).

On Microsoft Windows machines, additional steps are necessary before you can continue with the execution host installation. Follow the steps described in “[How to Install a CSP-Secured System](#)” on page 70-Steps 6a, 6b and 6c.

Note – If you are using Microsoft Windows machines, after the installation, each user has to register their Windows password with N1 Grid Engine using the `sgpasswd` client application. (See Appendix B for more information.)

- 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 `sgc-root` or use a network installation directory.**
- 3 If the `$SGE_ROOT` environment variable is not set, set it by typing:**

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

To confirm that you have set the `$SGE_ROOT` environment variable, type:

```
# echo $SGE_ROOT
```
- 4 Change directory (`cd`) to the installation directory, `sgc-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 successful installation you need a running Grid Engine qmaster. It is also necessary 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/nlge6`.

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

The Grid Engine root directory is:

```
$SGE_ROOT = /opt/nlge6
```

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/nlge6]` `>>`

Your `$SGE_ROOT` directory: `/opt/nlge6`

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

8 Type 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 type the cell names now.**
- **If you have decided not to use cells, then press the Return key to continue.**

Using cell >default<.

Hit <RETURN> to continue >>

Press <RETURN> to continue.

9 The install script checks to see whether the admin user already exists.

If the admin user already exists, the script continues uninterrupted. If the admin user does not exist, the script shows the following screen where you must supply a password for the admin user. After the admin user is created, press the Return key to continue with the installation.

Local Admin User

The local admin user sgeadmin, does not exist!

The script tries to create the admin user.

Please enter a password for your admin user >>

Creating admin user sgeadmin, now ...

Admin user created, hit <ENTER> to continue!

10 Press the Return key 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 >>

11 Specify whether you want to use a local spool directory.

For information on spooling, see [“Spool Directories Under the Root Directory” on page 18.](#)

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 configure a local spool directory for this host.

ATTENTION: The local spool directory doesn’t have to be located on a local drive. It is specific to the <local> host and can be located on network drives, too. But for performance reasons, spooling to a local drive is recommended.

FOR WINDOWS USER: On Windows systems the local spool directory MUST be set to a local harddisk directory.

Installing an execd without local spool directory makes the host unuseable. Local spooling on local harddisk is mandatory for Windows systems.

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 >host1< created.
```

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

12 Specify whether you want execd to start automatically at boot time.

You might not want to install the startup script if you are installing a test cluster or you would 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 >>
```

a. WINDOWS ONLY – Choose whether to display the GUI for Windows jobs.

An N1 Grid Engine Helper Service is included with the N1 Grid Engine 6.1 distribution. This service enables Windows jobs to display a GUI on the visible desktop of the execution host. The visible desktop is either the desktop of the user currently logged in on the execution host or the desktop of the next user who will log in. It is not the log in screen.

The Helper Service is a independent component loosely coupled with the execution daemon. The startup of the Helper Service is plugged in the Services dialog box in the Windows control panel. You can install only one Helper Service per host. There can be only one execution daemon installed per Helper Server.

The installation script asks during the installation of a execution host whether you want to see the GUI of Windows jobs.

SGE Windows Helper Service Installation

```
-----
If you're going to run Windows job's using GUI support, you have
to install the Windows Helper Service
Do you want to install the Windows Helper Service? (y/n) [n] >>
```

b. Start the execution daemon.

Grid Engine execution daemon startup

```
-----
Starting execution daemon. Please wait ...
starting sge_execd
```

Hit <RETURN> to continue >>

13 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.

14 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 about 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 96.](#)

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 about how to complete this task using QMON, see “Configuring Administration Hosts With QMON” in *Sun N1 Grid Engine 6.1 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 about how to complete this task using QMON, see “Configuring Submit Hosts With QMON” in *Sun N1 Grid Engine 6.1 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 *Sun N1 Grid Engine 6.1 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.](#)

1 Log in to the spooling server host as root.

2 If the \$SGE_ROOT environment variable is not set, set it by typing:

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

To confirm that you have set the \$SGE_ROOT environment variable, type:

```
# echo $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/nlge6.

Checking \$SGE_ROOT directory

The Grid Engine root directory is:

```
$SGE_ROOT = /opt/nlge6
```

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/nlge6] >>

Your \$SGE_ROOT directory: /opt/nlge6

Hit <RETURN> to continue >>

7 Type 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 an RPC Client/Server mechanism!
 In this case, qmaster will
 contact an 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 Type the directory path of your spooling directory.

You might need to change this path if this directory is NFS mounted, or if you do not have write permissions to this directory.

Enter the database directory
 or hit <RETURN> to use default [/opt/nlge6/default//spooldb] >>

creating directory: /opt/nlge6/default//spooldb

11 Start the RPC server.

Now we have to startup the rc script
 >/opt/nlge6/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/nlge6/default//spooldb

Please remember these values, during Qmaster installation
 you will be asked for them! 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**

Automating the Installation Process

This chapter describes how you can automate the software installation process for grid engine software.

- “Automatic Installation Overview” on page 55
- “Using the `inst_sge` Utility and a Configuration Template” on page 57
- “Automatic Installation With Increased Security (CSP)” on page 63
- “Automatic Uninstallation” on page 64
- “Automatic Backup” on page 65
- “Troubleshooting Automatic Installation and Uninstallation” on page 66

Automatic Installation Overview

You can use the `sge-root/inst_sge` utility to install and uninstall N1 Grid Engine master hosts, execution hosts, shadow host, and Berkeley DB spooling server hosts. You can also use this utility to backup automatically the N1 Grid Engine configuration and accounting data. You can use the `inst_sge` utility in interactive mode to supplant any of the commands that were described in [Chapter 2](#).

Note – The use of a Berkeley DB Spooling Server host does *not* provide high availability. In addition, the Berkeley DB Spooling Server has no authentication mechanism, and should only be used on a closed network with fully trusted users.

To simplify automatic installation and backup processes, use the configuration templates that are located in the `sge-root/util/install_modules` directory.

The automatic installation requires no user interaction. No messages are displayed on the terminal during the installation. When the installation finishes, a message indicates where the installation log file resides. The name of the installation log file is of the form `install_hostname_timestamp.log`. Normally, you can find information about errors during

installation in this file. In case of serious errors, the installation script might not be able to move the log file into the spool directory. In this situation, the log file is placed in the /tmp directory.

Special Considerations

The first step in performing an automatic installation is to set up a configuration file. You can find configuration file templates in the *sge-root/Util/install_modules* directory. Consider the following as you plan your automatic installation:

- To use automatic installation on remote hosts, the root user must be able to access those hosts through rsh or ssh without supplying a password.
- For local spooling, that is, spooling on the master host, no special configuration is needed. However, the directory where the spooling occurs must *not* be on an NFS version 3 volume. You *may* use an NFS version 4 volume for local spooling.
- To run the Berkeley DB spooling server on a host other than the master host, you must install and configure RPC services on this separate host.

To perform this step manually before you start the automatic installation, use the following command:

```
./inst_sge -db
```

You can also use the following command to install automatically the Berkeley DB Spooling Server:

```
% ./inst_sge -db -m -x -auto <full-path-to-configuration-file>
```

This command checks the SPOOLING_SERVER entry within the configuration file and starts the Berkeley DB installation on the server host.

Note – If you start the automatic installation on the master host, the entire cluster can be installed with one command. The automatic installation script accesses the remote hosts through rsh or ssh and starts the installation remotely. This process requires a well-configured configuration file, which each host must be able to read. That file should be installed on each host or shared through NFS.

Using the `inst_sge` Utility and a Configuration Template

To automate system installation, use the `inst_sge` utility in combination with a configuration file.

Note – You cannot use the auto installation procedure to install remotely a Windows execution host. You must run the auto installation procedure directly on the Windows execution host.

▼ How to Automate the Master Host Installation

Before You Begin You need to complete the planning process as outlined in [“Plan the Installation” on page 11](#).

In addition, you 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 cannot use this method of installation.

1 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
```

2 Edit your configuration template, using the values from the worksheet you completed in [“Plan the Installation” on page 11](#).

The configuration file template includes liberal comments to help you decide where appropriate information belongs. See [“Configuration File Templates” on page 59](#).

3 Log in as `root` on the system that you want to be the N1 Grid Engine master host.

4 Create the `sge-root` directory.

The `sge-root` directory is the root directory of the N1 Grid Engine software hierarchy.

5 Go to the `sge-root` directory and start the installation.

```
# cd sge-root
```

```
# ./inst_sge -m -auto full-path-to-configuration-file
```

The `-m` option starts the master host installation and installs the master daemon on the local machine. In addition, the `-auto` option sets up any remote hosts, as specified in the configuration file.

Note – You cannot install remotely a master host. You must always install a master host locally.

To prevent data loss or destroying already installed clusters, the automatic installation terminates if the configured `$SGE_CELL` directory or the configured Berkeley DB spooling directory already exists. If the installation terminates, the script displays the reason for the termination on the screen.

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`.

Tip – You can also combine options if you want to perform multiple installations with one command. For example, the following command installs the master daemon on the local machine and installs all execution hosts that are configured in the configuration file:

```
./inst_sge -m -x -auto full-path-to-configuration-file
```

Wait for notification that the installation has completed.

When the automatic installation exits successfully, it displays a message similar to the following:

```
Install log can be found in: /opt/nlge61/spool/install_myhost_30mar2007_090152.log
```

The installation log file includes any script or error messages that were generated during installation. If the `qmaster_spooling_dir` directory exists, the log files will be in that directory. If the directory does not exist, the log files will be in the `/tmp` directory.

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 ("").

Automating Other Installations Through a Configuration File

In addition to installing the master host, you can perform a variety of other automatic installations using a similar process. The actual form of the `inst_sge` command differs slightly, and different sections of the configuration file apply. This section provides some examples.

