



Sun™ ONE Grid Engine 5.3 Release Notes

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-5077-11
October 2002, Revision 01

Send comments about this document to: docfeedback@sun.com

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 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 of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

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 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

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

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, 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 licences de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE 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 CONTREFAÇON.



Adobe PostScript

Contents

About This Linux-Supported Release	1
Contents of This Software Package	2
Fresh Installation of Sun ONE Grid Engine 5.3 Software	2
About Upgrading an Existing System	3
Time Required for the Upgrade	4
Important Changes To Note	5
Names of Daemons and Commands	5
Daemon Startup Scripts	5
Environment Variables	6
Communication Daemon (commd) Service Name	6
Other File Name Changes	6
The Queuing System Interface (QSI)	7
About the Upgrade Procedure	7
▼ How To Perform the Upgrade	8
▼ How To Upgrade From Sun Grid Engine, Enterprise Edition 5.3beta2	18
▼ How To Install a Sun ONE Grid Engine, Enterprise Edition 5.3 Patch	19
Installation Rules	20
Installation Procedure	20
Installing Localized Packages	22

- ▼ How To Install Localized Packages Via `pkgadd` 22
- ▼ How To Install Localized Packages Via `tar` 23
- Security 23
- Irregularities Related to Documentation 23
- Irregularities Related to Various Locales 24
- Corrections and Additions to `man` Pages 26
 - `sgc_conf(5)` and `queue_conf(5)` 26
 - `access_list(5)` 26
 - `sgc_aliases` 27
 - `qmod` 28
 - `-e/-d/-us/-s *` 28
 - `-r` 28
 - Array Job Task Index Values Range 28
 - `qmake(1)` 28
 - ARCH 29

Sun ONE Grid Engine 5.3 Release Notes

Read this document carefully before you install the accompanying software. This document includes the following main sections.

- “About This Linux-Supported Release” on page 1
- “Contents of This Software Package” on page 2
- “Fresh Installation of Sun ONE Grid Engine 5.3 Software” on page 2
- “About Upgrading an Existing System” on page 3
- “Installing Localized Packages” on page 22
- “Security” on page 23
- “Irregularities Related to Documentation” on page 23
- “Irregularities Related to Various Locales” on page 24
- “Corrections and Additions to man Pages” on page 26

About This Linux-Supported Release

Sun ONE Grid Engine 5.3 is an updated Sun Microsystems product that fully supports the following version of the Linux operating system along, of course, with Sun’s Solaris™ Operating Environment:

- Linux kernel 2.4.x on the Intel architecture with `glibc 2.2.2` and later

Sun ONE Grid Engine release 5.3 is a new release of the Sun Grid Engine product, formerly available as version 5.2.3.

Contents of This Software Package

The Sun ONE Grid Engine 5.3 distribution contains the following top-level directory hierarchy.

- `3rd_party` – Contains information about freeware, public domain, and public license software used
- `bin` – Sun ONE Grid Engine executables
- `ckpt` – Sample checkpointing configurations
- `doc` – PDF-format documentation
- `examples` – Sample scriptfiles, configuration files and application programs
- `inst_sge` – Generic Sun ONE Grid Engine installation procedure
- `install_execd` – Driver for a simplified execution host installation
- `install_qmaster` – Driver for a simplified master host installation procedure
- `lib` – Required shared libraries
- `locale` – Contains localization catalogs
- `man` – On-line manual pages in nroff format
- `mpi` – Sample parallel environment interface for the MPI message passing system
- `pvm` – Sample parallel environment interface for the PVM message passing system
- `qmon` – Pixmaps, resource and help files for qmon, the graphical user interface
- `util` – Some utility shell procedures used for installation tasks and some template Sun ONE Grid Engine shutdown and boot scripts
- `utilbin` – Some utility programs mainly required during the installation

Fresh Installation of Sun ONE Grid Engine 5.3 Software

See `SGE53AdminUserDoc.pdf`, the *Sun ONE Grid Engine 5.3 Administration and User's Guide* included in PDF format in this distribution—in the `SDRMdoc` package or corresponding `tar.gz` file—for installation instructions. To upgrade an existing Sun Grid Engine or predecessor system, see the following section.

About Upgrading an Existing System

Use the instructions in the following sections to upgrade your older version of Sun Grid Engine or Sun Grid Engine, Enterprise Edition software with version 5.3 of the product. Note that these instructions apply both to Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 software, and that *many of these instructions apply only to Sun ONE Grid Engine, Enterprise Edition 5.3 software.*

Note – If you already have version 5.3 of either Sun Grid Engine or Sun Grid Engine, Enterprise Edition installed, and you are now installing the patch release of either product, see the section, “How To Perform the Upgrade” on page 8. If you are upgrading from Sun Grid Engine, Enterprise Edition 5.3beta2, use the procedure in the section, “How To Upgrade From Sun Grid Engine, Enterprise Edition 5.3beta2” on page 18.

These instructions mainly affect sites that have one of the following products installed.

