



Grid Engine Management Module

For Sun Control Station 2.2

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

Part No: 817-3606
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, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

Copyright 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 États-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, de Sun Microsystems, Inc. aux États-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 États-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.



050502@11223



Preface

This *Grid Engine Management Module for Sun Control Station 2.2 User's Guide* provides users with all the information they need to successfully install and use GEMM software. This manual is part of the Grid Engine 6 Update 4 distribution.

Who Should Use This Book

The document is intended for users familiar with the Sun Control Station product and the N1 Grid Engine 6 product. The user of this document should be the person administering the software.

How This Book Is Organized

Chapter 1 provides an detailed description of the GEMM product.

Chapter 2 contains some command line equivalents of the functionality provided by the GEMM GUI.

Chapter 3 Explains the contents of the `setup.conf` file.

Accessing Sun Documentation Online

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

Ordering Sun Documentation

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <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:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

TABLE P-1 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . Perform a <i>patch analysis</i> . Do <i>not</i> save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

Grid Engine Management Module

This manual contains information about accessing and using GEMM as well as managing N1GE versions.

Using GEMM

The Grid Engine Management Module (GEMM) for the Sun™ Control Station allows you to install and set up a grid. It also allows you to monitor the performance of the hosts in the grid. This document explains the features and services available through the Grid Engine control module.

GEMM supports the following operating systems and hardware platforms:

- Solaris 9 and 10 on SPARC, x86, and x64 hardware platforms
- RedHatLinux 7.3, RedHatLinux 8.0, RedHatLinux 9 on x86 hardware platforms
- FedoraLinux 1, FedoraLinux 2, FedoraLinux 3 on x86, and x64 hardware platforms
- RedHatLinux 2.1 WS, RedHatLinux 2.1 AS, RedHatLinux 2.1 ES on x86 hardware platforms
- RedHatLinux 3WS, RedHatLinux 3 AS, RedHatLinux 3ES on x86, and x64 hardware platforms
- SuSELinux 9.0 on x86 and x64 hardware platforms
- JDS 1, JDS 2 on x86 hardware

The GEMM module allows you to:

- manage different versions of N1 Grid Engine
- install a master host on the grid
- install additional compute hosts on the grid

- monitor and diagnose the performance of the grid
- uninstall the Grid Engine components from selected compute and access hosts
- uninstall the Grid Engine components from the master host and all compute and access hosts
- check the status of Grid Engine (from Station Settings > Active Monitor on the SCS.)
- configure the monitoring settings for the grid

The following sections describe each of these functions.

Note – In Grid Engine terminology, compute hosts are called execution hosts. Access hosts are called submit hosts.

Installing GEMM

GEMM is part of the Grid Engine Update 4 distribution and is included in the `Gemm/tar/n1ge-6_0u4-gemm.tar.gz` package. Use the following steps to install GEMM

1. Unpack the GEMM archive file.
2. Add GEMM to a Sun Control Station 2.2 system. See the Sun Control Station documentation for instructions on adding a new management module.

Accessing GEMM

You access the GEMM features by clicking on the Monitor menu item on the Sun Control Station main screen as shown in the following figure.

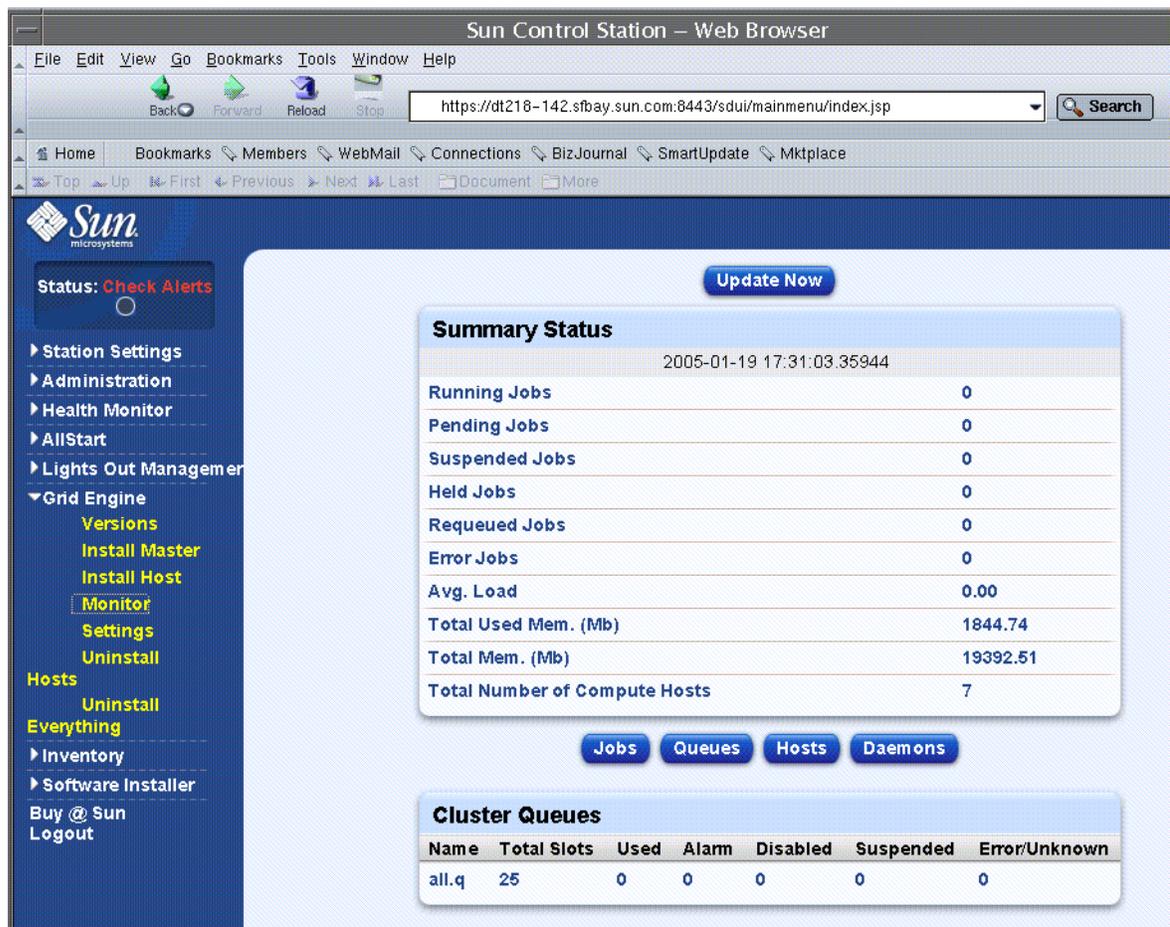


FIGURE 1-1 Monitor Main Page

Note – In most of the short procedures in this chapter, the first step is to click the Grid Engine item in the left menu bar and the second step is to click on a sub-menu item. To reduce the number of steps in each procedure, the menu commands are grouped together and shown in Initial Caps. Right-angle brackets separate the individual items. For example, select Grid Engine > Settings means to click Grid Engine in the left menu bar and then click the Settings sub-menu item.

Task Progress Dialog

When you launch a task (for example, when installing a master host or uninstalling a host), a Task Progress dialog appears in the user interface (UI). This dialog has a Status field indicating the current status of the task and a progress bar. When the progress bar displays 100%, the task has completed.

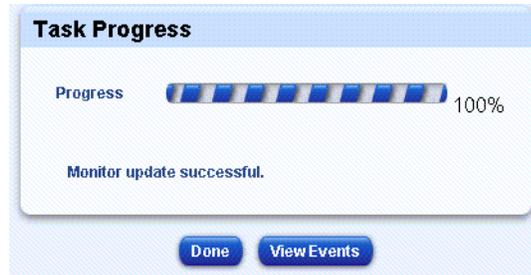


FIGURE 1-2 Task Progress Dialog

If you want to perform another task in the UI while the current task is underway, you can put the Task Progress dialog in the background. Simply click the Run Task In Background button located below the progress bar.

To return to the Task Progress dialog, select `Administration > Tasks` on the left. The Task table appears. If the task is still underway, a status message displays in the Duration column. Click on the *progress-bar* icon in this column to re-display the Task Progress dialog for this task.

Once the task is complete and the progress bar displays 100%, two buttons appear below the Task Progress dialog: Done and View Events.

- To view the list of events associated with the completed task, click `View Events`. The `Events For <Task>` table appears. If you then click the *up-arrow* icon in the top-right corner, the Tasks table appears.
- To return to the previous screen, click Done.

Managing Versions

GEMM allows you to upload one or more versions of N1 Grid Engine, and choose which one to deploy on the grid. To do these tasks, click the `Versions` menu item to display the Version Management page.

