



Sun N1™ Grid Engine 6

Workload Management and Grid Provisioning
for Compute-Intensive Requirements

Highlights

- Connect Microsoft Windows desktops as execute and submit hosts
- Supports virtually all commercial operation systems
- Interoperability with Sun™ Control Station 2.2 and 2.2.1 for single GUI grid management
- New optimized binaries for Solaris™ 10 OS x64
- Improved Accounting and Reporting Console (ARCo)
- Enterprise-grade scalability
- Flexible, extensible resource management capabilities
- Detailed statistics gathering for analysis and utility billing
- Sophisticated scheduler with advanced planning abilities
- Industry-standard APIs
- Supported integration with industry-standard Globus Toolkit through Sun iForce™ partners



Overview

The Sun N1™ Grid Engine 6 enables enterprises to build grids that make employees more productive. Enterprises can monitor and select the optimal usage of computer resources on most commercial operating systems and platforms. The latest refresh of the N1 Grid Engine 6 is update 4, released in May 2005.

Sun N1™ Grid Engine 6

The N1 Grid Engine 6 provides policy-based workload management and dynamic provisioning of application workloads.

It creates a grid of network-connected servers, workstations, and desktops; provides user access to the grid; and provides administrative and management interfaces. Computing tasks or jobs are distributed across the grid in accordance with resource requirements for the job, user requests for the job, and administrative/managerial policies. Usage accounting data is stored and made available so that it is possible to determine what resources were used in the execution of a job, and for whom the job was run. Administrative overhead for the grid is low, so depending upon the optimization of resources to workload, overall resource utilization in an N1 Grid Engine 6 grid can approach 100 percent.

N1 Grid Engine 6 is used in enterprises to create grids at both a department level and at a multidepartment or enterprise level. Through the pooling of departmental resources into larger enterprise grids, multiple users, teams, and departments can share common resources while working on different projects with different goals and schedules, providing maximum

resource availability to all users in a flexible, policy-based environment. Productivity can be dramatically increased compared to pregrid approaches. For example, Mentor Graphics, an electronic design automation software supplier, can run an application that previously took 13 hours on a workstation in 45 minutes, using a grid of 30 unused desktops.

In the May 2005 Update 4, N1 Grid Engine 6 added new binaries for the Solaris™ 10 OS x64 and the ability to connect Microsoft Windows desktops as execute hosts. The interoperability with Sun™ Control Station 2.2 and 2.2.1 allows enterprises to manage a large number of nodes from a single GUI.

Enterprise-grade scalability

The N1 Grid Engine 6 utilizes a new, multi-threading master daemon agent architecture, a redesigned internal agents communication, and an improved spooling process. These enhancements significantly reduce the latency in the system, while increasing capacity and adding new features. An N1 Grid Engine 6 master can now manage a grid of up to ten thousand hosts, meeting the scalability needs of even the largest grids.

The ability to incorporate thousands of Microsoft Windows desktops and servers in grids enables enterprises to significantly increase the utilization of compute resources. Hundreds and thousands of ubiquitous, previously idle, Microsoft Windows desktops are now put to work. The productivity increase can bring projects to completion faster than before, simply by employing idle resources.

The N1 Grid Engine 6 can be installed with three different profiles: Normal, High, and Maximum throughput. These scheduler profiles are useful for optimizing the system for the typical workload, and can be further fine-tuned at any time. For instance, many electronic design automation, biotechnology, and financial services have long sequences of intensive tasks with individually short execution times. N1 Grid Engine 6 can be optimized for this and other workload profiles.

Flexible, extensible resource management capabilities

The N1 Grid Engine 6 permits the administrator to use both built-in measures (such as system load, memory, and so on) as well as custom metrics (such as floating license, shared storage, and more) to manage the utilization of Grid hosts in a way that matches the site's usage policies, while preventing oversubscription of scarce resources. Configuration objects, such as cluster queues and host groups, enable the abstraction of job execution environments and host aggregations in a flexible and centralized manner, providing extremely simple management.

Hierarchical policies

The N1 Grid Engine 6 can automatically set the priorities of jobs to match a site's business policies. The basis for this prioritization can include: entitlements which are due to a user, department, or project; the urgency of a job based on the resources it requests, as well as time constraints; and a means to implement site-specific custom prioritization schemes. Weighting factors allow the administrator to determine precedence for how these policies are applied.