- CODINE, a product of Gridware, Inc.
Using the procedure in this section, you can upgrade CODINE versions 5.0.x and 5.1.x.
- Global Resource Director, a product of Gridware, Inc.
Using the procedure in this section, you can upgrade Global Resource Director 5.0.x and 5.1.x.
- Sun Grid Engine versions previous to version 5.3
Using the procedure in this section, you can upgrade Sun Grid Engine versions 5.2.x, and 5.3beta1. (To upgrade from Sun Grid Engine, Enterprise Edition 5.3beta2, see the section, “How To Upgrade From Sun Grid Engine, Enterprise Edition 5.3beta2” on page 18.)

Note – The update procedure does *not* support any upgrades from CODINE 4.x or Global Resource Director 1.x. Note also that you *cannot* “downgrade” your site’s software from any Global Resource Director version to Sun ONE Grid Engine 5.3, also known as the “baseline” product. Sites that are using the Global Resource Director product should upgrade to Sun ONE Grid Engine, Enterprise Edition 5.3 software.

Time Required for the Upgrade

The following are estimates of the amount of time you will need to carry out the various tasks involved in upgrading your software. Of course, these are only estimates and the particular nature of your current installation may require more or less time.

- Preparing the upgrade: from 30 minutes to two hours
- Shutting down, and verifying the shutdown of, daemons; one to three minutes per host
- Deleting old spool files:
 - Spool directories on shared file system: five minutes
 - Spool directories installed locally: one to three minutes per host
- Backing up the old configuration: five to 10 minutes
- Deleting the old version: five to 10 minutes
- Installing the new version: five to 10 minutes
- Running the upgrade procedure: five to 10 minutes
- Updating the local startup scripts and starting the daemons: one to three minutes per host

Depending on the complexity of your installation, you may have to carry out additional functional tests for the various Sun ONE Grid Engine, Enterprise Edition 5.3 objects, such as the following.

- Parallel environments
- Checkpointing environments
- Load sensor scripts
- Cluster and queue configuration settings such as `prolog`, `epilog`, `terminate`, `suspend` methods
- Command tests (`qsub`, `qrsh`, `qlogin`, `qsh`)

If you have modified original sample scripts and configuration examples—and you are referencing these files located in the original distribution—or if you decide not to use the compatibility mode for environment variables (which is suggested), these additional tests may take several hours to ensure the full functionality of the Sun ONE Grid Engine, Enterprise Edition 5.3 system.

The upgrade procedure will be easier to carry out if you have access to user root (`rsh` or `ssh`) without providing a password to all your execution hosts. This does not necessarily need to be from your `qmaster` machine, by the way.

Important Changes To Note

Among the important changes in this version of the Sun ONE Grid Engine 5.3 software relative to Sun Grid Engine versions 5.0 through 5.2.3 are the names of the daemons, environment variables, and communication service daemon.

Names of Daemons and Commands

The names of all Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 daemons now begin with the `sge_` prefix. Affected names are the following.

- `sge_commd`
- `sge_coshepherd`
- `sge_execd`
- `sge_qmaster`
- `sge_schedd`
- `sge_shadowd`
- `sge_shepherd`
- The name of the `{cod|grd}commdcntl` command has been changed to `sgecommdcntl`.

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

Daemon Startup Scripts

- The system-wide startup script which, in former versions of the product, was in `<codine_root>/<cell>/codine5` or `<grd_root>/<cell>/grd5` now has the following name: `<sge_root>/<cell>/rcsge`
- The per-machine startup script, which is often installed in `/etc/init.d/{codine5|grd5}`, now has the following name: `rcsge`
- The symbolic link, which is often installed in `/etc/rc2.d/S95{codine5|grd5}`, now has the following name: `S95rcsge`

You must *delete* the old local startup scripts and the symbolic links, and replace them with the new startup script. On Linux systems with the `insserv` binary, the number for the `S95rcsge` script will be determined dynamically.

Environment Variables

The name of environment variables which, in previous versions of the product, begin with the prefix, `CODINE_/COD_/GRD_`, now begin with the prefix, `SGE_`.

Some examples follow.

- Former `CODINE_ROOT/GRD_ROOT` is now `SGE_ROOT`
- Former `COD_CELL/GRD_CELL` is now `SGE_CELL`
- Former `COD_O_HOME` is now `SGE_O_HOME`
- Former `GRD_STDOUT_PATH` is now `SGE_STDOUT_PATH`

Note – The Sun ONE Grid Engine, Enterprise Edition 5.3 program supports a compatibility mode in which the old names of variables still can be used. The upgrade script will ask if you want to set this mode. However, it is possible that future versions of the product will *not* support this compatibility mode; therefore, use of this compatibility mode is *not* recommended. If the size of a typical user environment is already near its limit you may encounter problems with the proper setting of all environment variables when a job is started.

Communication Daemon (`commd`) Service Name

The name of the communication service daemon has changed from `codine_commd/grd_commd` to `sge_commd`. The former name of the daemon is no longer supported.

Other File Name Changes

All files and manual pages which had the `codine_/cod_/grd_` prefix are renamed and use the `sge_` prefix. Following is a list of affected file names.

- In the `<root_dir>/<cell>` directory:
 - The former `codine_aliases` and `grd_aliases` files are now named: `sge_aliases`
 - The former `cod_request` and `grd_request` files are now named: `sge_request`
- In the user's home directory and the submit directory:
 - The former `.cod_request` and `.grd_request` files are now named: `.sge_request`

The Queuing System Interface (QSI)

The Queuing System Interface (QSI) feature is no longer a part of the product.