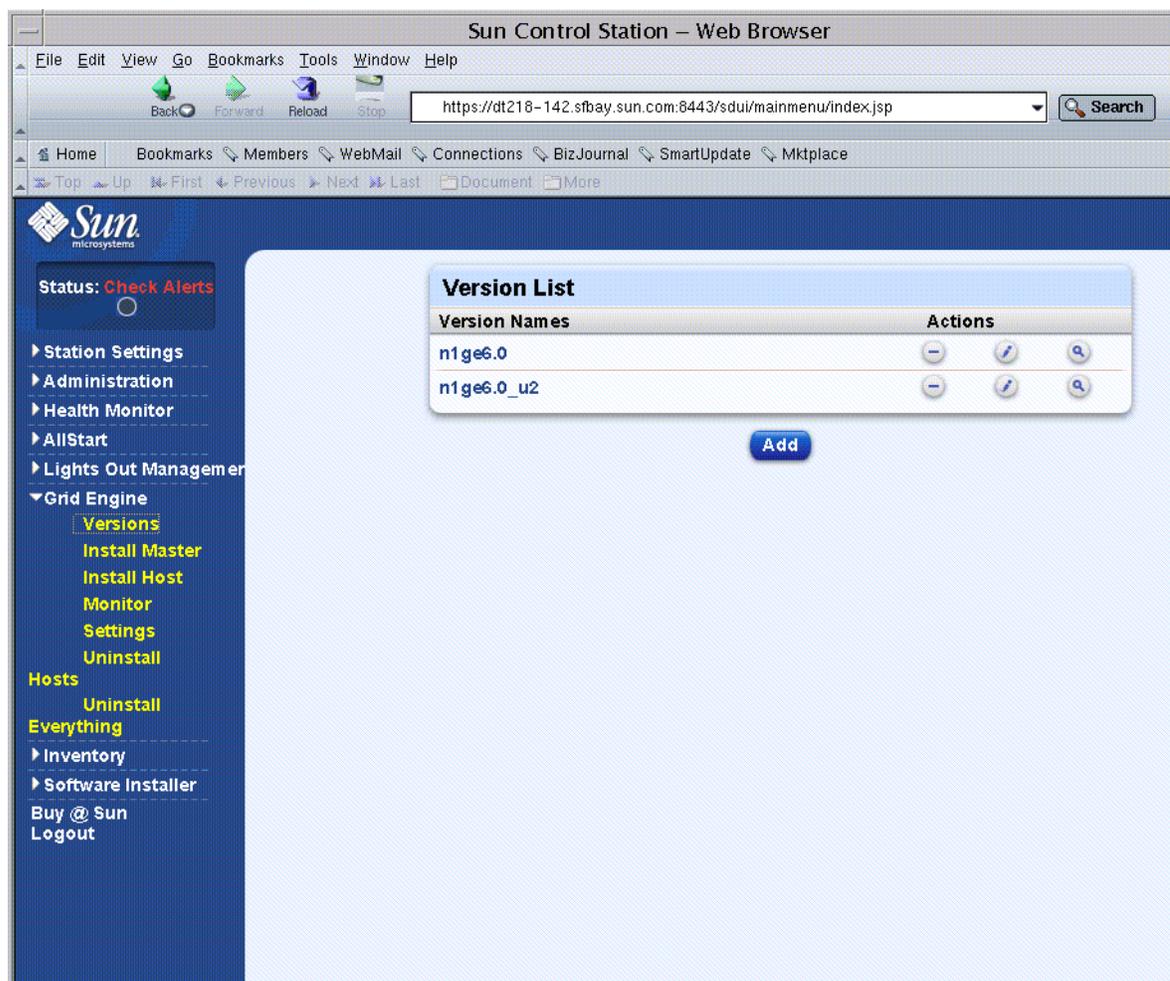


FIGURE 1-3 Version List Page

This page is where you can upload, modify and manage different versions of N1 Grid Engine software.

Note – You can only deploy one version at any given time. If you wish to deploy another version, you must first uninstall the grid from all the hosts.

Versions List Icons

In the Version list, each version has three icons:

- The Minus icon removes a version and all its files.

- The Modify icon lets you rename a version.
- The Inspect icon lets you add or remove individual files from a version.

Adding Grid Engine Versions

The Version Management main page displays a list of versions currently available on the SCS server. Initially, no versions are defined.

▼ To define a version

- Steps**
1. **Click the Add button which produces a dialog where you name a version.**
This version name can be anything, as long as it does not contain any non-whitespace or punctuation characters other than “-” and “_”.
 2. **After you name the version, click the Submit button.**
Once you submit the name, the Version list displays again with the newly-created version present. You can add more versions at any time.



FIGURE 1-4 Add Version Dialog

Adding Files to a Version

You must first add files to a version before you deploy it to the grid. Adding files consists of adding all N1 Grid Engine package files that are part of the given version.

Package File Criteria

The following criteria apply to the package files:

- GEMM requires N1 Grid Engine 6 Update 4 or later.
- All package files must be in `.tar.gz` format. Although N1 Grid Engine currently is made available in `.pkg` format for Solaris, the `.pkg` format files cannot be used by GEMM.
- For any given version, there must be the “common” package for that version, as well as all the “bin” (binary) packages which support the kinds of hosts in your grid. For example, if your grid consists of Solaris 9 SPARC hosts as well as Solaris

10 and x64 hosts, then you must include in the version these files:

- `<name>-bin-solaris-sparcv9.tar.gz`
- `<name>-bin-solaris-x64.tar.gz`
- `<name>-common.tar.gz`

where `-<name>` is the name given to that version, for example `n1ge-6_0`.

- In any given version, you cannot mix different update levels of N1 Grid Engine. All packages associated with a version must belong to the same update level, for example `n1ge-6_0`. The only exception is when you deploy N1GE 6 patches, which will be described ***.

▼ To Add Files

- Steps**
1. click the **Inspect** icon in the version list. You will see a list of files currently contained in that version.
 2. Clicking the **Add** button produces a dialog box where you can load version files one at a time. You can also upload files from the local browser using the **File** browser or from a remote URL.

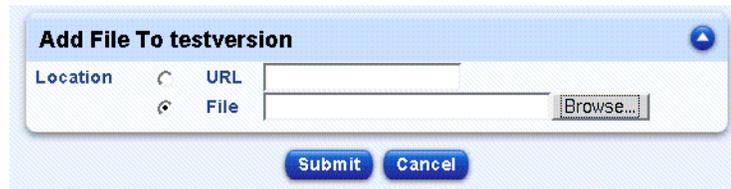


FIGURE 1-5 Add Files Dialog

Note – When you upload files from a remote URL, you can only specify a URL which can be accessed from the SCS server directly without going through a proxy server. You cannot specify a proxy server when using the Version Management web dialog. Please see the documentation for the command-line equivalent of Version Management, `gemmVersionMgmt.pl`, to learn how to upload files using a web proxy.

▼ To Remove Files

- Steps**
1. Select a file from the list of files.
 2. Click on the **Minus** icon for that file.

Deploying a Patched Version of N1GE 6 With GEMM

N1GE software updates are made available through the mechanism of patch files. You cannot use an N1GE patch alone; you must use it in conjunction with a full distribution of N1GE software. When you install a patch, it replaces various files in the existing full version.

There are two ways you can install N1GE patches:

- Install patch files on a live N1GE grid already running an existing full installation of N1GE 6. This procedure is described in the patch documentation but is *not* supported by GEMM.
- Install patch files at the same time as you install a fresh installation of an original, full, N1GE software distribution. You can use this technique when you are creating a new grid and want to install it with the latest N1GE updates. You also can use this approach when you want to use N1GE with the latest updates and don't mind getting rid of your old setup entirely (without worrying about saving old configurations or maintaining jobs currently in the systems). GEMM can handle this procedure automatically as described here.

1. Create a new version in the Version Manager of GEMM. Populate this version with N1GE files from an original full version, just as if you were going to deploy this version.
2. Get the desired patch files. When Sun Microsystems creates and releases patch updates, these files are made available on the SunSolve website (<http://sunsolve.sun.com> (<http://sunsolve.sun.com>)). For each patch release, there is one patch for the N1GE "common" package, as well as one patch for each architecture-specific package. Get all the patch files necessary for your particular environment.

Patch files are distributed in both `.pkg` format as well as `.tar.gz` format. Make sure to obtain only the `.tar.gz` form of the patches. These patch files are themselves contained in a ZIP archive; be sure to unzip the archive to extract the `.tar.gz` files.

3. Put these `.tar.gz` patch files into the previously created version, using either the Version Manager web UI or the command line. Now, You can use GEMM to deploy this version onto any Grid host, just as with an original, unpatched version of N1GE.

Be sure that only one patch level of N1GE is deployed across the grid. You should take care should to avoid mixing different patch levels in the same distribution. Also do not use patch files for only some but not all of the architecture-specific packages required for your environment.

Installing Grid Engine Hosts

To set up a compute grid, you must first select one of the managed hosts to be the *master host*. You can then set up additional *compute hosts*.

Note – After you install the N1GE6 software on a server and add hosts to the grid using GEMM, all the N1GE daemons on the hosts will be running, but you must submit jobs separately.

For documentation on the N1GE6 software, refer to the user manuals at the following URL:

<http://docs.sun.com/db/coll/1017.3> (<http://docs.sun.com/db/coll/1017.3>)

▼ Installing a Master Host

You can configure only one managed host to be the master host. If you have already configured a master and you select the Install Master sub-menu item, a message appears that you have already configured a master for the compute grid.

