

# Grid Computing: A New Technology for the Advanced Web

**Wolfgang Gentzsch, Sun Microsystems, Inc.**

Palo Alto, USA, August 23, 2001

Mankind is right in the middle of another evolutionary technological transition which once more will change the way we do things. And, you guessed right, it has to do with the Internet. It's called "The Grid", which means the infrastructure for the Advanced Web, for computing, collaboration and communication.

## **Grid Components: Networks, Computers, Software**

The Internet itself has dramatically changed over the last three decades. While, in the late Sixties, it was built mainly to provide scientists with an infrastructure for faster communication via electronic mail, it has rapidly grown and improved since then, mainly because of three driving elements: networks, computers, and software.

In the mid Nineties, George Gilder predicted the "Network Abundance". Every nine months, total network bandwidth doubles. For many years, every day, thousands of miles of fiberoptic cables are laid down, ready for deployment in private and business applications. While many of us still suffer from 56 bps (Bits/sec) telephone modem speed, researchers and enterprises have already access to networks with a bandwidth of 10 million bps, and some even up to one billion bps. Soon, we will see 100 billion bps and more. Thus, network bandwidth will grow by a factor of 5000 over the next 10 years. Just this July, the US National Science Foundation approved the \$53 million DTF Distributed TeraScale Facility project, a network with 40 billion bps, connecting research centers in San Diego (SDSC), Pasadena (Caltech), Urbana-Champaign (NCSA), and Chicago (ARNL). It seems that there is no limit.

Another building block of the Internet are the computers. Today, for example, there are over one hundred million PCs in homes and at work, plus some 10 million powerful compute servers, from midrange to high-end, used at ISP Internet Service Providers, or for high-performance scientific, engineering and commercial applications. Their performance doubles every 18 months, which was observed and predicted by former Intel Chairman Gordon Moore in 1965. Thus, computer performance will grow by a factor of 100 over the next 10 years, then breaking the Petaflops Performance Barrier. This will make today's handheld electronic devices soon very powerful nodes in the Internet.

The third and most complex Internet building block is software, either for running the networks and the computers and their intercommunication – then called the *middleware*, or for solving our day-to-day problems and running our business, called *application* software. The ever increasing benefit, resulting in the combination of the networked computers, the software to run them, and the people who use them, is called Metcalfe's Law, after Bob Metcalfe, who developed Ethernet at Xerox Parc, in 1973: "The usefulness of a network equals the square of the number of users."

These three laws of Gilder, Moore, and Metcalfe, respectively, and the technological evolution they describe, are currently converging into and enabling the Advanced Web, on top of the Internet infrastructure. Past Internet and World Wide Web mainly enabled information provision, retrieval and exchange, and some e-commerce.

The new Advanced Web adds a wide variety of opportunities, based on computing, collaboration and communication, for individuals, groups, research and engineering teams, and for the whole community. It will provide great services in our private, community, and business environments. Universal connectivity gives users immediate and easy access to any kind of information and service they want, helps them in solving problems and in making personal and business decisions, and allows them to easily offer their own services to anybody. The new Advanced Web changes the way we live and work.

Enter *The Grid*. The term has been derived from the "Power Grid" infrastructure which provides electricity to every wall socket. In our context, The Grid describes the technology infrastructure for the Advanced Web, for computing, collaboration and communication.




Figure 1: Evolution of the Grid

Cluster Grid	Campus Grid	Global Grid
Local clusters/farms deployed on a departmental or project basis, one owner	Merging cluster grids into one campus or enterprise grid, multiple projects and multiple owners	Merging campus grids into a global grid across organization, multiple sites

## Distributed Computing and the Grid

In the early Nineties, research groups started exploiting distributed computing resources over the Internet: scientists collected and utilized hundreds of workstations for parallel applications like molecular design and computer graphics rendering. Other research teams glued large supercomputers together into one virtual metacomputer, distributing subsets of a meta-application to specific vector, parallel and graphics computers, over wide-area networks, e.g. the computer simulation of multi-physics applications like the interaction of a fluid with a rotating propeller blade. Additionally the scope of many of these research projects was to understand and demonstrate the actual potential of the networking, computing and software infrastructure and to develop it further.

Multiprocessor Systems (MPs), Clusters, Grids are examples of distributed computing architectures. In MPs, processors are tightly coupled, through shared memory or high-speed interconnect (e.g. crossbar switch). Examples are PVPs (parallel vector processors). They are most suitable in HPC High Performance Computing, for parallel applications which rely on fast message passing communication among their parallel processes.

Clusters, on the other hand, are loosely coupled single or multiprocessor computers, interconnected through networks which are one or two orders of magnitude slower than MP interconnects. Examples are Beowulf clusters made of commercial off-the-shelf hardware and running Linux; or the Sun Technical Compute Farm (TCF), running Solaris/TM Operating System. They are mostly used for heavy throughput computing, distributing many (usually non-parallel) compute jobs onto the processors, collecting individual results back into one global result (space). Examples are in the film industry, rendering of thousands of frames to produce a movie, or the design and test simulations to build the next generation VLSI chip in EDA Electronic Design Automation. Or in bioinformatics, scanning hundreds of thousands of sequences in genomics and proteomics.