About the Upgrade Procedure

The full upgrade procedure involves the following tasks.

- Shutting down the cluster
- Backing up your old system
- Deleting old files and directories
- Unpacking the distribution
- Renaming the service
- Deciding about `admin_user`
- Running the update script
- Installing a new startup script on every host
- Starting the new Sun ONE Grid Engine system

Note – After backing up your old CODINE, Global Resource Director, or Sun Grid Engine cluster, it is advisable to delete all of the former version's files.

▼ How To Perform the Upgrade

1. Shut down the cluster.

Shut down your cluster before upgrading. There must be no running or pending jobs at `qmaster` or at the execution hosts. You should also make sure that there are no more running communication daemons (`cod_commd`, `grd_commd`, `sge_commd`) or a running scheduler daemon (`cod_schedd`, `grd_schedd`, `sge_schedd`) on the `qmaster` host.

To be absolutely safe, you can log in to every host of your older cluster and execute an appropriate `ps` command. You might want to issue the `grep` command for the string, `cod_`, `grd_`, or `sge_` in your `ps` output to identify any remaining processes from the older system.

When you execute the shutdown commands (as described in the following paragraphs), all components of the older system should be gone.

Enter the following commands to shut down your existing cluster.

```
# qconf -kej
# qconf -ks
# qconf -km
# {cod|grd|sge}commdctl -k
```

The first command kills all execution daemons and jobs. The second command kills the scheduler daemon. The third command kills the master daemon. The fourth command kills the communication daemon.

Note – You must execute the fourth command from *every execution host* as well as from the master host.

2. Make absolutely sure that there are no old jobs in your existing configuration by deleting the execution daemon spool directories and the spooled jobs in the `qmaster` spool directories.

Enter the following commands.

```
# rm -rf <execd_spool_dir>/<hostname>
# rm -rf <qmaster_spool_dir>/jobs
# rm -rf <qmaster_spool_dir>/job_scripts
# rm -rf <qmaster_spool_dir>/zombies
```

3. Back up the existing system

Before beginning the upgrade procedure, it is highly advisable to make a backup of your existing system.

To minimize the size of the backup, you may safely delete the following.

- The `execd` spool directories (no configuration is stored here) in:
`<execd_spool>/<hostname>`
- Old messages files of `qmaster` and scheduler: `<qmaster_spool_dir>/messages`
and `<qmaster_spool_dir>/schedd/messages`

Enter the following command to make the backup.

```
% tar cvf OLDSGE-BACKUP.tar <your_sge_root_dir>
```

4. Create a separate backup of your existing configuration.

Assuming everything is installed in standard locations, enter the following command to make the backup.

```
% tar cvf OLDSGE-CONFIG.tar $CODINE_ROOT/default/common \  
$CODINE_ROOT/default/spool/qmaster
```

Caution – Do *not* delete the old cluster configuration. Thus, you must *not* delete your `common` directory which is located in `<your_sge_root>/<cell>/common`. The directory path is usually `<your_sge_root>/default/common`. You also must *not* delete your `qmaster` spool directory. The path to your `qmaster` spool directory is defined in the global cluster configuration. Often, it is located in `<your_sge_root>/default/spool/qmaster`.

5. Check any references in the configuration to old files.

If you made any local changes to files of the distribution, make sure to make a separate backup of these files. Typically, these might be files that were modified and are now referenced in your cluster configuration or are used by jobs from your users. Among them are the following examples.

- Files that are in the `mpi/`, `pvm/`, or `ckpt/` directory, and are used by a parallel environment (PE) or checkpointing environment (CKPT)
- Wrapper commands that are in `mpi/` and are used by batch job scripts
- Load sensor scripts in `util/resources/loadsensors`

6. Delete the old distribution that was installed by way of the Solaris `pkgadd` procedure.

If you installed Sun Grid Engine, Enterprise Edition 5.2.x or 5.3beta on a Solaris system by way of the `pkgadd` procedure, it is typically safe to use the `pkgrm` command to remove the following packages (not all are usually available on your system):

- `SDRMdoc`
- `SDRMcomm`
- `SDRMsp32`
- `SDRMsp64`
- `SDRMsia`
- `SDRMedoc`
- `SDRMEcomm`
- `SDRMEsp32`
- `SDRMEsp64`

Enter the following command to see what packages are installed.

```
# pkginfo | grep SDRM
```

7. Delete the old distribution that was installed from tar.gz files.

If you installed a previous version by unpacking the tar distribution, it is typically safe to delete the following files and directories (not all files may exist in your existing CODINE, Global Resource Director, or Sun Grid Engine root directory).

Enter the following commands.

```
# cd <your_sge_root>
# rm -rf 3rd_party
# rm -f README*
# rm -f LICENSE*
# rm -f UPGRADE*
# rm -rf api
# rm -rf bin
# rm -rf catman
# rm -rf ckpt
# rm -rf doc
# rm -rf examples
# rm -f inst_codine
# rm -f inst_grd
# rm -f install_execd
# rm -f install_qmaster
# rm -rf locale
# rm -rf man
# rm -rf mpi
# rm -rf pvm
# rm -rf qmon
# rm -rf qsi
# rm -rf security
# rm -rf util
# rm -rf utilbin
```

8. Unpack the new distribution.

Do the following.

a. Log in to the machine where user `root` has read/write permissions in the `$SGE_ROOT` directory.

