Sun™ Grid Engine

Make the most of your compute resources

With installations at thousands of production sites worldwide, Sun™ Grid Engine is one of the most successful software products for managing workloads in high-performance computing. It enables organizations to build grids that help make employees more productive than ever before by dramatically reducing the time it takes to answer large, complex questions — from weeks or months to a matter of days.

Highlights

- Advance Reservation — request and reserve grid resources
- Multiclustering with Service Domain Manager — share resources across two or more clusters
- Array task dependency — create finer-grained dependencies among array jobs
- Improved scalability — enable greater throughput with scalability up to 63,000 core CPUs
- Enhanced functionality for the Solaris™ 10 OS
- Support for virtually all commercial operating systems in commercial use

In a network of many computers with users submitting jobs from multiple compute-intensive applications, it is a challenge to manage where each job should go. As one of the world's most widely used distributed resource management (DRM) software solutions, Sun Grid Engine manages compute resources in clusters of tens of thousands of CPUs based on policies and service-level agreements.

The popularity and ubiquity of Sun Grid Engine software is due in part to the open-source Grid Engine project, which was introduced in 2001. One of Sun's first open-source projects, it currently has thousands of production users worldwide. For organizations that need support not offered for open-source code, Sun offers the updated commercial version Sun Grid Engine, which includes additional innovations from Sun and the worldwide community as well as the assurance of support contracts, licenses, subscriptions, and advanced features.

The new Sun Grid Engine 6.2

Today, the new capabilities of Sun Grid Engine 6.2 software reflect the expansion of user needs from pure research to the commercial space, with applications now requiring the power of many computers to solve a single problem.

Sun Grid Engine 6.2 software offers scalability of up to 63,000 cores, enabling the far greater speed and higher level of productivity that more and more organizations require. To reach this level of scalability, Sun Grid Engine includes a new-scheduler-as-thread in the grid master node, improved interactive job support, streamlined communications, and other "large" cluster improvements. It gives users the ability to reserve grid resources in advance — including hosts, memory, and licenses — for a specified time window. While the essence of grid workload management software is maintaining the maximum utilization, there are some situations — such as a solar eclipse for astronomers or a heavy trading day for a securities dealer — when there is a need to have all the resources prepared in advance and not have to wait for the system to assign them. At the request of Sun's top commercial and research customers, this capability is included in Sun Grid Engine 6.2.

The Service Domain Manager (SDM) module included in Sun Grid Engine 6.2 software distributes resources between different services according to configurable service-level agreements (SLAs). SDM functionality enables users to manage resources for all types of scalable services. It also delivers multicluster capabilities, enabling grids at different locations to share resources. If the required resources are not available inside a specific grid, Sun Grid Engine can locate them elsewhere, increasing productivity and making the most of available compute resources.

An optional module of Sun Grid Engine 6.2 software is the Accounting and Reporting Console (ARCo), which enables users to gather live reporting data from the Sun Grid Engine system and store the data for historical analysis in the reporting database, which is a standard SQL database. The program, which supports the PostgreSQL, Oracle, and MySQL™ database systems, reads the raw data in the reporting file and writes it to the SQL reporting database, where it can be accessed by ARCo. In support of Sun Grid Engine software's multiclustering feature, ARCo can now create cross-cluster queries. And with Sun Grid Engine 6.2, the query response time is dramatically improved, enabling complete reports of activities and accounting of the shared resources.

What's new in Sun Grid Engine 6.2 software?

Flexible administration

Sun Grid Engine 6.2 software permits system administrators to use both built-in measures (such as system load and memory) as well as custom metrics (such as floating licenses and shared storage) to manage the utilization of grid hosts in a way that matches the site's usage policies while preventing oversubscription of scarce resources. This wide range of tools and parameters enables administrators to squeeze the maximum performance in the grids they manage and ensures that requests are prioritized appropriately.