While MPs and Clusters are single systems, usually in one single administrative domain, Computational Grids consist of clusters of networked MPs and/or Clusters, located in multiple different administrative domains, scattered over departments, enterprises, or distributed globally even over the Internet. Naturally, therefore, these grids involve a much higher degree of complexity, especially at the middleware layer, to run, administer, manage, and use these distributed computing resources, and on the application layer, to design, develop and run the appropriate software which efficiently deploys such grids.

**Table: Distributed Computing Taxonomy**

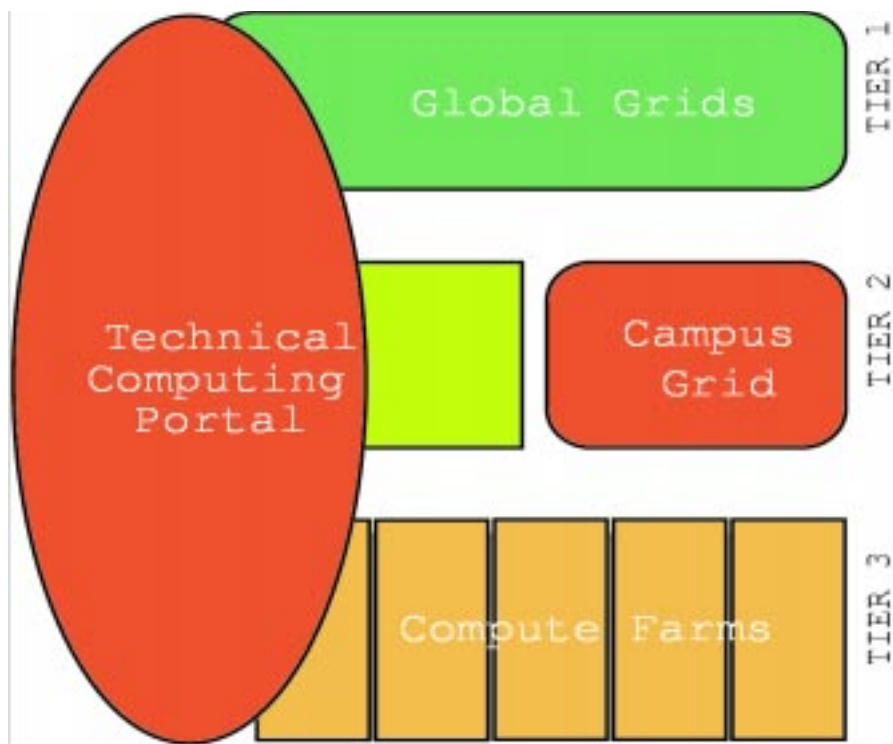
		Multiprocessor	Cluster	Campus & Global Grid*
1	Hardware	One box	Rack or boxes, one admin domain	Multiple boxes in multiple admin domains
2	Network	Tightly coupled	Loosely coupled	Via Intranet or Internet
3	Components	Proprietary	Common, off the shelf	Mix
4	Architecture	Homogeneous, single system	Homogeneous, single system image	Heterogeneous, multiple images
5	Operating System	Proprietary or standard Unix	Standard Unix + RMS**	Standard Unix + Grid RMS**
6	Communication	Low latency, high bandwidth	... In between ...	High latency, low bandwidth
7	RMS Resource Management	E.g. Solaris Resource Manager	Sun Grid Engine, PBS	GRD (Campus), Globus, Legion (Global)
8	Cost/price	Expensive buy	Inexpensive buy	"Rent" CPU cycles
9	Computing paradigm	Parallel and supercomputing	Throughput computing	Both; metacomputing, collaborative computing
10	Users	One/more groups	One group/owner	Groups, Communities
11	Examples	NEC SX5, Sun Starfire/tm	Beowulf, Sun TCF	Campus , Enterprise,...

\*) Current state-of-the-art

\*\*\*) Resource Management Software

In short, The Grid is a distributed computing architecture for delivering computing and data resources as a service, over the Internet, in much the same way that electricity is delivered over the power grid. It is the next logical step in the technology infrastructure, which connects distributed computers, storage devices, mobile devices, instruments, sensors, data bases, and software applications, and provides uniform access to the user community for computing, collaboration and communication. Examples of current grids are the NASA Information Power Grid (IPG); the DoD Distance Computing and Distributed Computing Grid (DisCom 2); the NSF NCSA National Technology Grid; NetSolve for accessing and sharing mathematical software; Nimrod for campus-wide resource sharing; SETI@Home for searching for extraterrestrial intelligence; the CERN DataGrid, processing Petabytes of particle data per year from its Large Hadron Collider experiment; or the APGrid connecting many computer centers in Asia and the Pacific Rim, in the near future.