This is either your file server or a machine where the NFS mount point is configured appropriately. Installing the distribution needs to be done by user `root`. The NFS clients must *not* mount the `$SGE_ROOT` directory with the NFS mount option, `-nosuid`. Otherwise, the `qrsh` command (and related commands such as `qmake` and `qtchsh`) will not work.

If you cannot mount the `$SGE_ROOT` directory without the `-nosuid` option, you can configure the path to your `qrsh` command in the global and local cluster configuration.

It is not necessary to allow read/write permissions for user `root` to install and run Sun ONE Grid Engine 5.3 or Sun ONE Grid Engine, Enterprise Edition 5.3 successfully.

9. (Optional for `pkgadd` installers only) Install the product with the Solaris `pkgadd` program.

This step is for installers who plan to install Sun ONE Grid Engine 5.3 for Solaris with the `pkgadd` method.

a. Remove all previous SDRM* packages or all directories of the distribution that are mentioned in Step 7.

Note that the default base directory (`$SGE_ROOT`) and default *adminuser* has changed to the following.

- Base directory: `/gridware/sge` (for Sun ONE Grid Engine 5.3) or `/gridware/sgeee` (for Sun ONE Grid Engine, Enterprise Edition 5.3)
- *Adminuser*: `sgeadmin`

Note – It is advisable to install this product by way of the Solaris `pkgadd` method if it's technically possible; that is, if the file server is a Solaris machine or the file system is mounted by a Solaris machine. This will allow you to manage the software and patches with standard mechanisms provided by Sun. If you are going to install this product for other binary architectures, you can easily add and unpack the `tar.gz` files for these architectures in your `$SGE_ROOT` directory.

b. Select your previous base directory and your previous *adminuser* name for the new installation.

10. (Optional for installers using `tar` files) If you downloaded the new distribution as `tar.gz` files, log in to the machine where user `root` has appropriate file permissions.

a. As root, enter the following commands.

```
# cd <your_sge_root>
# umask 022
# gzip -dc sge-<version>-common.tar.gz | tar xvpf -
# gzip -dc sge-<version>-doc.tar.gz | tar xvpf -
# gzip -dc sge-<version>-bin-<arch>.tar.gz | tar xvpf -
```

b. Set the file permissions.

Enter the following commands.

```
# cd <your_sge_root>
# util/setfileperm.sh <adminuser> <unixgroup> <your_sge_root>
```

In the commands above, *<adminuser>* is the UNIX user account under which the spool files by Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 should be created. See Step 12 for more information about the *adminuser* feature of the product. *<unixgroup>* is the UNIX group name.

<your_sge_root> is the absolute path where your `$SGE_ROOT` directory is located.

The following example illustrates the commands in this step.

```
# cd <your_sge_root>
# util/setfileperm.sh sgeadmin adm `/bin/pwd`
```

11. (Optional) Rename the `codine_commd`/`grd_commd` service.

If you are using the TCP service, `codine_commd` or `grd_commd` for defining the communication port in `/etc/services` or your NIS server, you must rename the service to: `sge_commd`

12. Decide about using the *adminuser* feature.

Ever since the CODINE 5.0 and Global Resource Director 5.0 versions, Sun Grid Engine products have supported the *adminuser* user feature. The purpose of this feature is to start and run Sun ONE Grid Engine daemons on NFS clients where user `root` has no read/write permissions.

For security reasons, however, many sites do not want to export their NFS file systems to the NFS clients and thereby grant read/write permissions for user `root`. If you want to configure Sun ONE Grid Engine 5.3 products to use the *adminuser*, do the following.

- a. **Create the *adminuser* account on all of your Sun ONE Grid Engine hosts (qmaster and execution hosts), or create the *adminuser* in your NIS passwd database.**

sgedadmin is the recommended user name. The password for the *adminuser* should have the same protection as the password for *root*. The password for the Sun ONE Grid Engine *adminuser* may not be given to any users who are not entitled to get the *root* password on your systems.

- b. **Log in as user *root* on your NFS file server or on an NFS client where user *root* has read/write permissions in the *\$SGE_ROOT* directory.**

- c. **If you did *not* install the Sun ONE Grid Engine 5.3 distribution by way of *pkgadd* on a Solaris computer, run the following script.**

```
$SGE_ROOT/util/setfileperm.sh <adminuser> <group> <your_sge_root>
```

This script sets the file permissions of your Sun ONE Grid Engine distribution.

- d. **Edit the file, *\$SGE_ROOT*/*<cell>*/common/configuration, to modify the *admin_user* configuration entry.**

Enter the *adminuser* name; for example (recommended):

```
admin_user sgedadmin
```

- e. **Recursively change ownership of the following directories to the *adminuser* that you identified in the previous step.**

- *\$SGE_ROOT*/*<cell>*/common/
- *<qmaster_spool_directory>*
- *<execd_spool_directory>/hostname*

For example, if you have identified *sgedadmin* as *adminuser*, you would enter the following commands.

```
# chown -R sgedadmin $SGE_ROOT/default/common
# chown -R sgedadmin $SGE_ROOT/default/spool/*
```

Note – If the spool directories of your execution daemons are not installed in the default location, *\$SGE_ROOT*/*default/spool/<hostname>*, or if they are installed on a local file system, you must log in to every execution host and change the owner of the execution daemon spool directory.

