



N1 Grid Engine 6 Release Notes

For Update 4

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

Part No: 817-5678-20
May 2005

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

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

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

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

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

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

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

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

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

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

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



050523@11223



Contents

N1 Grid Engine 6 Update 4 Software Release Notes	5
Accessing Documentation	5
Contents of This Software Package	6
Installing the N1 Grid Engine 6 Software	7
Changes in N1 Grid Engine 6 Update 4 Software	8
Support for Microsoft Windows Operating Systems	8
New Grid Engine Management Module (GEMM) for Sun Control Station	9
Support for Solaris 10 x64 (Solaris on Opteron systems 64-bit)	9
Other Functionality Delivered With This Update Release	10
New Features in N1 Grid Engine 6 Software	11
Accounting and Reporting Console (ARCo)	11
Resource Reservation	11
Cluster Queues	12
DRMAA	12
Scalability	13
Scheduler Enhancements	13
Automated Installation and Backup	13
qping Utility	13
Starting Binaries Directly	13
Resource Requests for Individual make Rules	14
Grid Engine System Binary Directory	14
Known Limitations and Workarounds	14
Known Limitations of N1 Grid Engine 6 Software	14
Known Limitations and Workarounds for the Microsoft Windows Platform	16
Known Limitations and Workarounds for GEMM	17

N1 Grid Engine 6 Update 4 Software Release Notes

- “Accessing Documentation” on page 5
- “Contents of This Software Package” on page 6
- “Installing the N1 Grid Engine 6 Software” on page 7
- “Changes in N1 Grid Engine 6 Update 4 Software” on page 8
- “New Features in N1 Grid Engine 6 Software” on page 11
- “Known Limitations and Workarounds” on page 14

Accessing Documentation

The distribution CD includes full documentation for a networked set of computer hosts that run N1™ Grid Engine 6 Update 4 software:

- N1GE6Update4_User_Guide.pdf – *N1 Grid Engine 6 User’s Guide*
- N1GE6Update4_Administration_Guide.pdf – *N1 Grid Engine 6 Administration Guide*
- N1GE6Update4_Installation_Guide.pdf – *N1 Grid Engine 6 Installation Guide*
- N1GE6Update4_GEMM_Guide.pdf – *Grid Engine Management Module for SCS 2.2*

You can access these files directly from the CD, in either PDF or HTML formats. The files are in the *cdmountpoint/N1_Grid_Engine_6u4/Docs* directory.

Contents of This Software Package

The Grid Engine 6 Update 4 software distribution is made up of the following components:

- The grid engine software binary packages, including all daemons, client programs, and libraries. You must load and install one binary package for each operating system architecture you intend to use.
- The grid engine software common package, containing install scripts, and other architecture-independent utilities.
- The grid engine software documentation package, containing these release notes, the installation guide, user guide and administration guide in PDF and HTML formats.
- The optional Accounting and Reporting Console (ARCo) software, which is made up of three separate packages:
 - The Sun Web Console package. You must select the package appropriate for the operating system architecture on which you plan to run the web console server.
 - The `dbwriter` package, written in Java and therefore available in only one version.
 - The ARCo module package, usable across different supported architectures.

Note – In order to operate ARCo, you also must set up a PostgreSQL or an Oracle database server . PostgreSQL and Oracle are not included in the N1 Grid Engine 6 software distribution. For more information, see Chapter 8, “Installing the Accounting and Reporting Console,” in *N1 Grid Engine 6 Installation Guide*.

- The optional Grid Engine Management Module for Sun Control Station package.