The Grid Engine module deploys only a *dedicated* N1GE6 master host. Unless you plan to have relatively low job throughput on your grid, you should not have the N1GE6 master host also act as a compute host. To add a host as a master host in the compute grid, you must first import the host into the SCS 2.2 framework. For more information, see “About Adding Managed Hosts” in the SCS 2.2 Release Notes.

Note – The SCS server cannot server as an N1 Grid Engine (N1GE) Master Host or Compute Host, since only SCS clients can have those roles. An SCS server cannot also be an SCS client at the same time. Thus, the SCS server has to be a different host than either the N1GE Master Host or Compute Hosts.

To install a master host for the grid:

- Steps**
- 1. Select Grid Engine > Install Master.**
The selector appears, displaying the list of managed hosts; see “Installing Grid Engine Hosts” on page 15.
 - 2. Click to highlight the managed host that you want to configure as the master host in the compute grid.**

3. **Pick a Version from the list presented**

The version picked at this step will be installed on the Master as well as all the hosts in the grid

4. **Click Install in the bottom right corner.**

The Task Progress dialog appears.



FIGURE 1-6 Install Master Dialog

▼ Installing Compute and Access Hosts

Once you have configured one of the managed hosts as the master host, you can add additional hosts to act as *compute hosts* or access hosts in the grid.

Note – To add a host as a compute host in the grid, you must first import the host into the Sun Control Station framework. For more information, see “About Adding Managed Hosts” in the SCS 2.2 Product Notes.

Note – Before you can add a compute host to a grid, you must first designate a master host. If you have not yet designated a master host, the system instructs you to do so. For more information see, “Installing Grid Engine Hosts” on page 15.

Steps 1. **Select Grid Engine > Install Compute Host.**

The selector appears, displaying the list of managed hosts; see the previous figure.

2. Click to highlight a host(s). You can also click **Select All** at the top to choose all hosts in the list.

You can pick the host(s) to be either compute or access hosts. Pick the desired button at the bottom of the page labelled “Install Compute Hosts” or “Install Access Hosts”

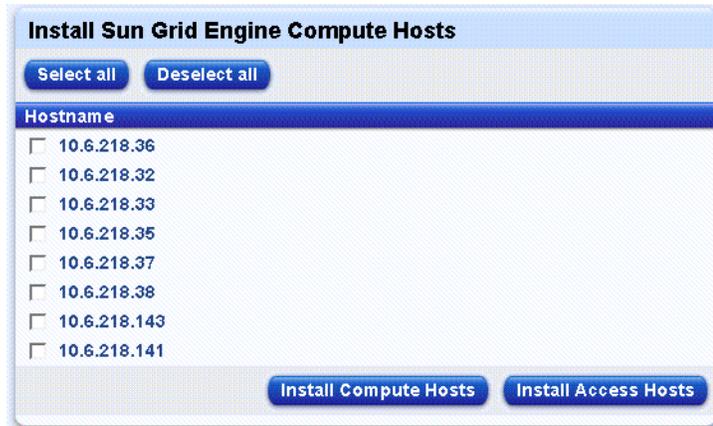


FIGURE 1-7 Install Compute Hosts Dialog

The Task Progress dialog appears. When the installation completes, a new dialog box appears which allows you to either finish the installation or view the installation events. If you choose **View Events**, a dialog similar to the following appears.

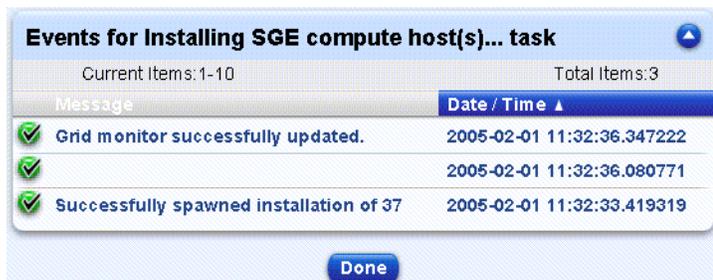


FIGURE 1-8 View Events Dialog

3. When you are finished installing hosts, click **Done**.

Monitoring the Grid

When you click the Monitor Grid menu item, a page with a high-level overview of the state of the grid appears.

This page has tables that allow you to:

- View Summary Status
- Examine Cluster Queue status
- Check Job Alerts
- Check Host Alerts
- Check Queue Alerts

Buttons on the main page let you go to pages where you can:

- View Job Details
- View Queue Details
- View Host Details
- Examine Daemon Log files

Also available from the SCS menu is the ability to quickly see the state of the Grid by choosing Station Settings >Active Monitor.

Viewing Summary Status

Summary Status	
2005-02-01 11:37:44.923173	
Running Jobs	29
Pending Jobs	96
Suspended Jobs	0
Held Jobs	0
Requeued Jobs	0
Error Jobs	0
Avg. Load	0.02
Total Used Mem. (Mb)	2356.98
Total Mem. (Mb)	23513.09
Total Number of Compute Hosts	8

FIGURE 1-9 Summary Status Table

The Summary Status table shows the total number of jobs in various states (pending, running, suspended, and so forth). It also shows the load averaged across all compute hosts and the total amount of used and installed memory summed over all compute hosts.

Updating Data

The subheading of this table contains a timestamp for when the data was obtained. By default, most monitoring data is automatically refreshed every minute. To display the most up-to-date database information in the tables, click the Monitor Grid menu item again. You can also reload the browser window. If the monitoring is not working properly for any reason, the subheading displays a warning and displays the timestamp for when the data was most recently obtained. This timestamp applies to all monitoring information displayed in GEMM, not just the Summary Status table.

Above this table is the Update button. Clicking this button retrieves the data immediately instead of waiting for the next one-minute interval. A progress bar shows the progress of the update. When the update completes, click the Done button to return to the main Monitor Grid page with the new data and updated timestamp.

If an update of the monitor is already in progress when you click the button, a message indicate this situation. As soon as the update in progress completes, the Update button will again be available to force a new update.

Viewing Jobs

You access the Jobs details page by clicking the Jobs button in the Summary Status table on the main Monitor page.

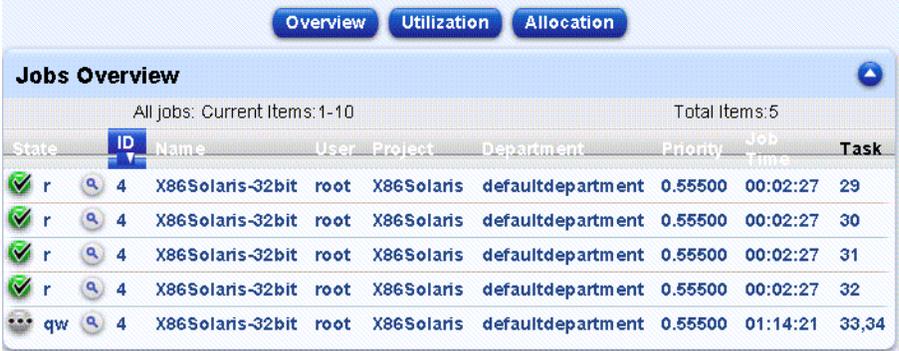
This page has a table which shows a summary of all current jobs in the system including jobs which are pending, running, suspended, held, or in an error state. Completed jobs are not listed. The top row of three buttons lets you see the list of jobs according to three different views: Overview, Utilization, and Allocation.

The initial view is always the Overview. Clicking any of the other buttons displays the other corresponding views. In all views, the back button on the table leads back to the main page.

Also present in all views at the bottom of the frame is the Filter, which you can use to limit the jobs displayed by providing configured criteria. Finally, the three buttons corresponding to the three different views are always shown at the top of each view, allowing you to move directly among the three views.

Using the Overview View

The Overview view shows an overall summary of the jobs.



The screenshot shows a web interface for 'Jobs Overview'. At the top, there are three buttons: 'Overview' (selected), 'Utilization', and 'Allocation'. Below the buttons is a table with the following data:

State	ID	Name	User	Project	Department	Priority	Job Time	Task
✓ r	4	X86Solaris-32bit	root	X86Solaris	defaultdepartment	0.55500	00:02:27	29
✓ r	4	X86Solaris-32bit	root	X86Solaris	defaultdepartment	0.55500	00:02:27	30
✓ r	4	X86Solaris-32bit	root	X86Solaris	defaultdepartment	0.55500	00:02:27	31
✓ r	4	X86Solaris-32bit	root	X86Solaris	defaultdepartment	0.55500	00:02:27	32
⋯ qw	4	X86Solaris-32bit	root	X86Solaris	defaultdepartment	0.55500	01:14:21	33,34

FIGURE 1-10 Jobs Overview Page

The columns in this table provide the:

- Job state, indicated by one or more letters plus a colored circle and icon
- Job ID
- Job name
- User who submitted the job
- Project under which the job was submitted
- Department of the submitter

- Priority of the job
- Job time, either the time spent pending, or for running jobs the time spent running
- Job task ID; for pending jobs, all task IDs are grouped together

Interpreting the Data

The icon scheme for the job state is:

- A Gray Icon means the job is pending
- A Green Icon means the job is running
- A Yellow icon indicates the job is suspended
- A Red Icon indicates that the job is in an error state

The letters shown for the job state are the same letters used by N1 Grid Engine to indicate the job state when you run the `qstat` command. For more information, see the *N1 Grid Engine Administration* manual.