Advance Reservation

Advance Reservation enables users to request and reserve grid resources in advance — like hosts, memory, and licenses — for a specified time window. The scheduler makes sure that the resources needed will be in place and available for the duration of the reservation. This feature enables users to coordinate grid resource availability with external factors such as individuals' schedules, equipment availability, or facility availability.

Multiclustering with Service Domain Manager

The new Service Domain Manager component of Sun Grid Engine software enables two or more Sun Grid Engine clusters to share resources. Using service-level objects (SLOs) as a guideline, SDM will migrate resources from underutilized, low-priority clusters into heavily loaded, higher-priority clusters.

SDM enables lateral scalability by adding new clusters alongside existing clusters without limiting flexibility or creating artificial resource boundaries. As workload conditions change, resources will migrate to where they are needed most.

Array task dependencies

Sun Grid Engine software's new array task dependencies enable users to increase productivity and decrease workflow times by creating finer-grained dependencies among array jobs. Previously, every task of an array job was blocked until each task of the previous array job in the workflow was completed. Now, individual tasks of one array job can depend on individual tasks of another array job. This new functionality allows array job workflows to execute with a much higher degree of parallelism, producing results faster and reducing time to market.

Improved scalability

With the ability to scale up to 63,000 core CPUs, Sun Grid Engine software enables throughputs not previously possible. For example, the Texas Advanced Computing Center's Ranger system, the fourth-largest supercomputer in the world, successfully runs a 60,000-core parallel application for a facial recognition application on Sun Grid Engine software and Sun hardware.

Enhanced functionality for the Solaris 10 Operating System

The Solaris™ 10 Operating System provides administrators with a wide range of tools to make their jobs easier, and with Sun Grid Engine 6.2 software, administrators can now leverage more of those tools automatically. Sun Grid Engine also runs on almost every commercial operating system in commercial use, including Linux, Microsoft Windows, and Mac OS X.

Licensing

With Sun Grid Engine, customers can choose between a subscription offering or a perpetual license.

An annual subscription of Sun Grid Engine software is available with either Standard Support or Premium Support. Subscriptions include both the license and the support for the duration of the subscription. The subscriptions are priced either per socket, with unlimited cores each, or per organization size (small, medium, or large), with unlimited sockets. For organizations wanting to spread the investment equally over a few budget years, subscriptions are the ideal choice.

Perpetual licensing is offered for either per socket or for unlimited sockets in an enterprise. Perpetual license entitlements are a one-time payment, but a support service contract will need to be purchased separately. Organizations purchasing a perpetual license must also purchase a support contract (which must be renewed each year) in order to have the same level of support as the subscription offering.

Sun supports Sun Grid Engine 6.2 software on any socket from any manufacturer as long as it is on one of the supported platforms.

Organizations using open-source Grid Engine software may also purchase a subscription and install the commercial version. The commercial version will have additional features not available in the open-source version that provide a range of benefits for enterprises needing a secure, scalable, and mission-critical cluster.

Sun Grid Engine companion software

Sun offers a complete portfolio of affordable, interoperable, and open software products that provide additional benefits to organizations deploying Sun Grid Engine and help maximize utilization and efficiency.

Solaris 10 Operating System

A free and open UNIX® operating system, the Solaris 10 OS is supported on more than 1,000 x64/x86 and SPARC® platforms. The Solaris OS includes unique features such as Dynamic Tracing (DTrace), ZFS™, Predictive Self-Healing, and Solaris Containers.

Although DTrace is not part of Sun Grid Engine software, it is the backbone of the Sun Grid Engine 6.2 master monitor. This monitor was designed to be used by Sun Grid Engine administrators as a first-aid kit for bottleneck analysis. It also helps with the tuning of the Sun Grid Engine master, which is the most performance-sensitive component in the grid environment. The monitor can be deployed even in the largest production environments with high throughput, since data sampling works almost contact-free.