The Grid Engine 6 Update 4 software distribution kit contains the following top-level directory hierarchy:

- `3rd_party` – Contains information about freeware, public domain, and public license software
- `bin` – Grid engine software executables
- `catman` – Online manual pages organized into admin and user commands
- `ckpt` – Sample checkpointing configurations
- `dbwriter` – DbWriter software used by the accounting and reporting console
- `doc` – Documentation in PDF and HTML formats
- `examples` – Sample script files, configuration files, and application programs
- `include` – DRMAA header file

- `lib` – Required shared libraries and DRMAA Java™ binding jar file
- `man` – Online manual pages in `nroff` format
- `mpi` – A sample parallel environment interface for the MPI message-passing system
- `pvm` – A sample parallel environment interface for the PVM message-passing system
- `qmon` – Pixmaps, resource, and help files for `QMON`, the graphical user interface
- `reporting` – Accounting and reporting console software
- `util` – Some utility shell procedures used for installation tasks and some template grid engine system shutdown and boot scripts
- `utilbin` – Some utility programs that are mainly required during the installation

Installing the N1 Grid Engine 6 Software

If you plan to add a new installation of the N1 Grid Engine 6 software or if you are just adding new packages to your N1 Grid Engine cluster installation (like Windows support, ARCo, or GEMM which have not been installed previously), see `N1GE6Update4_Installation_Guide.pdf`, the N1 Grid Engine 6 Installation Guide that is included on the distribution CD. If you already installed the N1 Grid Engine 6 software packages, you should install the patches which are available on <http://sunsolve.sun.com>. See the directions in the patch README documents on how to install the patches.

The patch matrix below lists the available patches for N1 Grid Engine 6 which are currently available (May 2005). Newer revisions of the patches or additional patches may be available at a later time. Please check <http://sunsolve.sun.com> for the availability of N1 Grid Engine 6 patches.

TABLE 1 Patches For Packages in Sun `pkgadd` Format

Package Name ¹	Operating System	Architecture ²	Patch-Id
SUNWsgee	Solaris Sparc, 32bit	sol-sparc	118094-04
SUNWsgeex	Solaris, Sparc, 64bit	sol-sparc64	118130-04
SUNWsgeei	Solaris x86	sol-x86	118131-04
SUNWsgeec	all	common	118132-04
SUNWsgeea	all	arco	118133-04
SUNWsgeed	all	doc	119846-01

1. See `pkginfo(1)`
2. N1 Grid Engine binary architecture string or *common* (Architecture independent packages), *arco* (Accounting and Reporting Console), and *doc* (Documentation)

TABLE 2 Patches for Packages in tar . gz Format

Operating System	Architecture	Patch-ID
Solaris, Sparc, 32bit	sol-sparc	118082-04
Solaris, Sparc, 64bit	sol-sparc64	118083-04
Solaris, x86	sol-x86	118084-04
Linux kernel 2.4/2.6	x86, lx24-x86	118085-04
Linux kernel 2.4/2.6	AMD64 lx24-amd64	118086-04
IBM AIX 4.3	aix43	118087-04
IBM AIX 5.1	aix51	118088-04
Apple MAC OS/X	darwin	118089-04
HP-UX 11	hp11	118090-04
SGI Irix 6.5	irix65	118091-04
all	common	118092-04
all	arco	118093-04
all	doc	119861-01

Changes in N1 Grid Engine 6 Update 4 Software

Along with many bug fixes, N1 Grid Engine 6 Update 4 includes the following changes.

Support for Microsoft Windows Operating Systems

N1 Grid Engine 6 Update 4 Windows client functionality (submit, administration and execution host) is now available for Microsoft Windows 2000 SP3 (or higher), Windows XP Professional SP1 (or higher) and Windows Server 2003. The N1 Grid Engine command line tools and execution host functionality are almost fully supported on these operating systems.

The support of N1 Grid Engine for Windows allows users to fully integrate Windows hosts into an existing N1 Grid Engine environment. Users are able to submit and monitor their jobs through the command line tools. Administrators can have full

control over a N1 Grid Engine cluster from a Windows host. The execution host functionality allows you to use Windows desktop machines and dedicated Windows compute servers for the execution of batch workload and interactive jobs.

Installation of N1 Grid Engine 6 U4 requires Microsoft Services For UNIX (SFU) 3.5 which provides tools and libraries to integrate Windows with UNIX. SFU 3.5 is available for no license fee and is supported by Microsoft. See <http://www.microsoft.com/windows/sfu/default.asp> for information and requirements about SFU as well as how to get SFU.