Sorting Rows

Jobs display ten rows at a time. You can see the entire list by using the pagination controls at the bottom of the table. By default, rows are displayed numerically by job ID, but you can use any column whose header is white to change the ordering of the rows. Clicking on a column header sorts the rows according to the values in that column. Clicking again on the column header reverses the sort. The sorting is preserved across pages if you click on a pagination button.

Job Details

Clicking the Inspect icon next to the ID of each job retrieves details about the job. A progress bar indicates the progress of this process. When the Done button appears, clicking it leads to a page with the details displayed for the chosen job. These details appear in three tables.

Job Details for Job 6	
Attribute	Value
Job ID	6
Exec File	job_scripts/6
Submission Time	Tue Feb 1 18:14:02 2005
Owner	root
User ID	0
Group	root
Group ID	0
Home Directory	/root
SGE Log Name	root
SGE Path Name	/gridware/sge/bin/ix24-x86 : /usr/kerberos/sbin : /usr/kerberos/bin : /usr/local/sbin : /usr/local/bin : /sbin : /bin : /usr/sbin : /usr/bin : /usr/X11R6/bin : /root/bin
SGE Shell Name	/bin/bash
Work Directory	/suites/sge
Submit Host	dt218-130
Account	sge
Merge	y
hard resource_list	arch=sol-sparc
Mail List	root@dt218-130
Notify	FALSE
Job Name	Solaris32-bit
Standard Out Path	/dev/null
jobshare	0
Shell List	/bin/sh
Environment	

FIGURE 1-11 Job Details Page

- The first table shows the job details, including various properties related to the jobs environment, resource requests, submit options, and so forth.
- The second table shows the current resource utilization for that job. If this information is not available, for example, because the job started too recently or the job is still pending, then this table is empty. For jobs with multiple tasks, the usage of each task appears on a separate line.
- The third table shows the scheduling information for that job.

The information displayed in these three tables corresponds directly to the output from the N1 Grid Engine 6 `qstat -j` command. For more information on job details, see the N1 Grid Engine 6 Administration manual.

Clicking the Back button of the first table returns you to the Overview page.

Using the Utilization View

You access the Utilization view of the job by clicking the Utilization button on the Jobs page.



State	ID	Name	Queue	CPU	Memory	Share	Run Time	NTickets	NUrgency	NPosix	Task
✓	6	Solaris32-bit	all.q@dt218-32	0:00:03:12	0.75572	0.00000	-06:33:07	0.5	0.5	0.5	1
✓	6	Solaris32-bit	all.q@dt218-32	0:00:03:15	0.76786	0.00000	-06:33:07	0.5	0.5	0.5	3
✓	6	Solaris32-bit	all.q@dt218-32	0:00:03:13	0.76183	0.00000	-06:33:07	0.5	0.5	0.5	5
✓	6	Solaris32-bit	all.q@dt218-32	0:00:03:14	0.76669	0.00000	-06:33:07	0.5	0.5	0.5	7
✓	6	Solaris32-bit	all.q@dt218-38	0:00:03:00	0.70795	0.00000	-06:33:07	0.5	0.5	0.5	2
✓	6	Solaris32-bit	all.q@dt218-38	0:00:02:59	0.70387	0.00000	-06:33:07	0.5	0.5	0.5	4
✓	6	Solaris32-bit	all.q@dt218-38	0:00:02:59	0.70698	0.00000	-06:33:07	0.5	0.5	0.5	6
✓	7	X86Solaris-32bit	all.q@dt218-141	0:00:03:07	0.31210	0.00000	-06:33:07	0.5	0.5	0.5	1
✓	7	X86Solaris-32bit	all.q@dt218-141	0:00:03:06	0.31112	0.00000	-06:33:07	0.5	0.5	0.5	2
✓	7	X86Solaris-32bit	all.q@dt218-141	0:00:03:07	0.31139	0.00000	-06:33:07	0.5	0.5	0.5	3

FIGURE 1-12 Jobs Utilization View

Unlike the Overview view, only running and suspended jobs appear. In the Utilization view, the columns are the:

- Job state, indicated by a colored circle and icon
- Job ID
- Job Name
- Queue instance where the job is running
- CPU utilization of the job
- Memory utilization of the job
- Calculated share
- Run time
- Normalized Ticket priority
- Normalized Urgency priority
- Normalized POSIX priority
- Job task ID; tasks belonging to the same job are never grouped

Note – If the CPU usage or memory usage values are blank, the usage information for that job has not yet been reported. Check back at a later time to see if the usage is then reported.

The description for the Overview page regarding the meaning of the icons for the job state is the same for this view, except that no letters are shown. The pagination of the table and the sorting based upon different columns all apply similarly to the Utilization View.

Job Diagnostics

An Inspect icon for the job Task ID is displayed for all jobs above the final column. Clicking this icon retrieves the current diagnostic information for that job. This diagnostic information corresponds to the data found in the job spool files in the jobs spool directory. A progress bar indicates the progress of this process. When the Done button appears, clicking it leads to a page with the status information displayed for the chosen job as in the following figure.

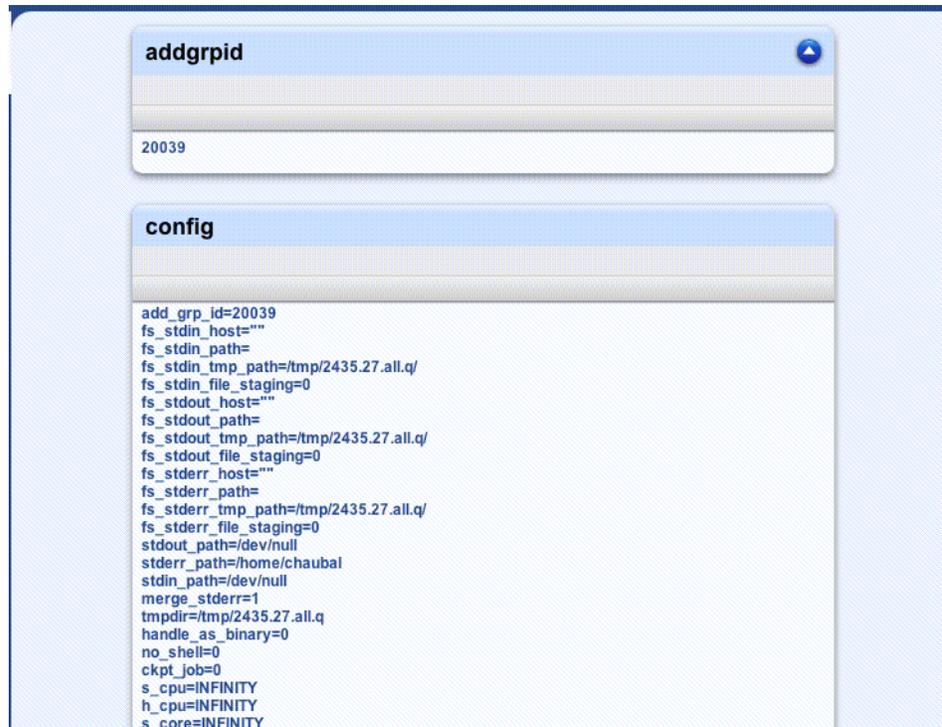


FIGURE 1–13 Job Diagnostic Details

Note – You can only obtain job diagnostic information if the job is running on a compute host that was deployed by GEMM. If the host on which the job is running was not deployed by GEMM, then clicking the Inspect icon results in an error message; clicking Done leads back to the Utilization view.

The Job Diagnostic details given in these tables include:

- addgrp_id
- config
- environment
- error
- exit status
- job_pid
- pe_hostfile
- pid

- trace

Interpreting the Tables

Each table corresponds to a different file from the job spool directory. For more information on the information in the job spool directory, see the N1 Grid Engine 6 Administration manual.

Clicking the back button of the `addgrpid` table returns you to the Utilization view.

Note – If a job has already completed by the time you click the Inspect button, or if the job completes during the information retrieval process, the information is lost and cannot be displayed. In this case, the progress bar will indicate a failure and clicking on the Done button leads back to the Utilization view.

Using the Allocation View

Clicking the Allocation button switches to the Allocation view of the jobs.

State	ID	Name	Tickets	Ovrde	Func	Tree	Posix	Urgency	Res	Wait	Ddln	Task
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	2
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	4
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	6
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	1
⚙	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	8-25:1
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	3
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	5
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	7
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	2
✓	6	Solaris32-bit	0	0	0	0	0.0	1000	1000	0	0	4

FIGURE 1-14 Jobs Allocation View

In this view, information is presented for all jobs and the columns provide details for the:

- Job state, indicated by a colored circle and icon
- Job ID
- Job name

- Total number of tickets for the job
- Number of override tickets
- Number of functional tickets
- Number of share tree tickets
- POSIX priority
- Total urgency for the job
- Resource contribution to the urgency
- Deadline contribution to the urgency
- Waiting time contribution to the urgency

The description of the icons for the job state on the Overview page apply here also, except that no letters are shown. The pagination of the table, and the sorting based upon different columns all apply similarly to the Allocation View.

