Sun Grid Engine: Towards Creating a Compute Power Grid

Wolfgang Gentzsch
Sun Microsystems Inc,
Palo Alto, California, USA
wolfgang.gentzsch@eng.sun.com

1. Introduction

This presentation provides an up-to-date overview of the Sun Grid Engine distributed resource management software and of future plans to enhance it towards and integrate it into Computational Grid Environments.
The main scope of many of today's Grid research projects is to understand and demonstrate the actual potential of networking, computing and software infrastructure for computational Grids and to build proof-of-concept middleware [1]. This has resulted in Internet infrastructure projects like Globus [2], Legion [3], and Punch [4] which enable users to combine a set of distributed resources into one integrated Grid Workbench to allow users to measure nature (e.g. with microscope, telescope or particle accelerator), process the measured data according to fundamental mathematical equations (e.g. the Navier-Stokes or Maxwell equations), and provide computer simulations and animations to study and understand these complex phenomena.

Most of the underlying sophisticated technologies for computational Grids are currently under development. Large research communities like the Global GridForum are coordinating all kinds of Grid-related activities. Research in resource management is underway in projects like EcoGrid [5], and the basic building blocks for commercial Grid resource managers exist, e.g. Sun's Grid Engine software [6].

2. Sun Grid Engine Distributed Resource Management

Sun Grid Engine is a new generation distributed resource management software which dynamically matches users' hardware and software requirements to the available (heterogeneous) resources in the network, according to policies usually defined by management in an enterprise.

Sun Grid Engine acts as the central nervous system of a cluster of networked computers. Via so-called daemons, the Grid Engine Master supervises all resources in the network to allow full control and achieve optimum utilization of the resources available. Grid Engine has been developed as an enhancement of Codine from former Genias GmbH and Gridware Inc, according to requirements from many early customers, like the Army Research Lab in Aberdeen, and BMW in Munich, where today Grid Engine manages over 800 powerful compute servers in each of these local enterprise Grids. With Grid Engine, average usage increased from well under 50% to over 90%, in both environments.

Sun Grid Engine aggregates the compute power available in dedicated compute farms, networked servers and desktop workstations, and presents a single access point to users needing compute cycles. This is accomplished by distributing computational workload to available systems, simultaneously increasing the productivity of machines and application licenses while maximizing the number of jobs that can be completed.

With Grid Engine software, technical enterprises can shorten time to market and discovery as well as complete research more quickly, while increasing innovation and quality, because they have:

- Increased available compute power
- Ability to manage more complex problems
- More flexible access to resources.

In addition, Sun Grid Engine software helps lower the costs of purchasing, installing, setting up and administering the computing environment because it allows:

- Maximized use of new/existing resources
- Lower administration costs
- Lower upgrade costs
- More efficient reuse of existing legacy resources.

Within the above framework, Sun Grid Engine contributes by virtue of its ability to:

Proceedings of the 1st International Symposium on Cluster Computing and the Grid (CCGRID '01)
0-7695-1010-8/01 $10.00 © 2001 IEEE

- Include all resources in the resource pool
- Help assure maximum possible resource utilization by keeping systems running at up to 98% productive utilization
- Match utilization with task priority.

Currently, Sun Grid Engine is available for Solaris and Linux clusters, extending the number of technical enterprises that can benefit from distributed resource management software. The Grid Engine for Linux builds on the success of Sun Grid Engine software for the Solaris Operating Environment -- which has been downloaded by over 1000 organizations in 70 countries since its launch in September 2000. Both the Solaris and Linux version of Sun Grid Engine software use the same source code base. Sun Grid Engine software can now be used in Solaris environments, Linux environments, and mixed Solaris/Linux environments. Consistent with Sun's own computing vision, Linux uses open standards and non-proprietary interfaces. Sun's Solaris Operating Environment and Linux are both driving growth, innovation, and success of UNIX and network computing. This leads to:

- More innovative and compelling new applications
- A large and growing talent pool of knowledgeable administrators and developers
- A broad range of computing solutions from multiple vendors.

Sun's vision for Grid Engine technology and the benefits it brings to technical computing, is to become mainstream, with its free web-download distribution, [6], open source model, and standards-based API. With availability on the Linux platform, Sun Grid Engine software can be deployed on the majority of UNIX systems on the market today, providing the massive scalability and agility needed to help companies keep pace with and capitalize on the Net Effect.

3. Sun Grid Engine and the Compute Power Grid

Sun's next step is to enhance Grid Engine, which currently is restricted to manage local networked computer resources, towards 'The Grid Broker'. The Grid Broker will be able to match the user's compute jobs with the available resources in a wide-area network, including accounting and invoicing users for the CPU power they consume and the quality of services they get, very much like todays electric power consumption, telephone usage or water supply. The ultimate Grid Broker, based on Genias/Gridware's GRD Global Resource Director, will match the user's requirements to the best fitting Application Service Provider (ASP) in the universe which optimally fulfills the user's hardware, software and service needs.

This GRID Broker belongs to the enabling technologies of the next Internet Age. The Internet, for a long time, has been used just for providing and getting information. Only recently, enabled by several important improvements in hardware infrastructure, security, authentication, and ease of access, the Internet is used for electronic commerce. And just now, the next revolutionary step in the Internet can be foreseen: The Grid Computing Infrastructure. In this infrastructure all kinds of dedicated Grids will be used for collaboration and collaborative computing in industry and research, for application simulation and animation, for real-time video, on-demand virtual reality presentations, and other services for consumers and producers.

References

[1] M. Baker, R. Buyya, D. Laforenza, *The Grid: International Efforts in Global Computing*, International Conference on Advances in Infrastructure for Electronic Business, Science,and Education on the Internet (SSGRR 2000), Rome, Italy, July 31 - August 6. 2000.

[2] I. Foster and C. Kesselman, *Globus: A Metacomputing Infrastructure Toolkit*., Intl Journal of Supercomputer Applications, Volume 11, No. 2, 1997.

[3] S. Chapin, J. Karpovich, A. Grimshaw, *The Legion Resource Management System*, Proceedings of the 5th Workshop on Job Scheduling Strategies for Parallel Processing, April 1999.

[4] N. Kapadia and J. Fortes, *PUNCH: An Architecture for Web-Enabled Wide-Area Network-Computing*, Cluster Computing: The Journal of Networks, Software Tools and Applications, September 1999.

[5] R. Buyya, J. Giddy, D. Abramson, *A Case for Economy Grid Architecture for Service-Oriented Grid Computing*, The 10th IEEE International Heterogeneous Computing Workshop (HCW 2001), San Francisco, USA, April 2001. http://www.buyya.com/ecogrid/

[6] Sun, Grid Engine, http://www.sun.com/Gridware