New Grid Engine Management Module (GEMM) for Sun Control Station

GEMM is a new addition to N1GE6 which provides a web-based interface for deployment, monitoring, and diagnostics of an N1 Grid Engine installation. It operates in the framework provided by Sun Control Station 2.2, a product which must be purchased separately.

Sun Control Station (SCS) 2.2 provides overall life-cycle management of servers, from bare-metal OS provisioning, to software and patch deployment, to basic health, inventory, and hardware monitoring, all in an easy-to-user web interface. GEMM adds to this the following capabilities:

- Deploy N1GE — Install and configure N1GE software on grid hosts, including a master host, compute hosts, and access hosts. Key features include:
 - deploy any supported version of N1GE. Initially, N1GE6u4 is supported. Future versions will be qualified for support.
 - work with a previously-installed master host.
- Monitoring — Provide high level monitoring of N1GE jobs, queues, and hosts. Key features include:
 - drill down for details on jobs, queues, and hosts
 - selectively filter job display to focus on jobs of interest
 - provide monitoring even for Grid hosts outside the SCS/GEMM framework.
- Diagnostics — Provide tools for doing first-level diagnostics of N1GE problems within the web interface. Key features include:
 - display job scheduling information
 - view spool files for running jobs
 - view messages files for `qmaster` and `execd` daemons.

Support for Solaris 10 x64 (Solaris on Opteron systems 64-bit)

Solaris 10_x64 (on AMD Opteron hardware) is now fully supported with this release.

Other Functionality Delivered With This Update Release

This list summarizes new and improved functionality which has been added to the N1 Grid Engine 6 software since it was released in June 2004.

- The Accounting and Reporting Console (ARCo) now uses the GUI elements of the Sun Web Console. This update improves the look and feel as well as the scalability of the web interface of ARCo.
- The scalability, submit rate, scheduling speed, status query speed of `qstat`, job turnaround times, and PE job start have been significantly improved for many typical use cases in comparison to the original N1 Grid Engine 6 release.
- The execution daemon installation is supported in Solaris 10 Containers (Zones).
- The DRMAA Java language binding is now available. The DRMAA Java language binding library is located in the `<sgc_root>/lib/drmaa.jar` directory. The documentation is in `<sgc_root>/doc/javadoes`.
- The `qping` utility has been significantly improved to diagnose N1 Grid Engine daemon communication. See the `qping(1)` man page for more information.
- The auto installation process has been improved. An auto installation and de-installation of the Berkeley DB (BDB) RPC server is now possible. The backup and restore procedure is now supported for the classic spooling option and the Berkeley DB RPC spooling option.
- Berkeley DB spooling on NFSv4 under Solaris 10 is supported.
- The BDB database now can be installed on a NFSv4 mounted filesystem on Solaris 10.

For performance reasons, it is recommended that you use NFSv4 BDB spooling only when the NFSv4 mount provides an excellent high speed connection to the file server.

- On Linux platforms, an LSB conforming "lock" file is created by the daemon startup script.
- New man pages for the utility binaries `gethostbyaddr`, `gethostbyname`, `gethostname`, `getservbyname` and `qping` have been added. The man page `sgc_h_aliases(5)` has been renamed to `host_aliases(5)`.
- It is now possible to avoid inheritance of the execution daemon environment to the job environment. It is also possible to avoid inheriting the variable `LD_LIBRARY_PATH` from the execution daemon environment (this feature can solve NFS issues in certain setups). See `sgc_conf(5)` for more information.
- New options to optimize the memory overhead and speed of `qstat`. See `sgc_qstat(5)` and `qstat(1)` for more information (`-u` and `-s` flags).
- The new `qconf -purge` switch easily allows the removal of all references to a host or hostgroup in a cluster queue. See `qconf(1)` under the `-purge` switch for more information.

- The spooling interval for the sharetree usage can now be configured to reduce I/O and improve performance on NFS mounted `qmaster` spool directories. See `sgc_conf(5)` for more information in the section about the `STREE_SPOOL_INTERVAL` parameter.
- Faster execution daemon reconnect in the Certificate Security Protocol (CSP) installation mode.
- Execution daemons now can reconnect faster to `qmaster` if the execution daemon or `qmaster` daemon has been restarted.