For more information on the meaning of each column, see the N1 Grid Engine Administration manual.

Filtering Jobs

In each of the three views Overview, Utilization, and Allocation, the Filter option appears below the job table.

Column	Search Type	Search Value
State	less than	
State	less than	
State	less than	

FIGURE 1-15 Filter Dialog

You use the filter to limit the jobs displayed to those matching a specified search condition. The filter lets you choose a column on which to filter, a search type to use, and a value on which to search.

You select the column and search type from a drop-down table, while you type the value into a text entry box. The drop-down table for column changes with each view depending on which columns are being displayed. The type of search can be one of: equals, not equals, less than, less than or equal to, greater than, and greater than or equal to.

You can define up to three filters at one time; the effects of multiple filters are combined together to provide the final result. After you set up the desired filter, click the Filter button to redisplay the current view with the filter applied. Pagination is still active and will maintain the filter across pages. Clicking the Clear button restores the unfiltered view.

The following figure shows you how a sorted jobs utilization page would look.

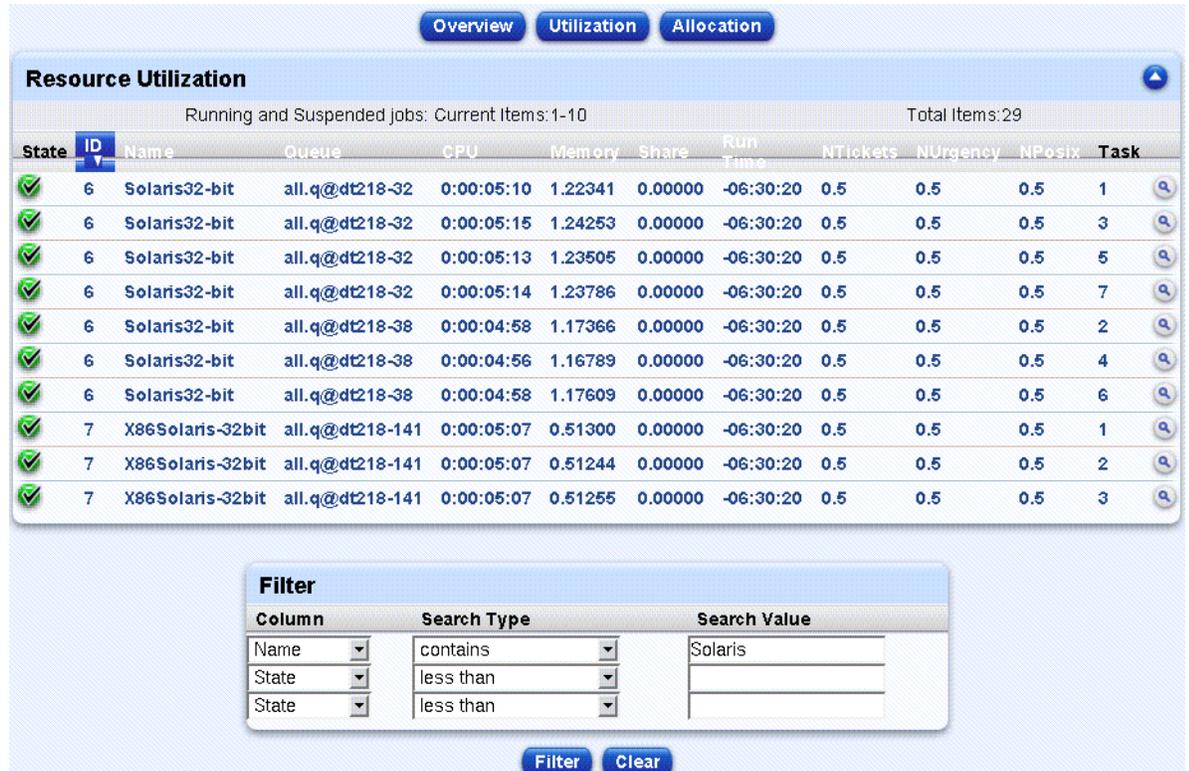


FIGURE 1-16 Filter Sorted Page

Note – When you choose the Job State as a search column, the search value is compared against the job status letter code as displayed in the Overview view, even though these letters are not displayed for the Utilization and Allocation view.

Viewing Queue Details

You access the Queue Details page by clicking the Queue button in the Summary Status table on the main Monitor page.

Queue Details				
Current Items: 1-10			Total Items: 10	
State	Queue	Status	Used Slots	Total Slots
✓	all.q@dt218-130	Running	0	2
✓	all.q@dt218-141	Running	4	4
✓	all.q@dt218-143	Running	0	4
✓	all.q@dt218-144	Running	0	4
✗	all.q@dt218-151	Alarm	0	1
✓	all.q@dt218-32	Running	0	4
✓	all.q@dt218-33	Running	0	4
✓	all.q@dt218-35	Running	0	4
✓	all.q@dt218-37	Running	0	4
✓	all.q@dt218-38	Running	0	3

FIGURE 1-17 Queue Details Page

Note that this table provides information on all queue instances on the currently selected master host, including instances on hosts that were not added by GEMM framework. The information appears in groups of ten rows at a time, with the ability to page back and forth between the rows.

Interpreting Data

For each queue instance, there are columns for the Queue instance name, the status, the total number of slots and number of used slots. The status is indicated by a colored circle and icon similar to the Job Alerts previously described. The only additional feature is a green icon to indicate queue instances that have no alert conditions. Clicking the Back icon in the table header returns you to the Monitor Grid main page.

Sorting Data

By default, rows display alphabetically by queue instance name but you can use any column whose header is written in white to change the ordering of the rows. Clicking on a column header sorts the rows according to the values in that column; clicking again on the column header reverses the sort. The sorting is preserved across pages if you click a pagination button.

Viewing Additional Details

The final column of each row has an Inspect icon. Clicking on this icon displays a table with the full details for that queue instance. The final entry in this table shows the timestamp when the data was obtained. For information on the meaning of the other table entries, consult the N1 Grid Engine 6 Administration manual. Clicking on the 0 icon for this table returns you to the Queue Details page.

Viewing Host Details

You access the Host Details page by clicking the Host button on in the Summary Status table on the main Monitor page.

Host Details -- SGE Proxy host: dt218-36					
Current Items:1-10			Total Items:10		
Hostname	Arch	Load/CPU	Used Mem	Total Mem	Free Swap
dt218-141	sol-x86	1.16	181.0	2046.976	2046.976
dt218-143	lx24-x86	1.23	103.254	2010.112	2046.976
dt218-144	lx24-x86	1.36	651.172	2020.352	1787.904
dt218-151	lx24-amd64	2.24	2537.472	3922.944	1992.704
dt218-32	sol-sparc	1.27	416.0	4096.0	2046.976
dt218-33	sol-sparc64	1.43	486.0	4096.0	2049.024
dt218-34	sol-sparc	1.44	369.0	4096.0	2046.976
dt218-35	sol-sparc64	1.24	486.0	4096.0	2049.024
dt218-37	sol-sparc64	1.39	445.0	4096.0	2049.024
dt218-38	sol-sparc	1.59	202.0	1024.0	2046.976

FIGURE 1-18 Host Details View

This page displays a table with the state of all the compute hosts that are members of the grid. The title of the table also indicates which host is currently chosen as the Proxy Host.

Note that this table has information on all compute hosts reporting to the currently-chosen master host, including those that were not added by GEMM framework.

Interpreting Data

The information appears in groups of ten rows at a time, with the ability to page back and forth between the rows. For each host, there are columns for the Hostname, Architecture, Load per CPU, Memory in use, Total Memory, and Swap Space in use. The status is also indicated by a colored circle and icon similar to the Host Alerts table with an additional green icon to indicate hosts that have no alert conditions. Clicking the Back icon in the table header returns you to the Monitor Grid main page.

Sorting Data

By default, rows display alphabetically but you can use any column whose header is white to change the ordering of the rows. Clicking on a column header sorts the rows according to the values in that column; clicking again on the column header reverses the sort. The sorting is preserved across pages if you click a pagination button.

Seeing Additional Details

The final column of each row has an Inspect icon. Clicking on this icon displays a table where full details for that host appear. The final entry in this table shows the timestamp when the data was obtained. For information on the meaning of the other table entries, please consult the N1 Grid Engine 6 Administration manual. Clicking the Back icon on this table returns you to the Host Details page.

Viewing Grid Engine Daemon Logs

You access the Grid Engine Daemon Logs page by clicking the Daemons button on in the Summary Status table on the main Monitor page.

Grid Engine Daemon Logs		
Current Items: 1-10		Total Items: 10
Host	Qmaster	Execd
dt218-130		
dt218-141		
dt218-143		
dt218-144		
dt218-32		
dt218-33		
dt218-35		
dt218-36		
dt218-37		
dt218-38		

1-10

FIGURE 1-19 Grid Engine Daemons Log View

The Logs page contains a table which displays the names of all compute hosts that were deployed by GEMM, plus the name of master host if it was deployed by GEMM.