- To install a shadow host, use the following form of the command:

```
inst_sge -sm -auto full-path-to-configuration-file
```

Tip – To install more than one shadow host, enter the host names in the `<SHADOW_HOST>` parameter section within the configuration file.

- You can install a separate execution host installation if the master host was installed without identified compute hosts or if you need to add additional compute hosts. For the execution host installation, you also need to have a configuration file.

To install all configured execution hosts, use the following form of the command:

```
inst_sge -x -auto full-path-to-configuration-file
```

- To install the Berkeley database server, use the following form of the command:

```
inst_sge -db -auto full-path-to-configuration-file
```

Configuration File Templates

Configuration file templates are located in the `sge-root/util/install_modules` directory.

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/nlge61"

# 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="6444"

# 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="6445"

# 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"

# ADMIN_USER, if you want to use a different admin user than the owner,
# of SGE_ROOT, you have to enter the user name, here
```

EXAMPLE 3-1 Example Configuration File (Continued)

```
# Leaving this blank, the owner of the SGE_ROOT dir will be used as
admin user
ADMIN_USER=""

# The dir, where qmaster spools this parts, which are not spooled by DB
#(mandatory for qmaster installation)
QMASTER_SPOOL_DIR="/opt/nlge61/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/nlge61/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 method 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/nlge61/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
```

EXAMPLE 3-1 Example Configuration File (Continued)

```
# 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 prompting 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 ordinary file to 644.
SET_FILE_PERMS="true"
```

EXAMPLE 3-1 Example Configuration File *(Continued)*

```
# 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"

# This is a Windows specific part of the auto isntallation template
# If you going to install windows executions hosts, you have to enable the
# windows support. To do this, please set the WINDOWS_SUPPORT variable
# to "true". ("false" is disabled)
# (mandatory for qmaster installation, by default WINDOWS_SUPPORT is
# disabled)
WINDOWS_SUPPORT="false"

# Enabling the WINDOWS_SUPPORT, recommends the following parameter.
# The WIN_ADMIN_NAME will be added to the list of SGE managers.
# Without adding the WIN_ADMIN_NAME the execution host installation
# won't install correctly.
# WIN_ADMIN_NAME is set to "Administrator" which is default on most
# Windows systems. In some cases the WIN_ADMIN_NAME can be prefixed with
# the windows domain name (eg. DOMAIN+Administrator)
# (mandatory for qmaster installation)
WIN_ADMIN_NAME="Administrator"

# This parameter set the number of parallel installation processes.
# The prevent a system overload, or exeeding the number of open file
# descriptors the user can limit the number of parallel install processes.
# eg. set PAR_EXECD_INST_COUNT="20", maximum 20 parallel execd are installed.
PAR_EXECD_INST_COUNT="20"
```

Automatic Installation With Increased Security (CSP)

The automatic installation also supports the *Certificate Security Protocol (CSP)* mode described in [Chapter 4](#). To use the CSP security mode, you must fill out the CSP parameters of the template files. The parameters are as follows:

```
# This section is used for csp installation mode.
# CSP_RECREATE recreates the certs on each installation, if true.
# In case of false, the certs will be created, if not existing.
# Existing certs won't be overwritten. (mandatory for csp install)
CSP_RECREATE="true"

# The created certs won't be copied, if this option is set to false
# If true, the script tries to copy the generated certs. This
# requires passwordless ssh/rsh access for user root to the
# execution hosts
CSP_COPY_CERTS="false"

# csp information, your country code (only 2 characters)
# (mandatory for csp install)
CSP_COUNTRY_CODE="DE"

# your state (mandatory for csp install)
CSP_STATE="Germany"

# your location, eg. the building (mandatory for csp install)
CSP_LOCATION="Building"

# your organisation (mandatory for csp install)
CSP_ORGA="Organisation"

# your organisation unit (mandatory for csp install)
CSP_ORGA_UNIT="Organisation_unit"

# your email (mandatory for csp install)
CSP_MAIL_ADDRESS="name@yourdomain.com"
```

To start the installation, type the following command:

```
inst_sge -m -csp -auto template-file-name
```

Note – Certificates are created during the installation process. These certificates have to be copied to each host of the installed cluster. The installation process can do this for you; however, you need to perform the following steps to allow the installation process appropriate permissions to copy the certificates:

1. Use rsh/rcp or ssh/scp on each host.
 2. Provide the root user with access to each host over ssh or rsh, without entering a password.
-

Automatic Uninstallation

You can also uninstall hosts automatically.

Note – Uninstall all compute hosts before you uninstall the master host. If you uninstall the master host first, you have to uninstall all execution hosts manually.

To ensure that you have a clean environment, always source the `$SGE_ROOT/$SGE_CELL/common/settings.csh` file before proceeding.

Uninstalling Execution Hosts

During the execution host uninstallation, all configuration information for the targeted hosts is deleted. The uninstallation attempts to stop the exec hosts in a graceful manner. First, the queue instances associated with the target host of the uninstallation will be disabled, so that new jobs will not be started. Then, in sequence, the following actions are done on each running jobs: checkpoint the job; reschedule the job; do forced rescheduling of the job. At this point, the queue instance will be empty, and the execution daemon will be shut down, then the configuration, global spool directory or local spool directory will be removed.

The configuration file template has a section for identifying hosts that can be uninstalled automatically. Look for this section:

```
# Remove this execution hosts in automatic mode
EXEC_HOST_LIST_RM="host1 host2 host3 host4"
```

Every host in the `EXEC_HOST_LIST_RM` list will be automatically removed from the cluster.

To start the automatic uninstallation of execution hosts, type the following command:

```
% ./inst_sge -ux -auto full-path-to-configuration-file
```

Uninstalling the Master Host

The master host uninstallation removes all of the N1 Grid Engine configuration files. After the uninstallation procedure completes, only the binary files remain. If you think that you will need the configuration information after the uninstallation, perform a backup of the master host. The master host uninstallation supports both interactive and automatic mode.

To start the automatic uninstallation of the master host, type the following command:

```
% ./inst_sge -um -auto full-path-to-configuration-file
```

This command performs the same procedure as in interactive mode, except the user is not prompted for confirmation of any steps and all terminal output is suppressed. Once the uninstall process is started, it cannot be stopped.

Uninstalling the Shadow Host

To start the automatic uninstallation of the shadow host, type the following command:

```
% ./inst_sge -usm -auto full-path-to-configuration-file
```

Uninstalling the Shadow Host

To start the automatic uninstallation of the shadow host, type the following command:

```
% ./inst_sge -usm -auto full-path-to-configuration-file
```

Automatic Backup

The automatic backup procedure backs up configuration and accounting data in much the same way as the interactive backup procedure. You can run the automatic backup procedure as a cron job if you want to schedule unattended or periodic backups. The automatic backup requires a configuration file, for which a template is located in the `sge_root/util/install_modules/backup_template.conf` file.

Comments within the configuration file template indicate what values to use for your environment.

Starting an Automatic Backup

After you set up the configuration file, type the following command to start the automatic backup:

```
% ./inst_sge -bup -auto full-path-to-configuration-file
```

To prevent overwriting existing backup files, a date/time combination is added to the end of the backup directory name that is specified in the configuration file.

EXAMPLE 3-2 Backup Configuration File

```
#-----
# Autobackup Configuration File Template
#-----
# Please, enter your SGE_ROOT here (mandatory)
SGE_ROOT="/opt/gridengine"
# Please, enter your SGE_CELL here (mandatory)
SGE_CELL="default"
# Please, enter your Backup Directory here
# After backup you will find your backup files here (mandatory)
# The autobackup will add a time /date combination to this dirname
# to prevent an overwriting!
BACKUP_DIR="/opt/backups/ge_backup"
# Please, enter true to get a tar/gz package
# and false to copy the files only (mandatory)
TAR="true"
# Please, enter the backup file name here. (mandatory)
BACKUP_FILE="backup.tar"
```

Troubleshooting Automatic Installation and Uninstallation

The following errors might be encountered during auto-installation.

Problem: If the *sge_cell* directory exists, the installation terminates to avoid overwriting a previous installation.

Solution: Remove or rename the directory.

Problem: If the Berkeley database spooling directory exists, the installation terminates to avoid overwriting a previous installation.

Solution: This directory must be removed or renamed in order to proceed. Make sure that the ADMINUSER has permissions to write into the location where the Berkeley database spooling directory is located. The ADMINUSER will be the owner of the Berkeley database spooling directory.

Problem: The execution host installation appears to succeed, but the execution daemon is not started, or no load values are shown.

Solution: Verify that user root is allowed to rsh or ssh to the other host, without entering a password.

If your network does not allow user root to have permissions to connect to other hosts through rsh or ssh without asking for a password, the automatic installation will not work remotely. In this case, log in to the host and use the following command to start the automatic installation locally on each host:

```
% ./inst_sge -x -noremove -auto /tmp/install_config_file.conf
```


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 43](#), and [“Registering Administration Hosts” on page 50](#).

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.

- 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 | portSGE_QMASTER_PORT}/cell/private`.
- User keys and certificates are stored in `/var/sgeCA/{sge_service | portSGE_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 SGE 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 43](#).

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 are 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 73](#).

- 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
```

Note – On a Windows execution host, the `tar` utility cannot restore the ownerships and permissions. In this case, the Administrator must set the ownerships and permissions manually.

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 are 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 73](#).

▼ 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.