New Features in N1 Grid Engine 6 Software

The original N1 Grid Engine 6 provides the following new features.

Accounting and Reporting Console (ARCo)

The optional ARCo enables you to gather live accounting and reporting data from a grid engine system and store the data in a standard SQL database. ARCo also provides a web-based tool for generating information queries on that database and for retrieving the results in tabular or graphical form. ARCo enables you to store queries for later use, to run predefined queries, and to run queries in batch mode, for example, overnight.

For details, see Chapter 5, “Accounting and Reporting,” in *N1 Grid Engine 6 User’s Guide*, and Chapter 8, “Installing the Accounting and Reporting Console,” in *N1 Grid Engine 6 Installation Guide*.

Resource Reservation

The grid engine system scheduler supports a highly flexible resource reservation scheme. Jobs can reserve resources depending on criteria such as resource requirements, priority, waiting time, resource sharing entitlements, and so forth. The scheduler enforces reservations in such a way that jobs with highest urgency receive the earliest possible resource assignment. Resource reservation completely avoids well-known problems such as job starvation.

With respect to resource requirements, a job’s importance can be defined on a per resource basis for arbitrary resources, as well as for administrator-defined resources such as third party licenses or network bandwidth. Reservations can be assigned across the full hierarchy of grid engine system resource containers: global, host, or queue.

For more information, see the `sge_priority(5)` man page.

Cluster Queues

N1 Grid Engine 6 software provides a new administrative concept for managing queues. It enables easier administration while maintaining the flexibility of the Sun Grid Engine 5.3 queue concept.

A *cluster queue* can extend across multiple hosts. Those hosts can be specified as a list of individual hosts, as a host group, or as a list of individual hosts and host groups. By adding a host to a cluster queue, the host receives an instance of that cluster queue. A *queue instance* corresponds to a queue in Sun Grid Engine 5.3.

When you modify a cluster queue, all of its queue instances are modified simultaneously. Even within a single cluster queue, you can specify differences in the configuration of queue instances, depending on individual hosts or host groups. Therefore, a typical N1 Grid Engine 6 software setup will have only a few cluster queues, and the queue instances controlled by those cluster queues remain largely in the background.

For further details, see the `queue_conf(5)` man page.

DRMAA

N1 Grid Engine 6 software includes a standard-compliant implementation of the Distributed Resource Management Application API (DRMAA), version 1.0. DRMAA 1.0 is a standard draft for review at Global Grid Forum. It provides a standard API for the integration of applications with Distributed Resource Management System, such as N1 Grid Engine 6 software, with external applications like ISV codes or graphical interfaces. Major functions provided by DRMAA include job submission, job monitoring, and job control. N1 Grid Engine 6 software includes an implementation for the C-language binding of DRMAA. Details are available in the `drmaa_*(3)` man pages and on the DRMAA home page <http://www.drmaa.org/>.

Scalability

N1 Grid Engine 6 software implements a number of architectural changes from previous releases in order to support increased scalability:

- Spooling of persistent status information for the `sge_qmaster` can now be done using the high-performance Berkeley DB database instead of the previous file-based spooling.
- The `sge_qmaster` is multithreaded to support concurrent execution of tasks on multi-CPU systems.
- The Sun Grid Engine 5.3 communication system has been replaced. The communication system is now multithreaded and no longer requires a separate communication daemon.

Scheduler Enhancements

Different scheduling profiles can be selected for setups ranging from high throughput and low scheduling overhead to full policy control. The setups can be selected during the `sge_qmaster` installation procedure. In addition, a series of enhancements has improved scheduler performance greatly.

Automated Installation and Backup

The N1 Grid Engine 6 software installation procedure can be completely automated to facilitate installation on large numbers of execution hosts, frequently recurring reinstallation of hosts, or integration of the installation process into system management frameworks. For more information, see the file `doc/README-Autoinstall.txt`.