Sun Shared Visualization software

Based on the open-source VirtualGL project, Sun Shared Visualization software provides users on a variety of client platforms the ability to remotely access and share 3D-accelerated applications.

Sun Scalable Visualization software

Based on open-source software, Sun Scalable Visualization software enables systems to work together on a single application to scale in compute density, display resolution, and polygonal complexity.

Sun StorageTek™ QFS software

A robust file system that delivers maximum scalability, Sun StorageTek™ QFS software provides high-performance, heterogeneous data access over a SAN and is ideal for collaborative environments where large data volumes require shared access.

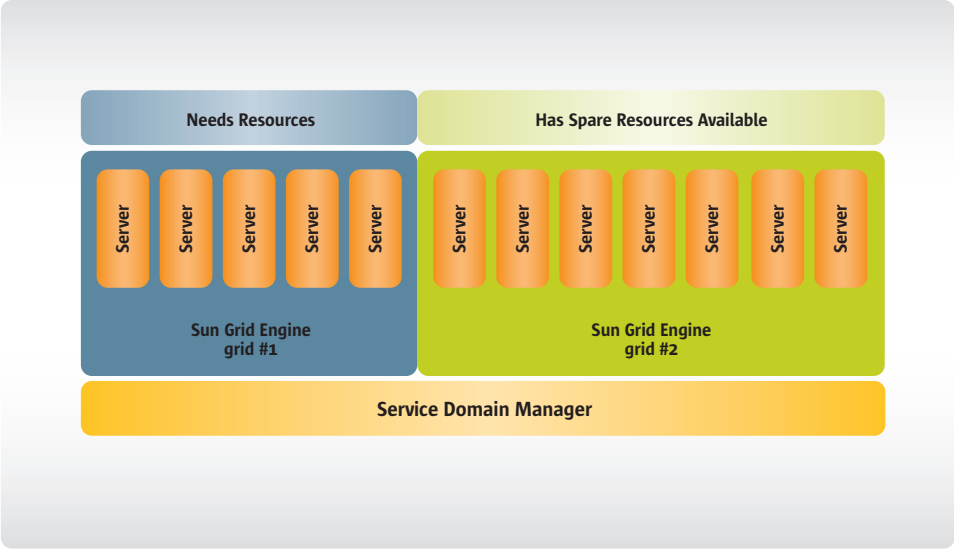

Greater scalability is possible through Sun Grid Engine 6.2 software's Service Domain Manager, which migrates resources among Sun Grid Engine clusters to accommodate workload.

Sun StorageTek Storage Archive Manager

To help organizations effectively manage and utilize data according to business requirements, Sun StorageTek Storage Archive Manager (SAM) software provides data classification, centralized metadata management, policy-based data placement, protection, migration, long-term retention, and recovery.

Lustre™ file system

Designed to meet the demands of the world's largest high-performance compute clusters, the object-based Lustre™ file system scales to tens of thousands of nodes and petabytes of storage with groundbreaking I/O and metadata throughput.

Sun HPC ClusterTools™ software

Offering a comprehensive set of capabilities for parallel computing, Sun HPC ClusterTools™ software is an integrated toolkit that enables developers to create and tune message-passing interface (MPI) applications that run on high-performance clusters and symmetric multiprocessors (SMPs).

Sun Studio

Sun Studio is the premier development environment for the Solaris OS. With its optimizing C, C++, and Fortran compilers, combined with the world-class NetBeans™ IDE and other performance tools, it delivers outstanding performance for single-threaded and multithreaded application development on the latest Sun hardware platforms and now also includes a complete development toolset for the Linux OS.

Solaris Cluster

Solaris Cluster is a multisystem, multisite disaster recovery solution that manages the availability of applications, services, and data across local, regional, and vastly dispersed datacenters. It includes a special agent that makes Sun Grid Engine highly available for disaster recoveries.