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

```

▼ How to Renew Certificates

- 1 **Change to `$SGE_ROOT` and become `root` on the master host (assuming that `$SGE_CELL` is the default):**

```
# tcsh
# source $SGE_ROOT/default/settings.csh
```
- 2 **Edit `$SGE_ROOT/util/sgeCA/renew_all_certs.csh`, and change the number of days that the certificates are valid:**

```
# extend the validity of the CA certificate by
set CADAYS = 365
# extend the validity of the daemon certificate by
set DAEMONDAY = 365
# extend the validity of the user certificate by
set USERDAY = 365
```
- 3 **Run the changed script (the default for all extension times is 365 days from the day the script is run).**

```
# util/sgeCA/renew_all_certs.csh
```
- 4 **Replace the old certificates against the new ones on all hosts that installed them locally (that is, under `/var/sgeCA/...`, see the execution daemon installation).**
- 5 **If users have copied certificates and keys to `$HOME/.sge`, they have to repeat `$SGE_ROOT/util/sgeCA/sge_ca -copy` to have access to the renewed certificates.**

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
```

Upgrading From a Previous Release of N1 Grid Engine Software

This chapter describes the steps necessary to upgrade your existing N1 Grid Engine software to N1 Grid Engine 6.1 software.

Note – The upgrade procedure is only able to upgrade your software from version 6.0 update 2 or higher. If you are running an older version of the N1 Grid Engine software, such as 5.3 or 6.0, you should upgrade to version 6.0 update 10 and then upgrade again to version 6.1.

About Upgrading the Software

The upgrade procedure is non destructive. The upgrade procedure installs the N1 Grid Engine 6.1 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.

Note – The `LD_LIBRARY_PATH` variable is not set in N1 Grid Engine 6.1 software. Remove the existing `LD_LIBRARY_PATH` settings from 6.0 before you start a 6.1 installation.

Before you begin the upgrade process, make sure that you source the existing `sge-root/sge-cell/common/settings.sh` or `sge-root/sge-cell/common/settings.csh` file.

Two upgrade scenarios are available:

- Upgrading to 6.1 using classic/Berkeley DB spooling
- Upgrading to 6.1 using Berkeley DB server spooling

▼ How to Upgrade to 6.1 Software Using Classic/Berkeley DB Spooling

1 Back up your cluster.

2 Shut down the cluster.

Type the following command:

```
% qconf -ke all -ks -km
```

3 Extract the new binaries to the \$SGE_ROOT directory.

Replace the old files with the new files.

4 Initiate the upgrade procedure.

Type the following command:

```
% ./inst_sge -upd
```

Follow the instructions and answer the questions presented to you. This update will re-create the rc files and settings files.

5 Restart the cluster.

Use the sgemaster and sgeexecd rc files.

Upgrade is complete.

▼ How to Upgrade to 6.1 Software Using Berkeley DB Server

64-bit SPARC only – For 64-bit Solaris systems, a 64-bit Berkeley DB server was not available. The 32-bit server was used instead. For doing a backup or a restore, the 32-bit db_dump and db_load binaries are recommended. The Solaris 32-bit packages have to be installed also. For all N1 Grid Engine releases and patches earlier than 6.0 Update 9, use the ./inst_sge -bup command for doing a backup. In N1 Grid Engine 6.0 Update 10, the backup still looks for the libdb-4.2.so file, which is not available. This leads to a backup failure. In this case, type the following command to enable the backup to function:

```
% touch $SGE_ROOT/lib/sol-sparc/libdb-4.2.so
```

1 Back up your cluster.

2 Shut down the cluster.

Type the following command:

```
% qconf -ke all -ks -km
```

3 Shut down the Berkeley DB server.

Type the following command:

```
% sgebdb stop
```

4 Extract the new binaries to the \$SGE_ROOT directory.

Replace the old files with the new files.

5 Restore the cluster configuration.**6 Initiate the upgrade procedure.**

Type the following command:

```
% ./inst_sge -upd
```

Follow the instructions and answer the questions presented to you. This update will recreate the rc-files and settings files.

7 Restart the cluster.

Use the sgemaster and sgeexecd rc files.

8 Restart the Berkeley DB server, master host, and execution hosts.

Upgrade is complete.

▼ How to Upgrade the Software from 5.3 to 6.0 Update 2

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 6.0 software concurrently with your older version of grid engine software, it is better to run the upgrade procedure when there are no running jobs.

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/nlge6`, and the administrative user is `sgeadmin`.

Grid Engine admin user account

The current directory

`/opt/nlge6`

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/nlge6`.

Checking `$SGE_ROOT` directory

The Grid Engine root directory is:

`$SGE_ROOT = /opt/nlge6`

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/nlge6]` >>

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 18](#).

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 >sgadmin< 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 section in 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/nlge6/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 inaccessible 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 equal, 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/nlge6/default/spool/qmaster
creating directory: /opt/nlge6/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 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 50.](#)

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 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/nlge6/default/spool/spooldb] >>

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

creating directory: /opt/nlge6/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 20](#).

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 18.

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/nlge6/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)

<administrator_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 6.1 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 43](#).

- 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 57](#).

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.

If necessary, you can use `qconf -mc` to reconfigure your complexes. For details, see Chapter 3, “Configuring Complex Resource Attributes,” in *Sun N1 Grid Engine 6.1 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 *Sun N1 Grid Engine 6.1 Administration Guide*.

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

1 Log in to the master host.

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

2 Verify that the daemons 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, restart the daemons.

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 97](#).

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

- 1 Log in to the execution hosts on which you ran the execution host installation procedure.
- 2 Verify that the daemons 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, restart the daemon.

```
# 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 97](#).

▼ 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.

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 *Sun N1 Grid Engine 6.1 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 these entries do not exist, add an entry to the machine's services database, giving it the same value as is configured on the master host.

b. 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.

- 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 any additional lines contain 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 *Sun N1 Grid Engine 6.1 User's Guide*.

5 After the job finishes executing, check your home directory for the redirected `stdout/stderr` files *script-name.ejob-id* and *script-name.ojob-id*.

job-id is a consecutive unique integer number assigned to each job.

In case of problems, see Chapter 9, “Fine Tuning, Error Messages, and Troubleshooting,” in *Sun N1 Grid Engine 6.1 Administration Guide*.

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 101
- “Removing the Software Using the `inst_sge` Utility and a Configuration Template” on page 102

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

- 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

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, 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, 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, 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
```


Installing the Accounting and Reporting Console

This chapter describes the steps to install the Accounting and Reporting console software (ARCo):

- “Setting Up the Database Software” on page 105
- “Installing the Accounting and Reporting Software” on page 112

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, a MySQL database, or an Oracle database.

Using PostgreSQL Database Software with N1 Grid Engine

Detailed information on the PostgreSQL database software can be found in the [Postgres documentation \(http://www.postgresql.org/docs/7.4/static/index.html\)](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.

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`.

```
tcpi_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 108](#).

▼ How to Set Up a PostgreSQL Database

- 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 a database user for reading the database.

If you do not use the arco_read user, you 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
```

Using Oracle Database Software With N1 Grid Engine

▼ How to Set Up an Oracle Database

- 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.**

Using MySQL Database Software with N1 Grid Engine

The Accounting and Reporting Console uses views. As a result, the console supports MySQL database version 5.0.x and higher. For more information on the MySQL database software, see the [MySQL documentation \(http://dev.mysql.com/doc/index.html\)](http://dev.mysql.com/doc/index.html).

Case Sensitivity in MySQL Database

In MySQL, databases correspond to directories within the data directory. Each table within a database corresponds to at least one file within the database directory. Consequently, the case sensitivity of the underlying operating system determines the case sensitivity of database and table names. Therefore, database and table names are case sensitive in most varieties of UNIX, and not case sensitive in Windows.

▼ How to Install MySQL Database on UNIX Systems

- 1 **Download the appropriate MySQL software for your system from <http://www.mysql.com>.**

The standard installation directory for UNIX systems is `/usr/local/mysql`. If you install the software into a different directory, you have to change the settings for the scripts provided in the package.

Note – ARCo is a Java web-based application and needs the Java™ DataBase connectivity (JDBC™) driver for converting JDBC calls into the network protocol used by the MySQL database. You can download the JDBC driver from <http://www.mysql.com/products/connector>.

- 2 **Create a symbolic link from the installation directory to `mysql`.**

```
# ln -s $installation_directory/mysql-standard-5.0.26-solaris10-i386 mysql
```

The `mysql` directory contains several files and subdirectories:

- The `bin` subdirectory contains client programs and the server.
 - The MySQL server and server startup scripts:
 - `mysqld` is the MySQL server.
 - `mysqld_safe`, `mysql.server`, and `mysqld_multi` are server startup scripts.
 - `mysql_install_db` initializes the data directory and the initial databases.
 - `mysqlmanager` is the MySQL instance manager.
 - Client programs that access the server:
 - `mysql` is a command-line client for executing SQL statements interactively or in batch mode.
 - `mysqladmin` is an administrative client.
 - `mysqlcheck` performs table maintenance operations.
 - `mysqldump` and `mysqlhotcopy` make database backups.
 - `mysqlimport` imports data files.
 - `mysqlshow` displays information about databases and tables.
 - Utility programs that operate independently of the server:
 - `myisamchk` performs table maintenance operations.
 - `myisampack` produces compressed, read-only tables.
 - `mysqlbinlog` is a tool for processing binary log files.
 - `pererror` displays the meaning of error codes.
- The `scripts` subdirectory contains the `mysql_install_db` script used to initialize the `mysql` database containing the grant tables that store the server access permissions.