13. **Log in as user *root* or as *adminuser* and run the update script.**

If you are using the *adminuser* feature, log in as the *adminuser* or run the update script with the *adminrun* command as described in Step 12a.

14. Set *SGE_ROOT*, *SGE_CELL* and other variables.

Note the following guidelines.

- Set your *SGE_ROOT* variable and *SGE_CELL* variable, if applicable. (If you used the default cell name, default, this is not necessary).
- Make sure that the setting of `$SGE_ROOT` does not contain any automounter prefixes such as `/tmp_mnt`.
- Make sure that `$SGE_ROOT` is set to a value such that this directory can be accessed from all your execution and submit hosts.
- You should set the `COMMD_PORT` variable if you are not using the `sge_commd` service. If you are using the `COMMD_PORT` variable, make sure to use an unused reserved port number.

```
# SGE_ROOT=<your_sge_root>; export SGE_ROOT (mandatory)
# SGE_CELL=<yourcell>; export SGE_CELL (depends on installation)
# COMMD_PORT=<portnumber>; export COMMD_PORT (depends on installation)
```

15. Run the update script.

Enter one of the following command sets.

```
# cd $SGE_ROOT
# util/sge_update.sh
```

Or:

```
# $SGE_ROOT/utilbin/<arch>/util/sge_update.sh
```

The script asks about your current product version and whether you want to upgrade to Sun ONE Grid Engine 5.3 (SGE) or Sun ONE Grid Engine, Enterprise Edition (SGEEE). The script will modify the following files.

In `$SGE_ROOT/<cell>/common`

- Delete:
 - `codine5|grd5`
 - `history/`
 - `license`
 - `qsi/`
 - `statistics`
- Rename:
 - From `codine_aliases|grd_aliases` to `sge_aliases`

- From `cod_request` to `sge_request` (if it exists)
- Update:
 - `configuration`
 - `product_mode`
 - `settings.csh`
 - `settings.sh`
- Create:
 - `rcsge` (the new startup script that replaces `codine5|grd5`)

In <qmaster_spool_directory>

- Delete:
 - `messages`
 - `jobs/`
 - `job_scripts/`
 - `zombies/`
 - `schedd/messages`
- Update:
 - `complexes/queue`
 - `exec_hosts/`
 - `schedd/`

A copy of your changed files and directories will be saved in `<SGE_ROOT>/<cell>/common/<YYYYmmdd-hh:mm:ss>`.

16. Locally remove the old per machine startup script and install the new script.

Depending on your operating system, the startup script is installed in `/etc/init.d/` or `/sbin/init.d/`, with a symbolic link in the corresponding `rc2.d` or `rc3.d` directory.

Enter the following commands.

```
# cd $SGE_ROOT
# util/update_commands/sge_startupscript.sh
```

These commands remove your old startup script and add the new startup script. Keep in mind that the procedure will *fail* to delete the old startup scripts *if you copied or renamed the script*.

Note – This procedure is *not* supported on IBM AIX and on Cray Unicos. Note also that on Linux systems with the `/sbin/insserv` binary (e.g., SuSe 7.1 or later), the script will install the new startup script with that mechanism (and try to remove the old startup script with `/sbin/insserv -r`).

You must replace the startup script on all execution hosts. If you have access for user `root` without providing a password, from one of your machines you may run the script initiated by the following commands.

```
# cd $SGE_ROOT
# util/sgeremoterun -all -- util/update_commands/sge_startupscrip.sh
```

By this method, you log in to all hosts for which an execution host in your `qmaster` spool directory is configured and run the script in the command line. The `sgeremoterun` script by default uses `rsh`; with the `-ssh` parameter, `remoterun` will use `ssh` to log in to the remote host.

To see all supported command line options, call `sgeremoterun` without any parameters.

Usage

```
util/sgeremoterun [-noexec] [-ssh] [-all] [-noqmaster] \  
                  [host1]... -- command
```

- `-noexec` – Do nothing, just print what would be done.
- `-ssh` – Uses `ssh` instead of `rsh`.
- `-sshpath` – Path where `ssh` is installed if not in `/bin:/usr/bin:/usr/local/bin`
- `-all` – Run command on all execution hosts found in `qmaster` spool directory.
- `-noqmaster` – Do not run command on current `qmaster` host.

17. Start the new Sun ONE Grid Engine or Sun ONE Grid Engine, Enterprise Edition system.

Enter the following command.

```
# $SGE_ROOT/$COD_CELL/common/rcsge -qmaster
```

If your `qmaster` host is also an execution host, you can also start the execution daemon by using the following command.

```
# $SGE_ROOT/$COD_CELL/common/rcsge -execd
```

To start Sun ONE Grid Engine execution daemons on all of your hosts, use the following commands.

```
# cd $SGE_ROOT
# util/sgeremoterun -all $SGE_ROOT/$COD_CELL/common/rcsge -execd
```

▼ How To Upgrade From Sun Grid Engine, Enterprise Edition 5.3beta2

You can upgrade from Sun Grid Engine, Enterprise Edition 5.3 beta2 to Sun ONE Grid Engine, Enterprise Edition 5.3 by changing the binaries only.