Detailed statistics gathering for analysis and utility billing

The N1 Grid Engine collects detailed data on resource utilization and job resource usage. Data is stored in a relational database in open, well-defined schema. (Currently supported databases are Oracle and PostgreSQL). Statistics are collected as frequently as once per minute, and include:

- Detailed job accounting information, such as submitting user, department, project, submit time, start and end times, job resource usage, and so on.
- System utilization for every compute host, as well as global resources configurable to include any metric of interest, even custom-defined ones.
- Aggregation of usage per user, department, project, and so on over any specified time range (daily, weekly, monthly, and more).

This information can be accessed using industry-standard APIs such as JDBC™, and thus can be used in third-party tools for analysis and utility-like chargeback and billing.

ARCo: Sun N1 Grid Engine accounting and reporting

The Accounting and Reporting Console (ARCo) tool comes with N1 Grid Engine 6. It allows the creation and display of reports, charts, and graphs of information collected by the N1 Grid Engine 6 system. All of this is provided in an easy-to-use Web interface. ARCo comes with several predefined reports, such as:

- Accounting per department
- Accounting per user
- Average job wait time per day

Authorized users can also edit the reports and create new ones. Basic reports can be designed using drop-down tables of parameters and conditions. More advanced reports are created by writing queries in SQL directly. Reports can be saved in PDF format for publication, or user formatted for input into a spreadsheet or other tool.

Sophisticated scheduler with advanced planning abilities

The scheduler in N1 Grid Engine 6 has sophisticated features normally found only in the most advanced HPC environments. Abilities such as resource reservation and backfilling ensure that critical resources are made available for important workloads, while simultaneously ensuring maximal utilization of these resources. A special feature is that reservations can be made for any resource, such as CPU, memory, software license, network card, and more.

Industry-standard APIs and Globus Toolkit integration

The Distributed Resource Management Application API (DRMAA) is an industry-standard, job-oriented API developed within the Global Grid Forum for the rapid integration of applications directly to a Distributed Resource Management system such as N1 Grid Engine 6. DRMAA provides a standard way of submitting a job, retrieving monitoring information, and initiating job control interactions such as a suspension.

In the N1 Grid Engine 6, DRMAA is implemented as a shared module, external to the application or the portal.

The DRMAA APIs are available for the Java™ and C/C++ language binding. The C/C++ language binding is also implemented as a shared module, external to the application or the portal. As new versions of the N1 Grid Engine are released, no code modifications or relinking will be necessary in applications and portals. With DRMAA, application builders may now present new, job-oriented services to users of their applications. This means that users no longer need to leave the application's environment to start or manage grid tasks.

For commercial-strength applications where grids are linked via the Internet in Wide Area Networks (WANs), Sun iForce partners offer an integration service with the industry-standard Globus Toolkit. See the N1 Grid Engine Web site for details.

Grid Engine Management Module

The Grid Engine Management Module (GEMM) provides a means to perform N1 Grid Engine 6 maintenance operations from within the context of other system management tools within

the N1 software product family. The current version of GEMM is designed to work with Sun Control Station (SCS) 2.2 and 2.2.1, a premier system management solution designed to be very lightweight and easy to use. It provides a comprehensive set of features, including:

- Web-based interface
- Ability to provision the Solaris OS and Linux to SPARC®, Opteron, and x64 or x86 hardware
- Health, performance, and inventory monitoring
- Software and package management

GEMM enables these tasks within the Sun Control Station 2.2 and 2.2.1 Web interface:

- Automatic deployment and uninstallation of N1 Grid Engine 6 to any managed host
- Monitoring of the health of the N1 Grid Engine 6, including jobs, hosts, queues, and daemons
- Web-based inspection of job and daemon status files, for diagnostics

Licensing

N1 Grid Engine 6 is licensed based on the number of total processors and master agents in a grid. Enterprise-wide licenses are available. All software licenses for N1 Grid Engine 6 are valid for Update 4. No additional licensing fees apply.

N1 Grid Engine 6 Update 4 offers attractive trade-up licenses to users wishing to incorporate more CPUs or Microsoft Windows hosts into the grid.

Enterprise-wide licenses are available and convenient. They allow multinational corporations at any location worldwide to use the N1 Grid Engine 6 for up to 120,000 CPUs and unlimited masters.