3 Add a login user and group for `mysqld`.

```
# groupadd mysql
# useradd -g mysql mysql
```

4 Change the ownership of program binaries to `root` and ownership of the data directory to the user that you use to run `mysqld`.

```
# chown -R root .
# chown -R mysql data
# chgrp -R mysql .
```

5 Create the MySQL grant tables.

```
# scripts/mysql_install_db --user=mysql
```

6 Start the MySQL server.

```
# bin/mysqld_safe --user=mysql &
```

7 Assign the root password.

```
# ./bin/mysqladmin -u root password 'new-password'
# ./bin/mysqladmin -u root -h ${hostname} password 'new-password'
```

More Information MySQL Installation Tips

- To start the MySQL server at boot time, copy support-files/mysql.server to /etc/init.d and link it to /etc/rc3.d/S99mysql and /etc/rc0.d/K01mysql. If MySQL is not installed in /usr/local/mysql., edit the file to change the basedir and datadir variables.
- Add the full pathname of this directory to your PATH environment variable so that your shell finds the MySQL programs properly.

▼ How to Set Up Accounting and Reporting Database in MySQL**1 Log in to the database as a superuser.**

```
# mysql -u root -p<password>
```

2 Create user and grant privileges for arco_read and arco_write.

```
mysql> GRANT ALL on *.* to 'arco_write'@'<hostname>' identified by '<password>' with GRANT OPTION;
mysql> GRANT ALL on *.* to 'arco_write'@'localhost' identified by '<password>' with GRANT OPTION;
mysql> GRANT ALL on *.* to 'arco_write'@'%' identified by '<password>' with GRANT OPTION;
```

```
mysql> GRANT SELECT,SHOW VIEW on <db_name>.* to 'arco_read'@'<hostname>' identified by '<password>';
mysql> GRANT SELECT,SHOW VIEW on <db_name>.* to 'arco_read'@'localhost' identified by '<password>';
mysql> GRANT SELECT,SHOW VIEW on <db_name>.* to 'arco_read'@'%' identified by '<password>';
```

3 Log in to the database as arco_write user.

```
# mysql -u arco_write -p<password>
```

4 Create the accounting and reporting database.

```
mysql> CREATE DATABASE ARCO;
```

5 Install the Accounting and Reporting Software.

Installing the Accounting and Reporting Software

▼ How to Set Up dbWriter

- 1 Extract the accounting and reporting software using either the `tar` method or the `pkgadd` method.

- **tar Method**

- # `cd sge-root`

- Type the following command as one string, with a space between the `-dc` and the path to the tar file.

```
# gunzip -dc
```

```
cdrom-mount-point/N1_Grid_Engine_6_1/ARCo/tar/nlge-6_0-arco-1_0.tar.gz | tar xvpf -
```

- **pkgadd Method**

- At the command prompt, type the following command and respond to the script questions.

- # `cd cdrom_mount_point/N1_Grid_Engine_6_1/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 *Sun N1 Grid Engine 6.1 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.

Note – The dbwriter software needs to be installed as user root.

```
# cd sge-root/dbwriter
% cd sge-root/dbwriter
% su
password:
# ./inst_dbwriter
# ./inst_dbwriter

Welcome to the GE installation
-----

GE 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 or cells.

```
Hit <RETURN> to continue >>

Checking $SGE_ROOT directory
-----

The GE root directory is:

    $SGE_ROOT = /opt/nlge6

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/nlge6] >>

Your $SGE_ROOT directory: /opt/nlge6

Hit <RETURN> to continue >>
```

GE cells

GE supports multiple cells.

If you are not planning to run multiple GE clusters or if you don't know yet what is a GE 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 GE commands.

Enter cell name [default] >>

Using cell >default<.

Hit <RETURN> to continue >>

6 Specify the location of your Java Software Development Kit.

Version 1.4.2 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/j2sdk1.4.2_06

7 Specify the connection parameters for the reporting database.

Setup your database connection parameters

Enter your database type (o = Oracle, p = PostgreSQL, m = MySQL) [] >> 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 Provide 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 Type 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 value should be the user name 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/nlge6/dbwriter/lib .....
found in /opt/nlge6/dbwriter/lib/pg73jdbc2.jar
```

OK, jdbc driver found

- **If no corresponding JDBC driver is found, you are asked to copy the JAR file of the database driver into the *\$ge-root/dbwriter/lib* directory.**

Note – The JDBC driver for PostgreSQL is provided with the installation. If you install the ARCo software on Oracle or MySQL, you should copy the appropriate JDBC driver into the \$SGE_ROOT/dbwriter/lib directory.

After you copy the JAR file to the correct location, press RETURN, and the search repeats.

```
Error: jdbc driver org.postgresql.Driver
      not found in any jar file of directory
      /opt/nlge6/dbwriter/lib
```

Please copy a jdbc driver for your database into
this directory!

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 setup 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 Database model initialization or upgrade

During the next step, the database tables are created or updated and the user `arco_read` is granted limited access to the database tables. The ARCo web application connects to the database using the user `arco_readuser`.

```
Query db version ... no sge tables found
New version of the database model is needed
```

```
Should the database model be upgraded to version 1? (y/n) [y] >>
```

The ARCo web application connects to the database with a user which has restricted access. The name of this database user is needed to grant the user access to the sge tables.

```
Please enter the name of this database user [arco_read] >>
Upgrade to database model version 1 ... Install version 6.0 (id=0) -----
Create table sge_job
Create index sge_job_idx0
Create index sge_job_idx1
create table sge_job_usage
Create table sge_job_log
Create table sge_job_request
Create table sge_job_request
Create table sge_queue
Create index sge_queue_idx0
Create table sge_queue_values
Create index sge_queue_values_idx0
Create table sge_host
Create index sge_host_idx0
Create table sge_host_values
Create index sge_host_values_idx0
Create table sge_department
Create index sge_department_idx0
Create table sge_department_values
Create index sge_department_values_idx0
Create table sge_project
Create index sge_project_idx0
Create table sge_project_values
Create index sge_project_values_idx0
Create table sge_user
Create table sge_user_values
Create index sge_user_values_idx0
Create table sge_group
Create index sge_group_idx0
Create table sge_group_values
Create index sge_group_values_idx0
```

```

Create table sge_share_log
Create view view_accounting
Create view view_job_times
Create view view_jobs_completed
Create view view_job_log
Create view view_department_values
Create view view_group_values
Create view view_host_values
Create view view_project_values
Create view view_queue_values
Create view view_user_values
committing changes
version 6.0 (id=0) successfully installed
Install version 6.0u1 (id=1) -----
Create table sge_version
Update view view_job_times
Update version table
committing changes
version 6.0u1 (id=1) successfully installed
OK

```

12 Specify how often the dbWriter program should check the N1 Grid Engine log files for new data.

Please enter the interval between two dbwriter runs in seconds [60] >>

13 Specify the spool directory for the dbwriter.

The dbwriter log files are stored in this directory. At startup the dbwriter process writes its pid in the file dbwriter.pid in this directory.

Please enter the path of the dbwriter spool directory [\$SGE_ROOT/\$SGE_CELL/spool/dbwriter]>>

14 Type the location of the file containing the rules for calculating derived values.

Based on the accounting and reporting 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 enables 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] >>

15 Set 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] >>

16 Verify the settings.

If you answer n, you are given the option to repeat the setup.

All parameters are now collected

```
SGE_ROOT=/opt/nlge6
SGE_CELL=default
JAVA_HOME=/usr/java (java version "1.4.2")
DB_URL=jdbc:postgresql://my-host.my-domain:5432/arco
DB_USER=arco_write
INTERVAL=60
REPORTING_FILE=/opt/nlge6/default/common/reporting
DERIVED_FILE=/opt/nlge6/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/nlge6/default/common

17 The installation script prompts you to choose whether to install a boot time startup script for the dbwriter:

dbwriter startup script

We can install the startup script that will
start dbwriter at machine boot (y/n) [y] >> **y**
Installing startup script /etc/rc2.d/S96sgedbwriter
Installation of dbWriter completed

18 Start the dbWriter program.

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

▼ How to Install Sun Java Web Console

- 1 Check whether Sun Java Web Console is already available on your system, as is usually the case for the Solaris 10 software and on newer Solaris 9 versions.

As root, you can check using the following command:

```
# smcwebserver -V
Version 2.2.1
```

ARCo for N1 Grid Engine 6.1 software requires Sun Java Web Console 2.2.x.

Note – ARCo for N1 Grid Engine 6.1 software does not support Sun Java Web Console 3.0 or later.

- 2 If you need to install the console, extract the web console package under a temporary directory.

```
# cd /tmp
# umask 022
# mkdir swc
# cd swc
```

Type the following command as one string, with a space between the xvf and the path to the tar file.

```
# tar xvf
cdrom_mount_point/N1_Grid_Engine_6_1/SunWebConsole/tar/swc_sparc_2.2.1.tar
```

- 3 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/sh -c "No Access User" noaccess
```

Note – On some operating systems, using `/bin/csh` as the default shell might cause the installation to fail. If ARCo will not install, and you see the error message `Server not started! No management applications registered`, try setting the default shell as `/bin/sh` for `noaccess`.

- 4 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
```

- 5 Run the Sun Java Web Console setup script.

```
# ./setup -n
<...>
```

Installation complete.

Starting Sun(TM) Web Console Version 2.2.1...

See /var/log/webconsole/console_debug_log for server logging information

The web console is installed but not started until after the ARCo console installation. Once the console is installed, you can use the following commands to stop, start, or restart the console at any time:

```
# /usr/sadm/bin/smcwebserver start
# /usr/sadm/bin/smcwebserver stop
# /usr/sadm/bin/smcwebserver restart
```

▼ How to Install the Accounting and Reporting Console

- 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  GE ARCo reporting module installation
```

```
-----
```

```
The installation will take approximately 5 minutes
```

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

The install script checks for the existence of the SGE_ROOT directory and SGE_CELL name. If you already created these parameters, the script continues.