When you perform the upgrade, it is permissible to have pending jobs in the system. It is also permissible to have most types of running jobs in the system, provided that you take special care of the `sge_shepherd` binary. However, the upgrade procedure will not permit any running jobs of the following types in the system:

- `qmake`
- `qrsh`
- `qtcsh`
- `qlogin`
- Tightly integrated parallel jobs

1. **Make a backup of your old binaries, distribution, and configuration.**
2. **Shut down your cluster, making sure there are no running `sge_command` processes.**
Follow steps a and b.

- a. Enter the following command, and then wait one minute before continuing with the next step.

```
# qconf -ke all -ks -km
```

- b. After waiting one minute, enter the following command on the `qmaster` host.

```
# $SGE_ROOT/util/shutdown_commd.sh -all
```

3. Verify that no SGE daemons are running.

Do not kill any `sge_shepherd` if there are running jobs.

4. Enter the following commands to rename your `sge_shepherd` binaries.

Do not copy the binary.

```
# cd $SGE_ROOT/bin
# mv <arch>/sge_shepherd <arch>/sge_shepherd.sge53b2
```

5. Unpack the distribution.

6. Set the file permissions with the following command.

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

7. Start the new Sun ONE Grid Engine, Enterprise Edition 5.3 system on your `qmaster` host and your execution hosts.

See Step 17 in the previous procedure, “How To Perform the Upgrade.”

▼ How To Install a Sun ONE Grid Engine, Enterprise Edition 5.3 Patch

Note – The instructions in this section are for those who are installing a patch only. If you are installing the full product distribution which includes the most recent patches, you do not need to perform this procedure.

Two types of patches are available:

- Patch in `tar.gz` format

- Patch in Sun Microsystems patch format to be installed with `patchadd`

A patch in `tar.gz` format usually contains all binaries, including those that were not changed by the patch. A patch in `patchadd` format contains only the files that were changed by the patch.

These installation instructions assume that you are running a homogenous Sun Grid Engine cluster in which all hosts share the same directory for the binaries. If you are running a Sun Grid Engine cluster in a heterogenous environment—a mix of 32-bit and 64-bit binaries for Solaris and/or other operating systems—it is only necessary to shut down the daemons for the architecture for which the patch is applied. If you installed the binaries on a local partition, you need only to stop the Sun Grid Engine daemons for the host on which you are installing the patch.

Installation Rules

By default, there may be no *running* jobs when you install the patch. While there may be *pending* batch jobs, there can be no pending *interactive* jobs (`qrsh`, `qmake`, `qsh`, `qtcsh`).

It is possible to install the patch with running batch jobs. To avoid a failure of the active `sge_shepherd` binary, you must move the old `sge_shepherd` binary and copy it back prior the installation of the patch.

You can never install the patch with running interactive jobs, running `qmake` jobs, or other jobs that use the tight parallel Sun Grid Engine integration support.

If the patch contains a new `sge_commd` binary—which is always the case if you are installing the patch in `tar.gz` format—you also must move away the old `sge_commd` binary.

Installation Procedure

1. Enter the following command to disable all queues (so that no new jobs can be started).

```
# qmod -d '*'
```


2. (Optional—only needed if you have running jobs that should continue to run when the patch is installed.) Enter the following commands.

```
# cd $SGE_ROOT/bin
# mv solaris64/sge_shepherd solaris64/sge_shepherd.sge53
# cp -p solaris64/sge_shepherd.sge53 solaris64/sge_shepherd
```

It is important that you first move the binary, and then copy it back to the original location using the `-p` switch of the `cp` command.

3. Shut down and then restart the `qmaster` and scheduler daemon and all execution daemons on all Sun Grid Engine hosts.

Note – It is only necessary to shut down the communication daemons (`sge_commd`) if the patch contains a new `sge_commd` binary.

- a. Enter the following command, and then wait 30 seconds before proceeding to the next step.

```
# qconf -ke all
```

- b. After waiting 30 seconds, enter the following commands.

```
# qconf -ks
# qconf -km
```

4. (Optional) If the patch contains the `sge_commd` binary, follow these next steps. Otherwise, skip to .

- a. As `root` on your `qmaster` machine, enter the following command.

```
# $SGE_ROOT/util/shutdown_commd.sh -all
```

- b. Use the `ps` command to verify that the `qmaster` and scheduler daemons (`sge_qmaster`, `sge_schedd`), the execution daemon (`sge_execd`), and communication daemons (`sge_commd`) on all of your hosts are stopped.

5. Install the patch with the `patchadd` process or by unpacking the `tar.gz` files in `$SGE_ROOT`.

6. Restart your Sun ONE Grid Engine cluster.

a. Enter the following command on your qmaster machine.

```
# /etc/init.d/rcsgc
```

(If the above path to the startup script is incorrect for your operating system, enter the correct path to the script instead.)

b. Repeat this step on all your execution hosts.

Enter the following command to enable your queues.

```
# qmod -e '*'
```

If you renamed the `sge_shepherd` binary, you may safely delete the old binary when all jobs that were running prior to the patch installation are finished.

Installing Localized Packages

Note – Before installing localized packages, you must have installed the Sun ONE Grid Engine 5.3 packages.