N1 Grid Engine 6 software also includes an automatic backup script that backs up all cluster configuration files.

qping Utility

A new `qping` utility enables you to query the status of the `sge_qmaster` and `sge_execd` daemons.

Starting Binaries Directly

The `qsub` command now supports the `-shell {y | n}` option, which is used with the `-b y` option, to start a submitted binary directly without an intermediate shell.

Resource Requests for Individual make Rules

In dynamic allocation mode, the `qmake` command can now specify resource requests for individual make rules.

Grid Engine System Binary Directory

The environment variable `SGE_BINARY_PATH` is set in the job environment. This variable points to the directory where the grid engine system binaries are installed.

Known Limitations and Workarounds

The following sections contain information about product irregularities discovered during testing, but too late to fix or document.

Known Limitations of N1 Grid Engine 6 Software

This N1 Grid Engine 6 software release has the following limitations:

- The stack size for `sge_qmaster` should be set to 16 MBytes. `sge_qmaster` might not run with the default values for stack size on the following architectures: IBM AIX 4.3 and 5.1, and HP UX 11.
- You should set a high file descriptor limit in the kernel configuration on hosts that are designated to run the `sge_qmaster` daemon. You might want to set a high file descriptor limit on the shadow master hosts as well. A large number of available file descriptors enables the communication system to keep connections open instead of having to constantly close and reopen them. If you have many execution hosts, a high file descriptor limit significantly improves performance. Set the file descriptor limit to a number that is higher than the number of intended execution hosts. You should also make room for concurrent client requests, in particular for jobs submitted with `qsub -sync` or when you are running DRMAA sessions that maintain a steady communication connection with the master daemon. Refer to your operating system documentation for information about how to set the file descriptor limit.
- The number of concurrent dynamic event clients is limited by the number of file descriptors. The default is 99. Dynamic event clients are jobs submitted with the `qsub -sync` command and a DRMAA session. You can limit the number of dynamic event clients with the `qmaster_params` global cluster configuration setting. Set this parameter to `MAX_DYN_EC=n`. See the `sge_conf(5)` man page for

more information.

- The ARCo module is available only for the Solaris Sparc, Solaris Sparc 64 bit, Solaris x86, Solaris x64, Linux x86, and Linux 64 bit kernels.
- ARCo currently supports only the following database servers: PostgreSQL 7.3.2, 7.4.1, 7.4.2, and Oracle 9i. Postgres 8.0.1 has been successfully tested on Solaris. An integration with MySQL will be provided once MySQL supports views.
- Only a limited set of predefined queries is currently shipped with ARCo. Later releases will include more comprehensive sets of predefined queries.
- Jobs requesting the amount INFINITY for resources are not handled correctly with respect to resource reservation. INFINITY might be requested by default in case no explicit request for a certain resource has been made. Therefore it is important to request that all resources be explicitly taken into account for resource reservation.
- Resource reservation currently takes only pending jobs into account. Consequently, jobs that are in a hold state due to the submit options `-a time` and `-hold_jid joblist`, and are thus not pending, do not get reservations. Such jobs are treated as if the `-R n` submit option were specified for them.
- Berkeley DB requires that the database files reside on the local disk, if `qmaster` is not running on Solaris 10 and uses a NFSv4 mount (full NFSv4 compliant clients and servers from other vendors are also supported, but have not yet been tested.) If the `sgc_qmaster` cannot be run on the file server intended to store the spooling data (for example, if you want to use the shadow master facility), a Berkeley DB RPC server can be used. The RPC server runs on the file server and connects with the Berkeley DB `sgc_qmaster` instance. However, Berkeley DB's RPC server uses an insecure protocol for this communication and so it presents a security problem. Do *not* use the RPC server method if you are concerned about security at your site. Use `sgc_qmaster` local disks for spooling instead and, for fail-over, use a high availability solution such as Sun Cluster, which maintains host local file access in the fail-over case.
- Busy QMON with large array task numbers. If large array task numbers are used, you should use "compact job array display" in the QMON Job Control dialog box customization. Otherwise the QMON GUI will cause high CPU load and show poor performance.
- The automatic installation option does not provide full diagnostic information in case of installation failures. If the installation process aborts, check for the presence and the contents of an installation log file `/tmp/install.pid`.
- On IBM AIX 4.3 and 5.1, HP/UX 11, and SGI IRIX 6.5 systems, two different binaries are provided for `sgc_qmaster`, `spooldefaults`, and `spoolinit`. One of these binaries is for the Berkeley DB spooling method, the other binary is for the classic spooling method. The names of these binaries are `binary.spool_db` and `binary.spool_classic`.