- 3 **Enter the path to your Java Software Development Kit.**

Version 1.4.2 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/j2sdk1.4.2_06
```

- 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, the database user for accounting and reporting should have only 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/nlge/reporting/WEB-INF/lib ...
found in /opt/nlge/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

The name of the database schema depends on your database. For a PostgreSQL database, the database schema name should be public. For Oracle, the name of the database schema should be the same as the name of the user account which is used by the dbwriter (arco_write).

Please enter the name of the database schema [arco_write] >>

Searching for the jdbc driver oracle.jdbc.driver.OracleDriver

OK, jdbc driver found

Should the connection to the database be tested? (y/n) [y] >> y

Test db connection to 'jdbc:oracle:thin:@my-system:1521:arco' ... OK

Note – The JDBC driver should be also located in \$SGE_ROOT/reporting/WEB-INF/lib directory.

7 Enter the login names of users who are allowed to store the queries and results.

Configure users with write access

```
Users: sge-admin
Enter a login name of a user (Press enter to finish) >> user2
Users: sge-admin user2
Enter a login name of a user (Press enter to finish) >>
-----
```

Note – After installation, you can add or delete authorized users by editing the config.xml file. See the information at the end of this chapter.

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 2.2.1...

9 If a previous version of ARCo is installed, you will be asked to remove it:

GE reporting module already registered at Java Web Console

The GE reporting modules can only be installed
if no previous version is registered.

Should the GE reporting module com.sun.grid.arco_6.0-Maintrunc be unregistered? (y/n) [n] >> **y**

10 Install pre-defined queries.

If the query directory already does not exist, it will be created. The example queries will be installed in *sge-root/reporting/database/example_queries/queries*. Existing queries will not be replaced.

Directory /var/spool/arco does not exist, create it? (y/n) [y] >> **y**

Create directory /var/spool/arco

```

Create directory /var/spool/arco/queries
Copy examples queries into /var/spool/arco/queries
Copy query Accounting_per_Department.xml ... OK
Copy query Accounting_per_Project.xml ... OK
Copy query Accounting_per_User.xml ... OK
Copy query Average_Job_Turnaround_Time.xml ... OK
Copy query Average_Job_Wait_Time.xml ... OK
Copy query DBWriter_Performance.xml ... OK
Copy query Host_Load.xml ... OK
Copy query Job_Log.xml ... OK
Copy query Number_of_Jobs_completed.xml ... OK
Copy query Queue_Consumables.xml ... OK
Copy query Statistic_History.xml ... OK
Copy query Statistics.xml ... OK
Create directory /var/spool/arco/results

```

11 Confirm that reporting module is set up.

ARCo reporting module setup

Copying ARCo reporting file into /opt/nlge6/default/arco/reporting

Setting up ARCo reporting configuration file. After registration of the ARCo reporting module at the Java Web Console you can find this file at

```
/var/opt/webconsole/webapps/reporting/config.xml
```

Hit <RETURN> to continue >>

12 Confirm that reporting module is registered to the web console and the console starts.

Registering the GE reporting module in the Java Web Console

```
Registering com.sun.grid.arco_6.0-Maintrunc.
```

Hit <RETURN> to continue >>

Java Web Console startup

Starting Java Web Console ... OK

Hit <RETURN> to continue >>

GE ARCo reporting successfully installed

13 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 takes effect the next time the console is started or restarted. The possible log levels are `WARNING`, `INFO`, `FINE`, `FINER`, and `FINEST`.

14 Connect to the Sun Java Web Console by accessing the following URL in your browser. Replace *hostname* with the name of your master host.

```
https://hostname:6789
```

15 Login with your UNIX account.**16 Select N1 Grid Engine 6 Accounting and Reporting Console.**

▼ How to Add Authorized ARCo Users

During the installation of the ARCo reporting module, you are asked to enter a list of users who should have write permissions to the ARCo system. Only those users are allowed save modifications on ARCo.

1 Add users to the appropriate file.

Depending on the web console version you are using, the list of authorized users is stored in `/var/opt/webconsole/webapps/reporting/config.xml` or `/usr/share/webconsole/webapps/reporting/config.xml`.

2 After editing this file, restart the Sun Java Web Console:

```
# smcwebserver restart
```

Microsoft Services For UNIX

Overview

Microsoft Windows Services for UNIX (SFU) makes it possible to integrate some Windows operating systems into existing UNIX environments. It provides components that simplify network administration and user management across the UNIX and Windows platform.

SFU provides Interix a complete, high-performance UNIX environment that provides UNIX shells like `csh` or `ksh`, several hundred tools and utilities, and a complete set of development tools and libraries which make it possible to port and use your UNIX-based applications to the Interix subsystem.

Windows Services for UNIX provides:

- In combination with Sun N1 Grid Engine the ability to integrate Windows hosts into N1GE clusters. This means that the execution and client environment of N1GE can be used on Microsoft Windows hosts.
- Services and commands to access the network file system (NFS). This makes it possible to share files between the UNIX and Windows environment
- A user mapping service which might access account and password services on UNIX and Windows systems (PCNFS, NIS)
- "single sign-on" capability for the Windows and UNIX environment. Passwords can be synchronized and authentication credentials can be mapped between the UNIX and Windows operating systems.
- The ability to execute UNIX shell scripts and applications to run on Windows-based computers in full-featured UNIX environments.

The following information describes the necessary steps to accomplish the installation of SFU. This document describes:

- The unsupported N1GE functionality.
- The SFU system requirements.

- Information to consider before SFU installation.
- The location of documentation which troubleshoots some common post installation and operation problems which might occur with SFU.

Unsupported N1GE Functionality

The following parts/applications which are part of a standard N1GE installation on UNIX operating systems are not supported in a Microsoft Windows environment and therefore cannot be used on Windows Hosts:

- Master and Scheduler (`sge_qmaster`, `sge_shadowd`, `sge_schedd`)
- Graphical User Interface (`qmon`)
- DRMAA
- `qsh` client command

System Requirements

The following requirements apply to the installation of Microsoft Services For UNIX.

- The system must have Internet Explorer version 5.0 or later installed prior to running SFU setup.
- The minimum hard disk requirements for installing SFU depends on which components you are installing. The maximum disk space requirement to install SFU components is 360 MB. The minimum disk space required is 20 MB. SFU must be installed on a partition that is formatted with the NTFS file system.
- You cannot install SFU on a system running Microsoft Services for Network File System (for example, Microsoft Services for NFS is a component of Windows Storage Server 2003).
- You must install the latest Windows service pack before installing SFU and N1GE then continue to install additional Windows service packs as they become available.
- Data Execution Prevention (DEP) must be disabled. DEP is not compatible with some parts of SFU and might cause segmentation faults. See <http://support.microsoft.com/kb/875352> for more informations about DEP. To disable DEP, follow these steps:
 1. Right click on the "My computer" icon on the desktop.
 2. Select "Properties".
 3. In the Properties dialog box, select the "Advanced" tab and click the "Settings" button in the "Startup and Recovery" section.
 4. In the next dialog box, click the edit button to edit the boot command line of your Windows installation.
 5. Add `"/noexecute=alwaysoff"` or modify an existing `"/noexecute"` option.

You can find more details concerning SFU requirements at
<http://www.microsoft.com/windows/sfu/default.asp>
(<http://www.microsoft.com/windows/sfu/default.asp>)

Services for UNIX Installation

Microsoft's SFU is required to install N1GE successfully. You can download SFU from the <http://www.microsoft.com>. Search the site for "Windows Services for Unix" to find the current download information.


1. Get the SFU distribution media.

If you downloaded SFU, just execute the application to unzip the files into a directory. This directory must be located on a file system that has at least 480 MBytes free space.

2. Login to the Windows system with the Administrator account.
3. Start the setup.exe application which was unpacked previously.




4. Enter your name and organization.



5. Accept the license agreement for SFU.



6. Now you have to choose between the standard installation and custom installation. The standard installation is recommended.



If disk space is limited, you might also choose the custom installation path but make sure that at least following components will be installed:

Utilities -> Base Utilities


Interix Gnu components -> Interix GNU utilities

Remote connectivity components -> Telnet Server and Windows Remote Shell

If you intend to use NFS shared file systems then you also need: Authentication tools for

NFS -> User Mapping and Server for NFS Authentication.

7. Depending on the underlying Windows operating system you might be asked two questions concerning the SFU security settings. The following screen shows the recommended selections:



The following information is a basic description concerning these two options. If you need further information please consult Microsoft's SFU documentation.

Enable suid behavior for Interix programs — According to the POSIX standard, a file has permissions that include bits to set both a UID (`setuid`) and a GID (`setgid`) when the file is executed. If either or both bits are set on a file, and a process executes that file, the process gains the UID or GID of the file. When used carefully, this mechanism allows a non-privileged user to execute programs that run with the higher privileges of the filer's owner or group. When used incorrectly, however, this behavior can present security risks by allowing nonprivileged users to perform actions that should only be performed by an administrator. For this reason, Windows Services for UNIX Setup does not enable support for this mechanism by default.

You should enable support for `setuid` behavior because N1GE will be running programs that require this support. Even if you do not enable support for `setuid` behavior when installing Windows Services for UNIX, you can enable it later.

Changing Default Behavior to Case Sensitivity — You might be required to choose whether to change the default behavior of object names, such as file names, to being case sensitive. The choice you make will affect system security as well as how Windows Services for UNIX functions. With Microsoft Windows, the names of most objects (such as files and directories) are case preserving, but case insensitive. So, you cannot have two files in the same directory named `sample.txt` and `Sample.txt` because Windows regards the names as identical for the purposes of identifying files. However, the UNIX operating system is fully case sensitive. So, UNIX systems distinguish between object names when the only difference between those names is the case of the objectname characters. Therefore, `sample.txt` and `Sample.txt` could appear in the same directory and the UNIX system would distinguish between them when performing operations on the files. For example, the command `rm`

S*.txt would delete Sample.txt but not sample.txt. In order to implement typical UNIX behavior, the Server for NFS and the Interix subsystem are normally case sensitive when working with file names.

This behavior can present security issues, particularly for users who are accustomed to the caseinsensitive conventions of Windows. For example, a Trojan horse version of edit.exe named EDIT.EXE could be stored in the same directory as the original. If a user were to type edit at a Windows command prompt, the Trojan horse version (EDIT.EXE) could be executed instead of the standard version. If case sensitivity is enabled, Windows users should be made aware of this possibility.

With Windows XP(Professional) and the Windows Server 2003 family, the default behavior of subsystems other than the Win32 subsystem is to be case preserving but case insensitive. In previous versions of Windows, such subsystems were fully case sensitive by default. In order to support standard UNIX behavior, Windows Services for UNIX Setup allows you to change the default Windows XP and Windows Server 2003 family behavior for non-Win32 subsystems when installing the base utilities (which installs the Interix subsystem) or Server for NFS. If you enable case sensitivity and then subsequently uninstall Server for NFS and the base utilities, Windows Services for UNIX Setup will restore the default, case-insensitive behavior of non-Win32 subsystems.

8. Configure User Name Mapping

User Name Mapping acts as a single clearinghouse that provides centralized user mapping services for the NFS client of Interix. User Name Mapping provides a map between the Windows users and groups on the NFS client, and the corresponding UNIX users and groups on the NFS server. In principle, these user and group names might not be identical, but for users who intend to use N1GE the names *must* be identical.

User Name Mapping lets you maintain a single mapping database for the entire enterprise. This feature makes it easy to configure authentication for multiple computers running Windows Services for UNIX. In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping letting you associate multiple Windows accounts with a single UNIX account. This feature can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission. You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps.

For information about simple and advanced maps, see "Simple and advanced maps" in "Help for Services for UNIX"


Note – Note: After the installation has finished, you can find "Help for Services for UNIX" here:

Start -> Programs -> Services for UNIX -> Help for Services for UNIX

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a Windows-based system running Server for PCNFS.

Note – User Mapping is part of SFU and not of N1GE. Please consult Microsoft Documentation and/or support to setup user mapping correctly.

Your selection in this dialog depends on the hosts and services which are currently provided



in your Windows environment and also in the UNIX environment. Otherwise, if there is no such server in your environment then you should select Local User Name Mapping Server.

Note – You should install SFU and enable the User Name Mapping service on your host which acts as Domain Controller for your windows environment. All other hosts should contact that Remote User Name Mapping Server.

If you choose Local User Name Mapping Server then you might either select Network Information Services (NIS) to access your passwd and group NIS-maps. Otherwise you have to select l if you will provide the files yourself.

9. Depending on your previous selections, you have either to enter the NIS Domain name and NIS Server name or the path of the passwd and group files.



Following, you will find an example for the files which have the standard UNIX format. This means that you can also use your /etc/passwd and /etc/group files from your UNIX environment.

C:\Unix\etc\passwd