▼ How To Install Localized Packages Via `pkgadd`

Use the following procedure to install localized software by way of the `pkgadd` method.

- At the command prompt, enter the first command and then choose and enter one of the three that follow, depending on which locale you want.

```
# unzip sge-5_3p2-solsparc-zh.zip (This is the Simplified Chinese localized package)
# unzip sge-5_3p2-solsparc-fr.zip (This is the French localized package)
# unzip sge-5_3p2-solsparc-ja.zip (This is the Japanese localized package)

# pkgadd -d . SUNWzhsgc (Simplified Chinese)
# pkgadd -d . SUNWfrsgc (French)
# pkgadd -d . SUNWjasgc (Japanese)
```

▼ How To Install Localized Packages Via tar

- At the command prompt, enter the following commands.

In the example, <tar_dir> is the abbreviation for the full directory and <xx> is one of ja, zh, or fr.

```
# cd <sge_root>
# zcat <tar_dir>/sge-5_3p2-solsparc-XX.tar.gz | tar xvpf -
```

Security

The Sun ONE Grid Engine 5.3 system can be operated in two modes with respect to security. The default mode is based on standard UNIX security and is intended for operation in a private secure network, typically behind a firewall for access through a reasonably trusted user base. Deployment across the Internet or in an insecure environment may expose your site to considerable security risks.

The second mode utilizes an OpenSSL-based embedded security framework. It authenticates accounts and services via certificates and encrypts Sun ONE Grid Engine 5.3 internal communication. Refer to the *Sun ONE Grid Engine 5.3 Administration and User's Guide* for detailed information on how to activate both modes.

Irregularities Related to Documentation

Issue – Though documented, the Queue type, “Transfer,” is not available in this release.

Workaround – No workaround exists for this problem.

Issue – The *Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 Reference Manual* incorrectly describes the `qconf -mqattr` switch as the “deprecated form of the `-mattr` switch,” when it is actually the deprecated form of the `-rattr` switch.

Workaround – Later versions of the manual will include this correction. Refer to the `man` pages, which have been corrected.

Issue – The *Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 Reference Manual* does not leave out a description of the `all` keyword for `queue_list` in `sge_pe(5)` and `checkpoint(5)` sections.

Workaround – Later versions of the manual will include this correction. Refer to the man pages, which have been corrected.

Irregularities Related to Various Locales

Issue – This software supports only `ja`, `zh`, and `fr` locales as the localized environments in the Solaris Operating Environment.

Workaround – If your current locale configuration is not `ja`, `zh`, or `fr`, explicitly set your configuration to one of them. For example:

```
% setenv LC_ALL ja
```

Issue – Related to the above issue, when `qmon` is started in any locale other than `C`, `ja`, `zh`, or `fr` (such as `zh_TW` or `ja_JP.PCK`), it may not automatically switch the locale-related configuration to that used for `C` locale. You may see the following problems, depending on the version of platform and your current locale:

- Core dumped (Segmentation Fault)
- No label/button string is displayed
- Some of label/button strings are incorrectly displayed and others are displayed in English

Workaround – Before bringing up `qmon`, set your locale to one of `C`, `ja`, `zh`, or `fr` explicitly, as shown in the following example.

```
% setenv LC_ALL C
```

In addition, if your default/current locale is not one of four locales listed above, you may have to switch the locale configuration of other software such as CDE and other components of Sun ONE Grid Engine software correspondingly.

Issue – No localization support is available for the Linux operating system in this release. Localized versions of this product are available only for the Solaris Operating Environment, versions 2.6, 7, 8 and 9.

Workaround – No workaround exists for this problem.

Issue – Occasionally, a flaw in the way that dynamically generated messages are assembled may cause non-English messages to contain English words.

Workaround – No workaround exists for this problem.

Issue – 64-bit libraries are missing for all short European locales. With the Solaris 8 Operating Environment, the core library in 64-bit is provided for the `fr` locale, so you cannot run the 64-bit version of Sun ONE Grid Engine 5.3 software in the `fr` locale.

Workaround – Apply patch 110023-04, which you can find at the following URL:

■ <http://sunsolve.sun.com>

Issue – All *localized* Sun ONE Grid Engine 5.3 objects must have ASCII-only identifiers. Non-ASCII localized names for queues, jobs, share nodes and leafs, and so forth are not supported. Sun has not determined how the product may behave with non-ASCII localized names, but does not believe it would lead to a problem as critical as a system crash, for example. However, this cannot be guaranteed and you may expect that display problems in the `qmon` graphical user interface will arise in certain cases.

Workaround – This is a limitation of Sun ONE Grid Engine 5.3 localized software for this release. Use only ASCII characters for the identifiers of objects such as queues, user names, complex names, Share Node and Leaf, and so forth.

Issue – In some `qmon` windows in `ja` and `zh` locales, components such as buttons or labels aren't completely displayed when the windows are open for the first time.

Workaround – Close the windows and then reopen them.

Issue – Some help files in `ja` and `zh` locales are not displayed perfectly. Localized online help isn't displayed fully and bottoms of their screens are truncated.

Workaround – No workaround exists for this problem.

Issue – When you start 64-bit version of `qmon` in the `fr` locale, you'll see warning messages related to "fontlist."