Serious software made simple

Sun provides a complete portfolio of affordable, interoperable, and open software systems designed to help you maximize the utilization and efficiency of your IT infrastructure. Built from the secure, highly available foundations of UNIX® and Java technology, these systems deliver implementations that are preintegrated and backward compatible. Sun's portfolio consists of Solaris and Linux software for SPARC® and x86 platforms, Sun N1™ software, and the Sun Java System — five integrated software systems for the data center, the desktop, the developer, mobile devices, and identity implementations.

The Java System is a radical new approach that changes forever the way businesses acquire, develop, and manage software. Only Sun has the experience and the end-to-end portfolio to deliver such a unique and industry-revolutionizing strategy. With the Java System, network services and critical business applications are up and running faster, easier, and at a lower cost than ever before, so you can focus on innovation, competition, and bottom-line results.

About Sun Microsystems, Inc.

Since its inception in 1982, customers have continually turned to Sun to help them grow their business, lower their costs, and gain competitive advantage. Sun is a leading provider of industrial-strength hardware, software, services, and technologies that make the Net work.

Sun N1™ Grid Engine 6

N1 Grid Engine 6 system recommendations

Master host

- 100 MB of free memory minimum
- 500 MB of free disk space minimum

Execution host:

- 20 MB of free memory minimum
- 500 MB of free disk space minimum *
 - * If the host is running a supported Microsoft Windows OS or
- 100 MB of free disk space minimum **
 - **If the host is running a supported UNIX OS

Database server recommendations

- 200 MB to 750 MB of free memory; 10 GB of free disk space (minimum)
- Sun Web Console: 200 MB of free memory (minimum); 250 MB of free disk space (minimum)

Operating systems and platforms

Operating systems

- Apple Macintosh OS/X 10.2.x or later
- Compaq Tru64 Unix 5.0, 5.1
- Hewlett Packard HP-UX 11.x
- IBM AIX 4.3, 5.1
- Linux x86, kernel 2.4x, 2.6x, glibc >= 2.2¹
- Linux AMD64 (Opteron), kernel 2.4x, 2.6x, glibc >= 2.2
- Silicon Graphics IRIX 6.5
- Solaris 10, 9, 8, and 7 OS (SPARC Platform Edition)
- Solaris 10, 9, 8 OS (x86 Platform Edition)
- Solaris 10 OS (x64 Platform Edition)

- Microsoft Windows Server 2003^{2,3}
- Microsoft Windows XP Professional with Service Pack 1 or later^{2,3}
- Microsoft Windows 2000 Server with Service Pack 3 or later^{2,3}
- Microsoft Windows 2000 Professional with Service Pack 3 or later^{2,3}

Note¹: Includes Linux Red Hat release 4.0

Note²: The following Microsoft operating systems are NOT supported: Windows 95, Windows 98, Windows Millennium Edition, Windows XP Home Edition, Windows NT Workstation, and Windows NT Server.

Note³: Requires Microsoft Windows Services for UNIX 3.5

Supported databases

- PostgreSQL 7.3.x and 7.4.x
- Oracle9i

Web browsers

- Netscape™ 6.2.2, 7.x or later
- Mozilla™ 1.2 and later
- Internet Explorer 5 and later
- Firefox 1.0 or later

Sun Web Console Operating Systems

- Red Hat Linux x86
- Solaris 10, 9, and 8 OS (SPARC Platform Edition)
- Solaris 10, 9, and 8 OS (x86 Platform Edition)

Sun Control Station Operating System

- Server support
 - Red Hat Linux 8.0; Enterprise Linux 2.1 and 3
 - SUSE LINUX 9.0

Learn More

Get the inside story on the trends and technologies shaping the future of computing by signing up for the Sun Inner Circle program. You'll receive a monthly newsletter packed with information, plus access to a wealth of resources. Register today at sun.com/joinic.

- Agent support (Grid Engine Module Management OS support):
 - Solaris 9 OS (x86 Platform Edition)
 - Red Hat Linux 7.3, 8.0, and 9.0; Enterprise Linux 2.1 and 3
 - Fedora Core
 - SUSE LINUX 9.0

Note: Microsoft Windows platform is not yet supported. For an updated list of supported OS, consult the Sun Control Station datasheet at sun.com/software/controlstation/ds_controlstation.xml

For more information

To learn more about N1 Grid Engine 6, visit sun.com/gridware.