```
root:x:0:0:UNIX root user:/home/root:/bin/tcsh
user1:x:1002:100:Full name of user1:/home/user1:/bin/tcsh
```

C:\Unix\etc\group

```
root::0:
```


Note – Some NIS maps do not contain an entry for user root. If this is the case and you intend to create following mapping, Administrator <-> root.

You create this entry by using the following steps:

- a. First create a password file containing the root entry during this installation step.
- b. If the SFU installation is finished then start the Services for UNIX Administration application and create the mapping: Administrator <-> root.
- c. Switch to NIS mapping.
- d. Use simple mapping or add manual mappings.
- e. At this point the installation starts installing components. Wait until all components are installed.



10. When the installation process finishes, you may have to reboot the machine, depending on the version of Windows you are using.



11. Make sure that the Interix Subsystem Startup is started during boot time. If you intend to use NFS shares and user mapping then also start Client for NFS and User Name Mapping. Depending on the installation options and version of the Windows operating system, one or more of these services are disabled by default.

Post SFU Installation Tasks

There are several tasks you should do after you install the SFU software.

1. Before you start using SFU and before you actually install N1GE, please check that user mapping is working correctly by following these steps.
 - a. Open an Interix shell locally on the Interix host.
 - b. Use the `login` command to switch to a known user that is not the Administrator.
 - c. Verify the access permissions of NFS shares which should be accessible by that user.
 - d. Try to access these network resources. If a user cannot access a Network drive, most likely the User Name Mapping is not working correctly.
2. Check users home directories. To enable the automounting of the users home directories use this series of menus:

Control Panel -> Administrative Tools -> Computer Management -> Users -> Properties -> Profile

Click on connect to, select a drive letter, and enter the path of the users home directory in UNC notation: \\<server>\<share>\<user home>

Within the Interix subsystem, you might access all NFS shares through the special directory: /net/<server>/<share>

You might also create links to these directories to access the shares directly (for example, `ln -s /net/myserver/export/share00/home /home`).

3. Enable Administrator names on your machines.

Make sure that the Administrator accounts on all machines which should be enabled as executions hosts for N1GE use the same account name (e.g. Administrator).

Also make sure that this user has manager privileges in your N1GE cluster. If this is not the case then before the installation of the execution daemon, please add the privileges using this command: `qconf -am <Administrator>`

4. Set the CLI commands which start an editor. Make sure to set the EDITOR environment variable to `vi` or your preferred UNIX editor within the Interix subsystem before you start using UNIX commands.

5. Mount NFS shares.

There are two ways to mount NFS shares to the Interix host:

- The recommended way is using the auto mount functionality of Interix. All network shares that the Computer Browser service of Windows can find are automatically mounted to the `/net` directory of Interix. Some of them may be listed with `"ls /net"`, some of the may be not, but they are anyway accessible. The syntax of the auto mounts is `"/net/<server>/<share>"`, e.g. `"/net/myserver/home"`. A link to a auto mounted share can be created to make it accessible under exactly the same name as on a Unix host. E.g. `"ln -s /net/myserver/home /home"` will make the users Unix home directories accessible via `"/home/username"`. Automounted shares are available from boot time on, they are available for all users (who have, in general, the permissions to access these shares) and can't get lost by misconfiguration.
- Network shares can also be mapped to drive letters. This is done with the command `"nfsmount"`. The syntax is `"/usr/sbin/nfsmount -u: \\computername\sharename devicename"` e.g. `"/usr/sbin/nfsmount -u: \\myserver\home Z:"` (All backslashes have to be written two times, because the shell interprets a single backslash as an escape character.) This drive is now accessible via `"/def/fs/Z"`. A link can be created to this drive to have the same path as on a Unix host.

6. Enable Interix inet daemons.

Although the Windows remote shell service and the Windows telnet service might be running on the host, you should disable the Windows daemons and use the Interix daemons instead. To disable the Windows services, go to `"Control Panel"/"Services"` and stop the services, set the start type to `"disabled"`. To enable the Interix daemons, edit the `/etc/inetd.conf` file. It already contains lines for the daemons that are commented out. Remove the comment character (`#`) and save the file. Then run `"/etc/init.d/inet start"` to start the daemons. These daemons are necessary to make the N1GE commands `qrsh` and `qlogin` working.

If the user wants to do a passwordless rlogin to the host and wants to have access to network resources then, the user must register the password with the Interix command "regpwd". This is because the network component of Windows never trusts a passwordless login, with the registered passwords the rlogind can emulate a login with password."

7. Configure the users' home directories.

If the users' home directories are located on an NFS server, this must be configured in Windows. In the "profile" tab of the users properties dialog, select "Connect", a free drive letter and enter the path to the users home directory in the UNC notation \\server\share\directory, e.g. "\\myserver\home\Peter".

Troubleshooting SFU

The following section describes some common problems that users may encounter when installing and using N1GE in a Services for UNIX environment on a Windows system.

- Impossible connect to the Interix subsystem via telnet or rsh

Make sure that the correct services are started. The corresponding Windows services have to be disabled. The Interix versions of telnetd and rshd have to be started. You can do this task by removing the pound sign (#) from the following lines part of /etc/inetd.conf