Two additional columns are also shown. The first column, labeled Master, contains an Inspect icon for the master host. The second column, labeled execd, contains an Inspect icon for each compute host. Clicking these icons lets you retrieve the actual log message files.

Note – If the master host was not deployed by GEMM, no host in the table will have the Inspect icon for the Qmaster column. Similarly, if there are compute hosts that were not deployed by GEMM, these hosts will not appear in this table. Clicking the Back icon in the table header returns you to the Monitor Grid main page.

Retrieving Log Message Files

Grid Engine Logs for Execd dt218-130		
Current Items: 1-25		Total Items: 391
Time Stamp	Message	
01/20/2005 17:44:48	starting up 6.0u3	
01/20/2005 17:44:15	NULL rhost passed to sge_send_receive_gdi_request()	
01/20/2005 17:44:15	unable to read qmaster name: can't open "/gridware/sge/default/common/act_qmaster" for reading qmaster hostname	
01/20/2005 17:44:15	fopen("/gridware/sge/default/common/act_qmaster") failed: No such file or directory	
01/20/2005 17:43:55	NULL rhost passed to sge_send_receive_gdi_request()	
01/20/2005 17:43:55	unable to read qmaster name: can't open "/gridware/sge/default/common/act_qmaster" for reading qmaster hostname	
01/20/2005 17:43:55	fopen("/gridware/sge/default/common/act_qmaster") failed: No such file or directory	
01/20/2005 17:43:35	NULL rhost passed to sge_send_receive_gdi_request()	
01/20/2005 17:43:35	unable to read qmaster name: can't open "/gridware/sge/default/common/act_qmaster" for reading qmaster hostname	
01/20/2005 17:43:35	fopen("/gridware/sge/default/common/act_qmaster") failed: No such file or directory	
01/20/2005 17:43:15	NULL rhost passed to sge_send_receive_gdi_request()	
01/20/2005 17:43:15	unable to read qmaster name: can't open "/gridware/sge/default/common/act_qmaster" for reading qmaster hostname	
01/20/2005 17:43:15	fopen("/gridware/sge/default/common/act_qmaster") failed: No such file or directory	
01/20/2005 17:42:55	NULL rhost passed to sge_send_receive_gdi_request()	
01/20/2005 17:42:55	unable to read qmaster name: can't open "/gridware/sge/default/common/act_qmaster" for reading qmaster hostname	

FIGURE 1–20 Example Log Message File

Clicking an inspect icon retrieves and displays the `qmaster` and `execd` daemon messages file for the corresponding host. A progress bar indicates the progress of this process. When the Done button appears, clicking it displays the contents of the chosen messages file with each line appearing in its own row in a table. Rows display 25 at a time with the ability to page through them.

The rows display in reverse chronological order, so that the most recent message appears at the top of the list. Clicking on the Back icon for this table returns you to the Grid Engine Daemon Logs page. For more information on daemon messages, see the N1 Grid Engine 6 Administration manual.

Interpreting Messages

The first column of this table shows a colored circle and icon to indicate the severity of that message. A green circle indicates a message of type Info. A yellow circle indicates a message of type Warning or Critical. A red circle indicates a message of type Error. The second column shows the time stamp for the message and the third column shows the actual text of the message.

Viewing Cluster Queues



Cluster Queues						
Name	Total Slots	Used	Alarm	Disabled	Suspended	Error/Unknown
all.q	25	0	0	0	0	0

FIGURE 1-21 Cluster Queues Page

This table shows a summary of the state of all the cluster queues configured on the grid, indicating the numbers of slots in various states. For information on cluster queues, see the N1GE 6 Administration Guide.

Viewing Host Alerts

Host Alerts						
Proxy host: 10.6.218.130						
* Hostname	Load Per CPU	Used Mem.	Total Mem.	Swap Used	Date/Time	
⊗ dt218-141	0.84	133.0	2046.976	2046.976	2005-02-01 11:46:47.509011	
⊗ dt218-143	0.83	130.961	2010.112	2046.976	2005-02-01 11:46:47.796513	
⊗ dt218-32	0.83	290.0	4096.0	2046.976	2005-02-01 11:46:47.97635	
⊗ dt218-33	0.83	441.0	4096.0	2049.024	2005-02-01 11:46:48.127144	
⊗ dt218-35	0.83	440.0	4096.0	2049.024	2005-02-01 11:46:48.278172	
⊗ dt218-36	0.83	343.0	2048.0	2049.024	2005-02-01 11:46:48.426551	
⊗ dt218-37	0.83	466.0	4096.0	2049.024	2005-02-01 11:46:48.537191	
⊗ dt218-38	0.85	184.0	1024.0	2046.976	2005-02-01 11:46:48.723965	

FIGURE 1-22 Host Alerts Page

This table shows all hosts where the threshold for either the load or memory has been crossed. There are two types of alerts each indicated by a different colored circle and icon.

A warning alert is indicated by a yellow icon. This alert displays if the load goes above the load warning threshold or the memory goes below the memory warning threshold.

A critical alert is indicated by a red icon. This alert displays if the load goes above the load critical threshold or the memory goes below the memory critical threshold.

The Host Alerts table is empty if no hosts have crossed any threshold. You configure the values for the load and memory warning and critical thresholds on the Settings page.

Viewing Queue Alerts

Queue Alerts					
* Queue Name	Status	Used Slots	Total Slots	Date/Time	
⊗ all.q@dt218-151	au	0	1	2005-02-01 11:52:21.413	

FIGURE 1-23 Queue Alerts Page

This table shows queue instances that are not in the usual running state. There are three types of alerts each indicated by a different colored circle and icon.

- A red icon indicates the queue instance is in either the Unknown or Error state.
- A yellow icon indicates the queue instance is in either an Alarm or Suspended state.
- A gray icon indicates the queue instance is in a Disabled state.

The exact state of the queue instance is also given in the Status column. For more information on queue instance states, see the N1 Grid Engine 6 Administration Manual.

Viewing Job Alerts



Job Alerts				
* Job ID	Task	Job Name	Job Time	Deadline
There is currently no data				

FIGURE 1-24 Job Alerts Page

This table displays grid jobs which are not in the usual running state. There are two types of alerts each indicated by a different colored circle and icon.

- A red icon indicates the job is in an Error state.
- A yellow icon indicates the jobs pending time has exceed the pending time threshold.

You configure the values for the pending time threshold on the Settings page. For more information on job states, see the N1 Grid Engine 6 Administration manual.

Using Grid Active Monitor

You can quickly see the status of the Grid by using the SCS Active Monitor feature. Choose Station Settings >Active Monitor. and scroll down the page to the Base Services table shown in the following figure.

Base Services			
Status	Service	Vendor	Date/Time Status Last Changed
...	Domain Name Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.408462
✓	OpenSSH Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.538791
✓	Postgres Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.593137
...	FTP Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.744672
✗	Email Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.805
✓	Web Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.86605
...	Telnet Server	Sun Microsystems, Inc.	2005-03-24 11:30:05.920447
✓	Tomcat Server	Sun Microsystems, Inc.	2005-03-24 11:30:06.047404
✓	Grid Engine Monitor	Sun Microsystems, Inc.	2005-04-22 13:30:02.71376

FIGURE 1-25 Grid Active Monitor Table

When the status of the grid changes due to an event like a queue alert, the button next to the Grid Engine entry changes color in the following way:

- Green: N1GE is up and running fine.
- Yellow: the SCS cannot contact the proxy host or cannot obtain monitoring information from it but it is still possible that the master is running.
- Red: the proxy host indicates that the master is down.
- Grey: N1GE is not installed anywhere.

Viewing Settings

When you click the Settings menu item a table displays with all the configurable settings available in GEMM.

Settings

Settings in italics cannot be modified while Master is installed

<i>Load Warning</i>	1.00
<i>Load Critical</i>	3.00
<i>Free Memory Warning (MB)</i>	100
<i>Free Memory Critical (MB)</i>	10
<i>Maximum Job Pending Time (hr)</i>	120
<i>SGE Root</i>	/beta/dir/2.23/
<i>SGE Cell</i>	default
<i>Qmaster TCP Port</i>	836
<i>Execd TCP Port</i>	837
<i>Admin Username</i>	sgadmin
<i>Admin UID</i>	218
<i>Admin Homedir</i>	/beta/dir/2.23/
<i>NFS Server Name</i>	10.6.218.151
<i>NFS Mount Point</i>	/beta/dir/2.23//default/cc
<i>Linux NFS Mount Options</i>	defaults
<i>Solaris NFS Mount Options</i>	
<i>Proxy Host</i>	dt218-151
<i>Grid Engine Version</i>	n1ge6_case_2.23

Update
Choose Proxy

FIGURE 1-26 Settings Page

The parameters are grouped in four categories: Monitor Alert settings, N1GE settings, NFS mount settings and Proxy settings.

Changing Monitor Alert Settings

These settings affect the display of alerts in the GEMM Monitor. All these parameters must be set using decimal numbers. Any other type of input produces a formatting error.

Load Warning -- You use this parameter to specify the load warning threshold. If this threshold is exceeded, a load warning alert appears in the Monitor. The value is in terms of system load, as reported by the OS, divided by the number of CPUs.