Workaround – Use the 32-bit `qmon` version instead. To do this, enter the following commands.

```
% mv $SGE_ROOT/bin/solaris64/qmon $SGE_ROOT/bin/solaris64/qmon.solaris64
% ln -s $SGE_ROOT/bin/solaris/qmon $SGE_ROOT/bin/solaris64/qmon
% mv $SGE_ROOT/lib/solaris64/libXltree.so $SGE_ROOT/lib/solaris64/libXltree.so.solaris64
% ln -s $SGE_ROOT/lib/solaris/libXltree.so $SGE_ROOT/lib/solaris64/libXltree.so
```

Issue – When you install a localized version of this product with the `-csp` switch, part of the messages are not localized because the non-localized external binary, `openssl`, is used.

Workaround – No workaround exists for this problem.

Corrections and Additions to man Pages

The following section contains corrections and additions to the man pages that accompany this product. Note that these additions and corrections also apply to the *Sun ONE Grid Engine 5.3 and Sun ONE Grid Engine, Enterprise Edition 5.3 Reference Manual*.

`sge_conf(5)` and `queue_conf(5)`

- Add the following sentence to the descriptions of both entries:

The procedure's standard output and the error output stream are written to the same file used also for the standard output and error output of each job.

`access_list(5)`

- Replace with the following text:

Grid Engine File Formats `ACCESS_LIST(5)`

NAME

`access_list` - Sun ONE Grid Engine access list file format

DESCRIPTION

Access lists are used in Sun ONE Grid Engine to define access permissions of users to queues (see `queue_conf(5)`) or parallel environments (see `sge_pe(5)`). A list of currently configured access lists can be displayed via the `qconf(1) -su1` option. The contents of each enlisted access list can shown via the `-su` switch. The output follows the `access_list` format description. New access lists can be created and existing can be modified via the `-au` and `-du` options to `qconf(1)`.

Sun ONE Grid Engine, Enterprise Edition departments are a special form of access list and allow in addition assignment of functional shares and override tickets.

FORMAT

The following list of `access_list` parameters specifies the `access_list` content.

name

The name of the access list.

type

This parameter is only available in a Sun ONE Grid Engine, Enterprise Edition system. Sun ONE Grid Engine does not support this parameter. The type of the access list, currently one of `ACL` or `DEPT`, or a combination of both in a comma-separated list. Depending on this parameter, the access list can be used as access list only or as a department.

oticket

This parameter is only available in a Sun ONE Grid Engine, Enterprise Edition system. Sun ONE Grid Engine does not support this parameter. The amount of override tickets currently assigned to the department.

fshare

This parameter is only available in a Sun ONE Grid Engine, Enterprise Edition system. Sun ONE Grid Engine does not support this parameter. The current functional share of the department.

entries

The `entries` parameter contains the comma-separated list of those UNIX users or UNIX user groups that are assigned to the access list or the Sun ONE Grid Engine, Enterprise Edition department. Only symbolic names are allowed. A group is differentiated from a user name by prefixing the group name with a `@` sign. Pure access lists allow enlisting any user or group in any access list.

With Sun ONE Grid Engine, Enterprise Edition departments, each user or group enlisted may be enlisted in only one department, to ensure a unique assignment of jobs to departments. To jobs whose users do not match with any of the users or groups enlisted under `entries`, the `defaultdepartment` is assigned, if existing.

SEE ALSO

`sge_intro(1)`, `qconf(1)`, `sge___pe(5)`, `queue___conf(5)`

COPYRIGHT

See `sge_intro(1)` for a full statement of rights and permissions.

sge_aliases

- Replace corresponding text with the following corrections.

FILES

`<sge_root>/<cell>/common/sge_aliases` - global aliases file

`$HOME/.sge_aliases` - user local aliases file

SEE ALSO

`sge_intro(1)`, `qsub(1)`

COPYRIGHT

See `sge_intro(1)` for a full statement of rights and permissions.

qmod

- Add the following descriptions to the `qmod` entry.

`-e/-d/-us/-s *`

When one of the above command-line switches is used to specify a queue (note that `-s` and `-us` can also be used for jobs), the following wildcards can be used.

* (asterisk)

? (question mark)

The syntax for these wildcards follows the rules for wildcards in a UNIX shell (this means it does not follow the rules for a regular expression).

`-r`

Only an admin user may use `qmod -r` to force rescheduling of a job.

Array Job Task Index Values Range

The minimum number of a range is 1. The maximum number $2^{32}-1$.

The number of submitted array job tasks may not exceed the range of the `max_aj_tasks` parameter of the global cluster configuration.

qmake(1)

- Add the following description of the `ARCH` environment variable.

If no resource request (`QS_NAME` command-line option `-l`) is specified, `qmake` will use the environment variable `ARCH` to request the same architecture for task execution as has the submit host. If `ARCH` is set, the architecture specified in `ARCH` will be requested by inserting the option `-l arch=$ARCH` into the command line options. If `ARCH` is not set, the `make` tasks can be executed on any available architecture. As this is critical for typical `make` (compile) jobs, a warning will be output.

ARCH

The architecture of the submit host. If this variable is set in the submission environment, `qmake` will request the given architecture for job execution (see `DESCRIPTION` above).

-