```
#telnet stream tcp nowait NULL /usr/sbin/in.telnetd in.telnetd -i
```

```
#shell stream tcp nowait NULL /usr/sbin/in.rshd in.rshd -a
```

If it is still not possible to connect to the machine, please check your firewall configuration. Donot block connections to corresponding ports:

Service	Ports
ftp	20, 21
ssh	22
telnet	23
rsh	514

- The wrong default login shell is started. Why?

.rhost and host.equiv authentications fails if new users accounts are created and if passwords of existing users are changed. The command regpwd has to be called. After that, use the steps to register passwords correctly.

- Why is the access to NFS mounted home directories to slow?

User Name Mapping might be the cause for this problem. If you have a large number of user maps, installing User Name Mapping on a Domain Controller improves performance by reducing network traffic. It is possible to create a User Name Mapping server pool. This method means you can use DNS round-robin to create a pool of computers running User

Name Mapping. This will provide improved performance on wide area networks as well as provide fail over when one of the servers is no longer available

- How can I map user root if it does not exist in the NIS maps?

First create a passwd file which contains an entry for the user root. Then, explicitly map the root account (no basic mapping) using the created passwd file. Finally, change the mapping to use the NIS maps. Note that the previous root mapping will persist.

- NIS Server cannot be contacted during the SFU installation

Interrupt the SFU installation and make sure that there is no other service or application running which already configures or uses the NIS server. If this is the case, then disable this service for the time of the SFU installation.

- The Interix Subsystem of SFU and/or the User Mapping is not enabled after reboot

Make sure that following services are automatically started after machine reboot: Interix Subsystem Startup, User Name Mapping

- Also if you use NFS mounted directories that also enable the service per default: Client for NFS

- Queues stick in u - state for a very long time

After the installation/restart of an execution host the corresponding queues have attached the unknown (u) state for a very long. This is normal behavior for Windows machines. After a full load report interval the u-state should be gone. If this is not the case, then check that the sge_execd has really been started on the corresponding machine.

User Management For N1GE on Windows Hosts

Overview

Every user of the N1GE execution environment of a Windows machine must have a user account that has the same name as on the UNIX hosts. User accounts contain information about the user, including name, password, various optional entries that determine when and how users log on. and how their desktop settings are stored.

Following sections describe how you would use Windows user management to support N1GE. Windows machines are referred to here using three different terms. The following table lists the terms and the operating systems which might run on each corresponding host:

Windows Host	Microsoft Windows 2000, Microsoft Windows XP, Microsoft Windows 2000 Server, Microsoft Windows Server 2003
Windows Server	Microsoft Windows 2000 Server, Microsoft Windows Server 2003
Windows Workstation	Microsoft Windows 2000, Microsoft Windows XP

Managing Users on Windows Hosts

It is possible to administer user accounts on all Windows hosts individually. Each Windows Host has an authentication centre which validates usernames and corresponding user rights. User accounts which are defined on a Windows workstation are referred to here as local user accounts or local users.

Each Windows Host has it's own local domain and each Windows Server has the ability to make that domain available to other hosts. Account names within a local domain and account names within a server domain can collide. To avoid such collisions, you must specify the correct user account by providing the domain name as a prefix to the user account name followed by a '+' character.

Windows User Example

The following is an example that illustrates the potential complexity of Windows host accounts interacting with Windows Domain accounts. Suppose Windows Workstation host named *CRUNCH* has a local user account named *Peter*. This Windows Workstation is part of the domain named *ENGINEERING*. This domain is provided by a Windows Server which also has a user account named *Peter*. In this example, the *ENGINEERING* domain is the default domain of the host named *CRUNCH*. The following table shows the possible results of what would happen if a person tried to login to *CRUNCH*.

TABLE B-1 Using Domain Accounts

Login Name	Result
CRUNCH+Peter	Peter is logged in with his account as a local user of the machine <i>CRUNCH</i> .
ENGINEERING+Peter	Peter is logged in with the account provided by the Windows Server hosting the <i>ENGINEERING</i> domain.
Peter	This approach is equivalent to using <i>ENGINEERING+Peter</i> because <i>CRUNCH</i> has <i>ENGINEERING</i> as it's default domain. Otherwise, the local account would be used.

Each domain has a special user account that provides superuser access. The default name on English systems for that account is *Administrator*. For native Windows, the members of the Administrators group and of the Domain Admins group in the server domain also have superuser access. However, for Interix, only the user Administrator of the local domain is the superuser of the local host.

The local Administrator can start applications under an users account without knowing the password of this user, but because - in opposition to the Unix super user 'root' - even the local Administrator is not fully trusted by the network, the application wouldn't be able to access network resources then. This is why there is the sgepasswd tool to register the users passwords, as explained in "Using N1GE in a Microsoft Windows Environment."

UNIX User Management

UNIX has no equivalent to the Windows domain concept. With UNIX, each user has a local account and is authenticated as a local account even if the underlying account information lies on a LDAP or NIS server. The UNIX super user (*root*) is similar to the local Windows super user (*Administrator*). The UNIX super user can start applications and processes on behalf UNIX accounts without knowing each corresponding password.

Using N1GE in a Microsoft Windows Environment

The N1GE execution environment has to start jobs on behalf of the submitting user. The execution daemon (`sge_execd`) on UNIX hosts runs as the user `root` so it can start jobs on behalf of all users.

On Windows hosts the execution daemon runs as the local *Administrator* user so it can start jobs on behalf of users without knowing their password, but these jobs wouldn't have the permissions to access network resources. Only fully authenticated users can access network resources, and for a full authentication the users password is needed. Therefore all users who want to submit jobs to a Windows execution host have to register their password with N1GE. The execution daemon still needs to run as the local Administrator to have the permissions to do several administrative tasks. .

Registering Windows User Passwords

Users who want to start N1GE jobs on Windows execution hosts use the `sgepasswd` client application to register their Windows passwords with N1GE. The following example shows Peter who has a user account in the domain ENGINEERING. Because ENGINEERING is the principal domain of the Windows execution host CRUNCH, Peter does not need to register his password for a specific domain. This should be the default in any properly set up single domain environment. `sgepasswd` must be run on a non-Windows host.

```
> sgepasswd
  Changing password for Peter
  New password:
  Re-enter new password:
  Password changed
```

Only in multiple domain environments it could be necessary to register the password explicitly for a specific domain.

Using the `sgepasswd` Command

`sgepasswd` modifies the Grid Engine password file (`sgepasswd(5)`.) This file contains a list of usernames and their Windows passwords in encrypted form. You can use the `sgepasswd` to add a new entry for your user account or change your existing password. Additionally, the root user can change or delete the password entries for other user accounts. `sgepasswd` is only available on non-Windows hosts. It has the following syntax:

```
sgepasswd [[ -D domain ] -d user ]

sgepasswd [ -D domain ] [ user ]
```

This command has these options:

-D domain	By default sgepasswd adds or modifies the current UNIX username without a domain specification. You can use this switch to add a domain specification in front of the current user name. Consult your Microsoft Windows documentation to get more information about domain users.
-d user	Only root can use this parameter to delete entries from the sgepasswd (5) file.
-help	Prints a listing of all options.

Additionally, The following environment variables affect the operation of this command.

SGE_CERTFILE	Specifies the location of public key file. By default sgepasswd uses the file <code>\$SGE_ROOT/\$SGE_CELL/common/sgeCA/certs/cert.pem</code>
SGE_KEYFILE	If set, specifies the location of the private key file. The default file is <code>/var/sgeCA/port\${SGE_QMASTER_PORT}/\${SGE_CELL/private/key.pem</code>
SGE_RANDFILE	If set, specifies the location of the rand.seed file. The default file is <code>/var/sgeCA/port\${SGE_QMASTER_PORT}/\${SGE_CELL/private/rand.seed</code> .

Adding Windows Hosts to Existing N1 Grid Engine Systems

If you have a running N1 Grid Engine system on which Windows support is not enabled, you can enable the support manually. The following steps provide a Windows-enabled N1 Grid Engine system that allows additional Windows execution hosts.

▼ How to Add Windows Hosts Later

- 1 **Copy Windows binaries to the `<sge-root>` directory.**
- 2 **Type the following command:**
`qconf -mconf`
Set the `execd_params` to `"enable_windomacc=true."`
- 3 **Type the following command:**
`qconf -am <win_admin_name>`
- 4 **Run the following command:**
`$SGE_ROOT/util/sgeCA/sge_ca -init -days 365`
- 5 **For a CSP installation, run the following command:**
`$SGE_ROOT/util/sgeCA/sge_ca -user <win_admin_name>`

6 Type the following command:

```
qconf -ah <new_win_hosts>
```

7 Copy certificates to each Windows host.**8 Set the owner of the certificates to ADMINUSER**

Use a command similar to the following example:

```
chown -R foo:bar /var/sgeCA/port<SGE_QMASTER_PORT>
```

9 Run normal exec daemon installation on each execution host.

Other N1 Grid Engine Installation Issues

Checking Linux Motif Libraries

On newer Linux systems, the `libXm.so.2` Motif libraries are not always installed resulting in the inability to run the precompiled Linux `qmon` binary.

To correct this problem, use the following steps.

▼ Installing Linux Motif Libraries

- 1 **Check if the libraries are already present by using this command.**

```
% ls -l /usr/X11R6/lib/libXm*
```

If the `/usr/X11R6/lib/libXm.so.2` points to a `libXm.so.2.x` version you are done. Note that a symbolic link to `/usr/X11R6/lib/libXm.so.3` does not work.

If the libraries are not present then continue following these steps.

- 2 **Download the corresponding `openmotif` libraries from**
http://www.ist.co.uk/DOWNLOADS/motif_download.html **or from the SUSE 9.1 distribution**
(there is an additional rpm file called `openmotif21-*` available).
- 3 **Install the missing libraries as root. For Suse 9.1, you install the `openmotif21-*` package like any other package. For packages downloaded from**
<http://www.ist.co.uk>, **install the libraries as in the following example.**

```
# rpm -i --prefix /tmp/test --force \
    openmotif-2.1.31-2_IST-JDS2003.i386.rpm
# cd /tmp/test/OpenMotif-2.1.31/lib
# cp libXm.so.2.1 /usr/X11R6/lib
# cd /usr/X11R6/lib
# ln -s libXm.so.2.1 libXm.so.2
```

4 Test qmon.