Note -- Certain microprocessors with special features such as hyperthreading may be registered as having more than one CPU per physical CPU socket, depending upon factors such as the BIOS or PROM configuration.

Load Critical -- You use this parameter to specify the load critical threshold. If this threshold is exceeded, a load critical alert appears in the Monitor. Similar to the Load Warning parameter, you set this parameter in terms of the system load scaled by number of CPUs.

Memory Warning -- You use this parameter to set the memory warning threshold. If the value drops below this threshold, a memory warning alert appears in the Monitor. You set the parameter value in terms of megabytes of free virtual memory.

Memory Critical -- You use this parameter to set the memory critical threshold. If the value drops below this threshold, a memory critical alert appears in the Monitor. You set the value in terms of megabytes of free virtual memory.

Maximum Job Pending Time -- You use this parameter to specify the amount of time that a job spends pending after which a Job Pending alert appears in the Monitor. You set the value in hours.

Note – It is important that you set these five parameters to sensible values, according to the characteristics of your particular grid. Otherwise, an excessive number of alerts will appear on the Monitor main page, cluttering the display.

Changing N1GE Settings

The N1GE settings affect the way N1GE is installed onto the master, compute and access hosts. The N1GE administrator must determine the various parameter values suited to their local Grid environment.

Factors you should determine include the local namespace for users, TCP services, file directory structure, operating system, and so forth. The values have default options which are suitable for a generic installation. You should be familiar with the N1GE 6 product before changing any of these values. If you wish to change more advanced configuration settings, please see Chapter 3, *Using the Setup configuration file*.

Once you deploy the master host, you cannot edit these values which remain in effect for all further deployments of compute and access hosts. You can only edit the values again if you uninstall the master host. The following section describes each setting

SGE Root -- This setting is the root directory under which the N1GE files will be installed. Note that the files will be installed on all hosts in this directory.

SGE Cell -- This settings is the N1GE cell name used for the deployment.

Qmaster TCP Port -- This setting is the TCP port to use for the N1GE qmaster daemon.

Execd TCP Port -- This setting is the TCP port to use for the N1GE execd daemon.

Admin Username -- This setting is the username of the N1GE admin user.

Admin UID -- This setting is the UID of the N1GE admin user.

Grid Engine Version -- This parameter indicates the version of N1 Grid Engine that will be deployed on the compute and access hosts.

Changing NFS Settings

These settings affect the way the N1GE “common” directory for the chosen cell name is mounted on all access and compute hosts. The settings are described as follows.

NFS Server Name -- The name of the NFS server from which all compute and access hosts will mount the N1GE “common” directory. When you deploy the master host using GEMM, this parameter is set automatically to the master host. Once you deploy the master host you cannot edit this value and it remains in effect for all further deployments of compute and access hosts. You can only edit the setting again if you uninstall the master host.

NFS Mount Point -- The directory which is mounted from the NFS server for the N1GE “common” directory. When deploying the master host using GEMM, this is set automatically to <SGE_Root>/<SGE_Cell>/common, where <SGE_Root> and <SGE_Cell> are the values specified above. Once you deploy the master host you cannot edit this value and it remains in effect for all further deployments of compute and access hosts. You can only edit the setting again if you uninstall the master host.

Linux NFS Mount Options -- This setting is the options used when mounting the “common” directory onto a Linux compute or access host. The value in this field is inserted into the Linux /etc/fstab file on each host as:

```
<Servername>:<Mountpoint> <Mountpoint> nfs <Mountoptions> 0 0
```

where <Servername> and <Mountpoint> are the values specified above and <Mountoptions> are the specified Linux NFS mount options.

Note – This parameter cannot contain any spaces

Solaris NFS Mount Options -- This setting specifies the options used when mounting the “common” directory onto a Solaris compute or access host. The value in this field is inserted into the Solaris /etc/vfstab file on each host as:

```
<Servername>:<Mountpoint> - <Mountpoint> nfs -yes <Mountoptions>
```

where <Mountpoint> is the values specified above and <Mountoptions> is the specified Solaris NFS mount options.

Note – This parameter cannot contain any spaces.

Changing the Proxy Host

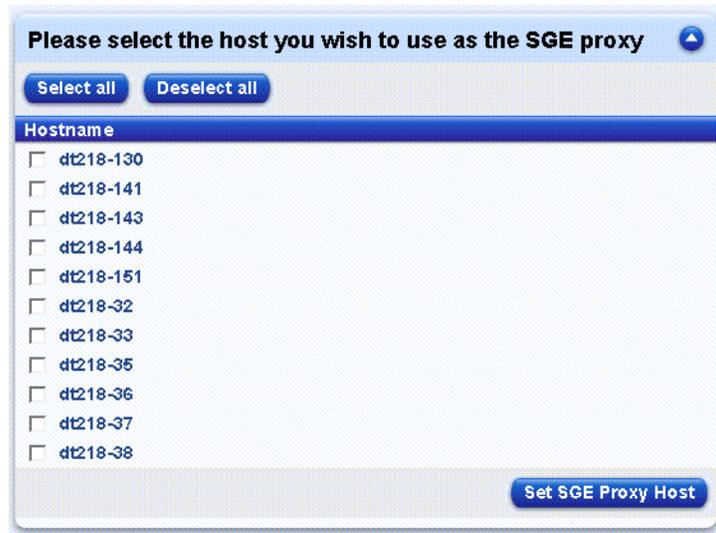


FIGURE 1-27 Change Proxy Host Page

Currently, there is only one proxy setting, which indicates the host on which monitoring commands are executed. If the master host has been previously deployed using GEMM, then the proxy host is set to this host and cannot be changed until the master is uninstalled. To choose the proxy host, click the Choose Proxy button at the bottom of the page. A table of all the hosts on which the GEMM framework has been installed. Select one host from this table.

Note – The host you chose must be an N1GE admin host; otherwise, install and uninstall of other hosts, as well as monitoring, could fail

Changing the N1GE Version



FIGURE 1-28 Change N1GE Version Page

To set N1GE version parameter, click the Choose Version button at the bottom of the page. This action presents a table from which you select a version by clicking its Inspect icon. The available versions are those uploaded in the GEMM Version management page. If you deployed the master host previously using GEMM, the version chosen at that time is displayed. Manual changes to this parameter are not allowed until you uninstall the master host.

Special Consideration: External Master Host

You can use GEMM for deployment and monitoring even with an N1GE master host not configured by GEMM. Possible scenarios include:

- There is an already-existing N1GE installation.
- You wish to deploy the master host on a platform not supported by the Sun Control Station framework.
- You need to install the master host in a configuration unsupported by GEMM, such as with a shadow host, or with high-availability cluster via Sun Cluster software.
- If you have an externally-configured master host, you can still use GEMM to deploy compute and access hosts, as well as for monitoring. However, you need to follow these steps:
 1. Collect the N1GE and NFS settings
 2. Establish a Proxy Host
 3. Deploy the Chosen Proxy Host as a Compute or Access

Collecting the N1GE and NFS settings

Once you have configured the master host and ensured that it is up and running properly, take note of all the values for the N1GE settings as well as the NFS settings. These settings are essentially the parameters you would use if you were to install an execution host manually and associate with the master host, including the choice of NFS options for mounting the N1GE common directory. For example, you might mount the common directory from the master host or you may need to mount it from a separate file server system or appliance. Note that the correct choice of the NFS settings for the N1GE common directory is a critical step, since the common directory contains a file which tells the compute and access hosts where to find the master host.

Part of this step is to ensure that the exact same version of N1 Grid Engine which is running on the master host has been uploaded to GEMM using the Version page. N1 Grid Engine will not function properly unless the same version, including update level is used on the master host and all compute and access hosts.

Once you have determined and set the N1GE and NFS settings, it is important not to modify them again. Otherwise, further compute and access host deployments could be corrupted and will not work.

Establishing a Proxy Host

In order for GEMM to deploy additional compute and access hosts and perform monitoring, you must choose a host from the Sun Control Station as an N1GE Admin Host. This host must remain as an N1GE admin host as long as GEMM is in use. You may choose a system which will be a compute host as well or you may choose a system which will only be an access host. This choice is determined by factors such as:

- Security concerns about a compute host having admin privileges --- this factor depends upon your established policy for using N1GE.
- Concerns about monitoring being impacted by compute tasks --- by default the monitoring command runs once a minute on the chosen host, which probably will not have a large impact unless the host is running a very resource-intensive job.
- Permanence - The host you choose must be one which you do not expect to take down ever during the course of running GEMM, otherwise monitoring and deployments will not work

Once you have decided which host to make the admin host, then perform these steps:

1. Set this host as the proxy host as previously described.
2. On the master host, add this host to the list of admin hosts. You can add the host using the N1GE GUI or add it from the command line by using the N1GE `qconf -ah` command.

Deploying the Chosen Proxy Host as a Compute or Access

At this point, click the Install Host menu item, select the chosen proxy, and install it as a compute or access host. You must select only the chosen proxy host and no other host in this step. You must wait to deploy additional hosts until the proxy host has been successfully established.