To change to the desired spooling method, modify three symbolic links before you install the master host. Do the following:

```
# cd sgc-root/bin/arch
# rm sgc_qmaster
```

```
# ln -s sge_qmaster.spool_classic sge_qmaster

# cd sge-root/utilbin/arch
# rm spooldefaults spoolinit
# ln -s spooldefaults.spool_classic spooldefaults
#ln -s spoolinit.spool_classic spoolinit
```

- Gathering of online usage statistics for running jobs, and dynamic reprioritization for such jobs, does not work on the following operating systems:
 - IBM AIX
 - HP/UX
 - Mac OS X

For a workaround, see the `sge_conf(5)` man page for information about how to adjust the execution host parameters `ACCT_RESERVED_USAGE` and `SHARETREE_RESERVED_USAGE`.

- PDF export in ARCo requires a lot of memory. Huge reports can result in a `OutOfMemoryException` when they are exported into PDF.

Workaround — Increase the JVM heap size for the Sun Web Console The following command the set max. heap size to 512 MB.

```
# smreg add -p java.options "... -mx512M ....."
```

A restart of the Sun Web Console is necessary to make the change effective as in this command:

```
# smcwebserver restart
```

- For DBWriter (part of ARCo) the 64-Bit support of the Java virtual machine needs to be installed on Solaris Sparc 64-bit and Solaris x64, and Linux 64-bit kernels.

Known Limitations and Workarounds for the Microsoft Windows Platform

- The installation of Services For UNIX 3.5 requires a good administrative understanding of the Windows platform and it's integration into a UNIX environment. Appendix A of the *N1 Grid Engine 6 Update 4 Installation Guide* provides an Overview of SFU. Further technical information and documentation about SFU can be found on the Microsoft web server at <http://www.microsoft.com/windows/sfu/default.asp>.
Username mapping, NFS mounts, and hostname resolving in SFU require special attention to successfully install the Grid Engine execution daemon, submit host functionality, and integration of Windows hosts into a N1 Grid Engine cluster.
- Installation of a Windows host in a N1 Grid Engine cluster with CSP mode enabled is currently not yet supported because it is untested. Full support for this functionality will be provided in June.

- The execution daemon cannot start jobs of users who are members of a Windows Domain. A fix for this problem is currently in progress and will be delivered in June 2005.
- On machines with more than one processor or Hyperthreading processors, only one processor is detected by the load sensor. This does not limit the use of multiple CPU's by jobs running under N1 Grid Engine on Windows. If a queue instance is created during the installation of a Windows execution host, the number of queue slots for this host can be changed after the installation.
- A Windows execution host cannot be installed with the auto installation procedure.

Known Limitations and Workarounds for GEMM

- If the system time on the N1GE Master Host and the SCS Server are not the same, the job times given in the Monitor will not be correct. In particular, if the time on the Master Host is ahead of the SCS Server, then it would be possible to see negative jobs times. Please ensure that the system time on the N1GE Master Host and the SCS Server are the same.
- A previous version of N1 Grid Engine 6, sold under the name "Sun Grid Engine Enterprise Edition 5.3", had an optional module that integrated with Sun Control Station 2.0 (the previous version of Sun Control Station 2.2). Customers using SGEEE 5.3 with SCS 2.0 who wish to upgrade to N1GE 6 with SCS 2.2 and GEMM should note that there is no automated upgrade procedure. You will need to uninstall SCS 2.0 manually from all the grid hosts, and then perform a fresh installation of SCS 2.2, followed by the installation of GEMM and subsequently N1GE 6.