```
% ldd 'which qmon'
```

Installing N1 Grid Engine on a System with IPMP

How to install the Grid Engine software on hosts with the Solaris Operating Environment IP Multipathing (IPMP) technology

What is IP Multipathing?

IP Multipathing is a technology that allows grouping of TCP/IP interfaces for fail over and load balancing purposes. If an interface within an IP Multipathing group fails, the interface is disabled and its IP address is relocated to another interface in the group. Outbound IP traffic is distributed across the interfaces of a group. For further details on IP Multipathing, refer to the Solaris Operating Environment documentation, which can be found at:

<http://docs.sun.com/db/doc/806-4075/6jd69oabu?a=view>. The IPMP features overview can be found at:

<http://www.sun.com/software/solaris/ds/ds-netmultipath/index.html>.

Issues between IPMP and Grid Engine

The only major issue is the error messages which occur when starting the Grid Engine daemons on a machine where the main interface is part of an IPMP group. This situation occurs when the IPMP load balancing distributes the connections across the interfaces in the group; therefore, the IP packets show up at the receiving end as coming from a different host rather than the one associated with the main interface. For example, on a machine with three interfaces named qfe0, qfe1, and qfe3, where the IP addresses for these interfaces are 10.1.1.1, 10.1.1.2 and 10.1.1.3 respectively, IPMP would need an extra address for each interface for testing. However that requirement is ignored in this example. Each of these addresses has a hostname associated with it. The hosts table looks like the following example:

```
10.1.1.1 sge
10.1.1.2 sge-qfe1
10.1.1.3 sge-qfe2
```

The machine's hostname is sge. When a connection is established from sge to another machine, it might go through sge, sge-qfe1, or sge-qfe2. Upon installation, Grid Engine will only recognize sge. When Grid Engine receives a connection request from sge-qfe2, it closes the connection because the request is not from one of the authorized (or known) nodes.

You solve this problem by using the `host_aliases` files (see the `sge_h_aliases` man page for details). You can use this file to "tell" Grid Engine that sge, sge1, and sge-qfe2 are all from the same machine. The `host_aliases` file in this case would look like this:

```
sge sge-qfe1 sge-qfe2
```

Note that if you make any changes to the `$SGE_ROOT/$SGE_CELL/common/host_aliases` file, you must stop and restart all running Grid Engine daemons (`sge_qmaster`, `sge_scheduler`, and `sge_execd`). To do this, login as root to all your Grid Engine hosts and enter these commands:

```
/etc/init.d/sgemaster stop
/etc/init.d/sgeexecd stop
    /etc/init.d/sgemaster start
/etc/init.d/sgeexecd start
```

Installing the Grid Engine Master Node With IPMP

There are two ways you can fix this problem: one way is to ignore the error messages during installation. This method is operating system independent (except for MS Windows) The other way is to temporarily disable IPMP on the interface associated with the machine's hostname. This method only works on systems with Solaris 8 or greater Operating Environments.

The ignoring the error messages procedure is as follows in these steps:

1. Run the `inst_sge -m` command while ignoring the error messages during the start up of the daemons.
2. Shutdown the daemons with the `/etc/init.d/sgemaster stop` and `/etc/init.d/sgeexecd stop` commands. Due to the networking errors, some daemons fail to shutdown and must be killed with the `kill -9` command. To see which daemons failed to shutdown use this command: `ps -e | grep sge_`.
3. Install the `host_aliases` file in the `$SGE_ROOT/$SGE_CELL/common` directory.
4. Restart the daemons with the `/etc/init.d/sgemaster start` and `/etc/init.d/sgeexecd start` commands.

The temporary disable method is as follows in these steps:

1. Identify the interface associated with the machine's hostname.
2. Verify that the interface has IPMP enabled with the `ifconfig <<interface>> | grep groupname` command.
3. Take note of the group name.
4. Disable IPMP with this command: `ifconfig <<interface>> group ""`.
5. Install the Grid Engine master node.
6. Install the `host_aliases` file in the `$SGE_ROOT/$SGE_CELL/common` directory.
7. Restart the daemons with the `/etc/init.d/sgemaster` and `/etc/init.d/sgeexecd` commands.

8. Re-enable IPMP using the following command: `ifconfig <<interface>> group <<IPMP group>>.`

Installing a Grid Engine on an Execution Host With IPMP

Once the `host_aliases` file is installed and the Grid Engine daemons are restarted, you can simply start the execution host installation without further problems.

Enabling Administrative and Submit Hosts with IPMP

You have two choices when enabling these hosts with IPMP:

- Follow the same procedure used for the execution host (updating the `host_aliases` file before installation.)
- Add all the hostnames associated with the administrative or submit host with one of the following commands:

```
qconf -ah <<hostname>> <<alias 1>> <<alias 2>> ...  
      (for the administrative host)
```

```
qconf -as <<hostname>> <<alias 1>> <<alias 2>> ...  
      (for the submit host)
```

Glossary

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.
administration host	Administration hosts are hosts that have permission to carry out administrative activity for the grid engine system.
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.
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.
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 <i>checkpoint</i> , thereby allowing the job to be aborted and resumed later without loss of information and already completed work. The process is called <i>migration</i> 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 <i>hosts</i> 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 <i>share</i> . 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 <code>sgexecd</code> .
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 <i>single</i> computational resource.
group	A UNIX group.
hard resource requirements	The resources that must be allocated before a job can be started. Contrast with <i>soft resource requirements</i> .
host	A system on which grid engine system functions occur.
interactive job	An interactive job is a session started with the commands <code>qssh</code> , <code>qsh</code> , or <code>qlogin</code> , which open an <i>xterm</i> window for user interaction or provide the equivalent of a remote login session.
job	A request from a user for computational resources from the grid.
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.
master host	The master host is central to the overall cluster activity. It runs the master daemon <code>sge_qmaster</code> and the scheduler daemon <code>sge_schedd</code> . By default, the master host is also an administration host and a submit host.
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.
share	The same as <i>entitlement</i> . 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 <i>hard resource requirements</i> .
submit host	Submit hosts allow for submitting and controlling <i>batch jobs only</i> . In particular, a user who is logged in to a submit host can submit jobs using <code>qsub</code> , can control the job status using <code>qstat</code> , and can use the grid engine system OSF/1 Motif graphical user interface <code>QMON</code> .
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 <i>access list</i> or a <i>department</i> .

Index

A

- abort installation procedure, 28, 51
- access list, 149
- account for installation, 16
- accounting and reporting
 - console
 - install, 120
 - dbWriter
 - enable reporting, 112
- accounting and reporting console, install, 120
- add administrative host, 50, 71
- add submit host, 50
- administration host, 149
- administrative host, 20
 - add, 50, 71
- administrative user, file handling, 16
- array, job, 149
- automounter, 14

B

- batch job, monitoring, 99
- Berkeley DB host, 18-19
- Berkeley DB installation procedure, 50-54

C

- campus grid, 149
- cell, 11, 14, 15, 149
- checkpointing, 149

- checkpointing environment, 149
- classic spooling, 18
- cluster, 149
- cluster grid, 149
- complex, 150

D

- database
 - for accounting and reporting, 105
 - Oracle, 109
 - PostgreSQL, 105, 108
- database host, 18-19
- database spooling host, 18-19
- date, 98
- dbWriter
 - for accounting and reporting
 - enable reporting, 112
 - install, 112
 - set environment variables, 112
- department, 150
- directory organization, 14
- disk space requirements, 13

E

- entitlement, 150
- environment
 - checkpointing, 149
 - parallel, 151
- /etc/services, 12, 16

execution host, 20, 150
 installation procedure, 43
execution host spool directory, 18

F

file access permission, 16
file handling, administrative user, 16
functional policy, 150

G

GEMM, 24
global grid, 150
grid, 150
 campus, 149
 cluster, 149
 global, 150
Grid Engine Management Module, 24
group, 150

H

hard resource requirements, 150
host, 150
 administration, 149
 Berkeley DB, 18-19
 execution, 150
 master, 151
 shadow master, 17-18
 submit, 152

I

id, equivalent user, 11
installation
 master host, 57-58
 qmaster, 57-58
 with root account, 16
 with unprivileged account, 16
installation account, 16

installation directory, 13-14
installation procedure
 abort, 28, 51
 Berkeley DB, 50-54
 Berkeley DB spooling server, 50-54
 execution host, 43
 master, 28
 spooling server, 51
interactive job, 149, 150

J

job, 150
 batch, 149
 class, 150
 monitoring, 99
 parallel, 151
job array, 149
job class, 150
job_id, 99

M

manager, 151
master, as execution host, 73
master host, 17, 151
 installing, 57-58
master installation procedure, 28
migration, 151

N

network services, 16-17
NIS, 12, 16-17, 97

O

operator, 151
Oracle database, accounting and reporting, 109
override policy, 151
owner, 151

P

- parallel environment, 151
- parallel job, 151
- permission, file access, 16
- policy, 151
 - functional, 150
 - override, 151
 - share-based, 151
- PostgreSQL database
 - accounting and reporting, 105, 108
 - create database, 108
 - create database owner, 108
 - create user, 108
- priority, 151
- project, 151

Q

- qconf -ah, 50
- qconf-ah, 71
- qconf-as, 50
- qmaster, installing, 57-58
- qmaster spool directory, 18
- qstat, monitor batch jobs, 99
- queue, 20-21, 149

R

- redirection
 - stderr, 99
 - stdout, 99
- requirements
 - hard, 150
 - soft, 152
- resource, 151
- resource requirements
 - hard, 150
 - soft, 152
- root account, 16
- root directory, 13
- rsh, 98

S

- schedd spool directory, 18
- scheduler, 21
 - profile, 21
 - high, 21
 - max, 21
 - normal, 21
- services, 12, 97
- services database, 16
- settings.csh, 97
- settings.sh, 97
- sgc_execd, 43
- SGE_EXECD_PORT, 97
- sgc_qmaster, 28
- SGE_QMASTER_PORT, 97
- sgc_schedd, 28
- shadow master host, 17-18
- shadow_masters file, 17
- share, 151
- share-based policy, 151
- share-tree, 152
- soft resource requirements, 152
- spool directories, 18
- spooling, classic, 18
- spooling server, installation procedure, 51
- stderr redirection, 99
- stdout redirection, 99
- submit host, 20, 152
 - add, 50
- Sun Web Console
 - for accounting and reporting
 - install, 119
- suspension, 152
- system requirements
 - disk space, 13

T

- ticket, 152

U

- unprivileged account, 16

usage, 152
user, 152
user id, equivalent, 11
user names, 15-16
user set, 152