MySQL software

MySQL Enterprise Server software is the most reliable, secure, and up-to-date version of MySQL software for cost-effectively delivering production database applications. The Sun Grid Engine ARCo module requires an external SQL database and fully supports MySQL database software.

Sun xVM Ops Center

A powerful tool for datacenter automation and virtualization management of servers, Sun xVM Ops Center provides a highly scalable, unified management platform for both physical and virtual environments.

System requirements

Operating systems

- Sun Solaris 10, 9, and 8 OS (SPARC platforms)
- Sun Solaris 10 and 9 OS (x86 platforms)
- Sun Solaris 10 OS (x64 platforms)
- Apple Mac OS X 10.5 (x86 platforms)
- Apple Mac OS X 10.4 (PowerPC platforms)
- Apple Mac OS X 10.4 (x86 platforms)
- Hewlett-Packard HP-UX 11.00 or higher (including HP-UX on IA64)
- IBM AIX 5.1, 5.3
- Linux x86, kernel 2.4, 2.6, glibc >= 2.3.2[1]
- Linux x64, kernel 2.4, 2.6, glibc >= 2.3.2[2]
- Linux IA64, kernel 2.4, 2.6, glibc >= 2.3.2[1]
- Microsoft Windows Server 2003[3, 4]
- Windows XP Professional with Service Pack 1 or later[3, 4]
- Windows 2000 Server with Service Pack 3 or later
- Windows 2000 Professional with Service Pack 3 or later

Master host

- 50 MB for each binary platform
- 80 MB of free memory minimum
- 100 MB of free disk space minimum

Execution host

- 50 MB for each binary platform
- 20 MB of free memory minimum
- 50 MB of free disk space minimum

Database server

- 50 MB for each binary platform
- Minimum 200 MB to 750 MB of free memory
- 10 GB of free disk space minimum

Sun Web Console

- 50 MB for each binary platform
- 200 MB of free memory minimum
- 250 MB of free disk space minimum

Databases supported for ARCo[5]

- PostgreSQL 8.0 through 8.3
- MySQL 5.0
- Oracle9i or Oracle Database 10g

Operating platforms supported for ARCo

- Sun Solaris 10, 9, and 8 OS (SPARC platforms)
- Sun Solaris 10, 9, and 8 OS (x86 platforms)
- Linux RPM distribution

Supported Sun Java Web Console version 3.0.x Web browsers

- Netscape™ 6.2 and above
- Mozilla 1.4 and above
- Internet Explorer 5.5 and above
- Mozilla Firefox 1.0 and above

Learn More

To learn more about Sun Grid Engine 6.2, visit sun.com/gridware or read the *Beginner's Guide to Sun Grid Engine 6.2 Installation and Configuration* white paper at sun.com/offers/details/Sun_Grid_Engine_62_install_and_config.html.

Service Domain Manager (SDM) platforms supported

- Solaris 10, 9, and 8 OS (SPARC platforms)
- Solaris 10 and 9 OS (x86 platforms)
- Solaris 10 OS (x64 platforms)
- Apple Mac OS X 10.5 (x86 platforms)
- Apple Mac OS X 10.4 (x86 platforms)
- Linux x86, kernel 2.4, 2.6, glibc >= 2.3.2
- Linux x64, kernel 2.4, 2.6, glibc >= 2.3.2

SDM Required Software[6]

- Sun Grid Engine 6.2

1. For Linux on x86 and IA64 platforms, the openmotif-2.2.2 package or higher is required.

2. For Linux on x64 platforms, the openmotif-2.2.3 packages or higher is required.

3. 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.

4. Requires Microsoft Windows Services for UNIX 3.5.

5. If you have multiple clusters, one dbwriter installation per cluster is needed, but only one Reporting installation is needed for all clusters. See ARCo documentation.

6. SDM is required for multicluster functionality. For more information, see SDM documentation.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

SunWIN #548275 Lit. #SWDS14648-0 10/08