Uninstalling Hosts

From the Grid Engine main page you have two uninstall choices. You can uninstall a particular host or hosts or you can uninstall everything.

▼ To Uninstall Hosts

You can remove one or more compute hosts from the compute grid.

When you uninstall a compute host, the N1GE software is shut down and removed from the selected hosts. The N1GE master host is instructed to remove those compute hosts from the N1GE compute grid.

Note – Before you start the uninstall procedure, ensure that no jobs are running on the compute hosts that you want to uninstall. Any jobs that are currently running on these hosts will be terminated. If the jobs are marked as “re-runnable”, they are automatically resubmitted to the N1GE compute grid for execution on another compute host(s). However, if they are marked as “not re-runnable,” then they are *not* rescheduled and are not automatically executed elsewhere.

Steps 1. **Select Grid Engine > Uninstall Host.**

The selector appears, displaying the list of hosts currently in the compute grid; see “Uninstalling Hosts” on page 44.

2. **Click to highlight a host(s). You can also click Select All at the top to choose all hosts in the list.**

3. **Click Uninstall Selected Nodes in the bottom right corner.**

The Uninstall Task Progress Dialog appears.

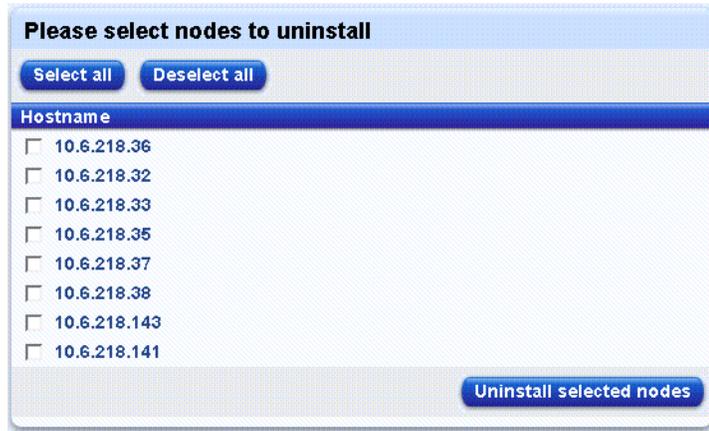


FIGURE 1–29 Select Nodes to Uninstall Page

▼ To Uninstall Everything

You can remove all components of the Grid Engine module from the master host and all compute hosts.

Before you uninstall everything, be aware that:

- all jobs (both running and suspended) are killed
- all pending jobs are lost
- all configurations and all records of previously run jobs are lost.

To uninstall everything:

- Steps**
1. **Select Grid Engine > Uninstall Everything.**
A screen appears, explaining the Uninstall Everything feature.



FIGURE 1–30 Uninstall Completely Dialog

2. **Click Uninstall Master Host and ALL Compute Hosts.**

The Task Progress dialog appears.

Command Line Equivalents

This chapter contains some command line equivalents for the actions you can perform using the browser GUI.

Command Line Interface (CLI)

Many of the tasks which can be performed from the UI can also be performed from the command line.

Manage Versions

To perform version management tasks use the following command with the appropriate argument.

Command: `/scs/sbin/gemmVersMgmt.pl [args]`

Arguments

- To list a version use: `-lv`
- To add a version use: `-av vers_name`
- To remove a version use: `-rv vers_name`
- To rename a version use: `-nv vers_name new_name`
- To list several version files use: `-lf vers_name`
- To add several version files: `-af vers_name`
- To remove several version files: `-il`

Examples

This example lists the versions in directory `/scs/data/gemm/versions`:

```
- /scs/sbin/gemmVersMgmt.pl -lv /scs/data/gemm/versions
```

This example renames the directory `vers_name` to a new name.

```
- /scs/sbin/gemmVersMgmt.pl -nv vers_name new_name
```

This example adds files to the directory `vers_name`:

```
- /scs/sbin/gemmVersMgmt.pl -af vers_name file1 file2 file3
```

This example removes files from the directory `vers_name`:

```
- /scs/sbin/gemmVersMgmt.pl -rf vers_name file1 file2 file3
```

Install a Master Host

To install a N1GE master host use this command

```
/scs/sbin/gemmInstallMaster.pl id version
```

Note – Now the version you install will be used for all compute and access hosts.

The *id* can be a hostname, an IP, or an appliance id. You can see a list of appliances using this command:

```
/scs/sbin/deviceInfo.pl -l
```

Example

```
- /scs/sbin/gemmInstallMaster.pl dt218-37 n1ge6.0
```

Install a Host

To install an N1GE compute host use the following command:

```
/scs/sbin/gemmInstallHost.pl host_type id1 id2
```

In this command, *host_type* is compute or access and *id* is an IP, a host, or an appliance.

You can get the appliances managed by GEMM using:

```
/scs/sbin/deviceInfo.pl -l
```

Example

```
- /scs/sbin/gemmInstallHost.pl compute dt218-34 dt218-36 dt218-38
```

Configure Settings

This command lets you list or change the settings.

```
/scs/sbin/gemmSettings.pl parameter = value
```

Examples

If you use the command without a *parameter = value* pair, you see a list of the present set of values as in the following example:

```
- /scs/sbin/gemmSettings.pl
```

```
Current configuration:
admin_homedir = /gridware/
sgadmin_uid = 218
admin_username = sgadmin
execd_port = 537
inst_version = nlge6.0
lnx_nfs_mt_opts = intr,softload
critical = 3.00
loadWarning = 1.00
maxPendTime = 120
memCritical = 10
memWarning = 100
nfs_mount_point = /gridware/sge/default/common
nfs_server_name = 10.6.218.37
proxy_id = 5
proxy_is_admin = Yq
master_id = 5q
master_port = 536
qmaster_ready = Y
sge_cell = default
sge_root = /gridware/sge
sol_nfs_mt_opts =
```

If you use the command with a *parameter = value* pair, it updates that parameter to the new value as in the following example.

```
- /scs/sbin/gemmSettings.pl maxPendTime=34
```

Uninstall Hosts

You use this following command to uninstall compute, access, and master hosts. The master will not be uninstalled until all other hosts have been uninstalled

```
Run - /scs/sbin/gemmUninstallHost.pl id1 id2
```

In this command, *id* is an IP, hosts, or appliance id.

Example

```
- /scs/sbin/gemmUninstallHost.pl dt218-36 dt218-38
```

Using the Setup configuration File

This chapter describes how you can do advanced configuration of your grid environment using the `setup.conf` file.

Overview

The Settings page allows you to configure the minimum set of parameters required for setting up a grid in your environment. If you wish to do advanced configuration of the grid for your environment, then it is possible to change the default values of an additional set of parameters. You should only do this if you are an advanced user of N1GE6; otherwise, the default values for these parameters should be proper for most grid installations.

Changing the Configuration File

The automatic installation is controlled by parameters located in the file `/scs/data/gemm/conf/setup.conf` located on the Sun Control Station server. To modify the values of these parameters, log in to the SCS server, and manually edit this file.



Caution – You can only modify the variables above the line which reads `# DO NOT MODIFY VARS BELOW THIS LINE!!!!`

Configurable Parameters

The modifiable parameters, their meaning and the default values are listed in the following sections.

EXECD_SPOOL_DIR_LOCAL="/var/spool/sge"

This variable defines the local spool directory used to write messages and information related to jobs processed on each compute host. You should set this parameter to a directory which is not used or modified by any other process. If this directory path does not exist on the host, it will be created at install time. When the host is uninstalled, this directory path is removed.

SPOOLING_METHOD="berkeley"

This variable defines the default spooling method. The options are "classic" and "berkeley". You should use the default "berkeley" value in most situations, because it provides the best performance. You should only use the value "classic" value for special circumstances. Please consult the N1GE Administration Guide for more details.

HOSTNAME_RESOLVING="true"

If this variable is true, the domain name is ignored during hostname resolution. If there are hosts in different domains with the same short hostname (for example, `hostA.domain1` and `hostA.domain2`), you should set this value should to "false".

DEFAULT_DOMAIN="none"

This parameter defines the name of the default domain, if you are using `/etc/hosts` or NIS for name resolution.

GID_RANGE="16000-16100"

In order to keep track of running jobs, Grid Engine assigns additional UNIX group IDs using this is range of values. The range should be as large as the maximum number of jobs that will ever be allow to run on a single host simultaneously. This range of group ids should not be used for any other purpose on any system. You should use numbers above 10000.

ADMIN_MAIL="none"

This value is the email address to which diagnostic messages are sent in case of problems (such as job failure or queue error).

`ADD_TO_RC="true"`

If you set this value to true, the RC scripts for N1GE daemons will be installed so that they start automatically at boot time.

`SET_FILE_PERMS="true"`

If you set this variable to "true", the file permissions of N1GE executables will automatically be set to the values required for proper operation.

`SCHEDD_CONF="1"`

This variable defines one of the three distributed scheduler tuning configurations (1=normal, 2=high, 3=max). This choice determines the initial values of certain scheduler parameters on the master host. You can change any value at any time after the initial installation. For more information on the three configurations, please consult the N1GE Administration Manual.