Figure 2: Three Grid Layers: The Global, Campus and Cluster Layer



A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to computational capabilities. In the near future, these grids will be used by computational engineers and scientists, experimental scientists, associations, corporations, environment, training and education, states, consumers, etc. They will be dedicated to on-demand computing, high-throughput computing, data-intensive computing, collaborative computing, and supercomputing, potentially on an economic consumer/supplier basis. Grid communities, among others, are national grids (like ASCI), virtual grids (e.g. for research teams), private grids (e.g. for a car manufacturer's CrashNet and its suppliers, for collaborative crash simulations), and public grids (e.g. Consumer networks).

These grids enable users to combine nearly any set of distributed resources into one integrated metacomputing workbench to allow users to measure nature (e.g. with microscope or telescope), process the data according to some fundamental mathematical equation (e.g. the Navier-Stokes equations), and provide computer simulations and animations to study and understand these complex phenomena.

Today, we see the first efforts to more systematically exploit these grid computing resources over the Internet. So called peer-to-peer computing projects, like SETI@home, Distributed.Net, and Folderol, let Internet users download scientific data, run it on their own computers using spare processing cycles, and send the results back to a central database. Recently, an academic project called Compute Power Market, has been initiated to develop software technologies that enable creating grids where anyone can sell idle CPU cycles, or those in need can buy compute power much like electricity or telephony today.

### **Grid Computing Challenges**

Most of the underlying sophisticated technologies for grids are currently under development. Prototype grid environments exist like public-domain projects Globus and Legion. Research in resource management is underway in projects like EcoGrid, and the basic building block for a commercial grid resource manager exists with Sun Grid Engine software.

The GGF (Global Grid Forum), founded in 1998, unites hundreds of computer scientists in working groups to discuss the one common grid architecture. Some of the challenges they are addressing are:

- developing application software for the grids
- identifying and accessing suitable computing resources in a distributed environment
- defining standard interfaces to enable communication among the different grid building blocks, and to facilitate application development
- guaranteeing authenticated access and secure data transfer
- providing service tools for monitoring, accounting, billing and reporting
- designing network protocols for message formats and exchange.

Security in grids is a particularly difficult problem. Resources being used may be extremely valuable and are often located in distinct administrative domains. The application being solved may be extremely sensitive, belong to a company's most valuable and secret assets. Therefore, the users need to have a "key" to the resources and the data, with uniform authentication and authorization. But there should be only one single sign-on, even in case of using hundreds of distributed resources to solve one complex problem. Existing software security standards like SSL and X.509 are being enhanced to achieve this. Current grid software environments like Globus and Legion already contain such a grid security infrastructure.

Another hard problem in Computational Grid environments is distributed resource management (DRM). Starting with networked workstations and with client/server environments, in the early Nineties, the objective then was to use these networked computers more efficiently, from an average of some 20% in unmanaged environments, up to 98% average usage in environments today controlled by DRM software like Sun Grid Engine, Condor, LSF, and PBS.

Figure 3: Every Grid Needs an Easy-to-Use Grid Access Portal

**Technical Computing Portal**  
powered by Sun Grid Engine

iPlanet™  
e-commerce solutions

iPlanet Portal Server 3.0      Home | Options | Content | Layout | Help | Log Out

**User Information** [ ? Edit X ]

**Welcome!**  
Frederic Pariente  
Last Update:  
April 26, 2001 4:25:00 PM PDT  
113 minutes left  
30 minutes max idle time

**SunTCP Job List** [ ? X ]

You have no running jobs.

[Submit new job...](#)

**Bookmarks** [ ? Edit X ]

Enter URL Below:

[Sun home page](#)  
[Install NetMail on Client](#)

**SunTCP Project List** [ ? X ]

- [Truck Benchmark](#) [edit](#) [delete](#)
- [Dummy project](#) [edit](#) [delete](#)
- [American option pricing](#) [edit](#) [delete](#)
- [Smith-Waterman genome blast](#) [edit](#) [delete](#)
- [Ecoli blast](#) [edit](#) [delete](#)

[Create new project...](#)

**SunTCP Application Portfolio** [ ? X ]

- [PAM-CRASH V2000D](#)
- [NCBI Blast 2.1.3](#)
- [Lassap 2.0](#)
- [American Put Option Pricer](#)
- [Sleeper demo application](#)
- [Sleep Simulation](#)

**Applications** [ ? Edit X ]

- [NetMail](#)
- [NetMail Lite](#)
- [NetFile](#)
- [NetFile Lite](#)

iPlanet Portal Server 3.0      Home | Options | Content | Layout | Help | Log Out

Sun Grid Engine, for example, is a distributed resource management software which dynamically matches users' hardware and software requirements to the available computing resources in the network, according to predefined usage policies usually prescribed by management in the enterprise. Sun Grid Engine acts much like our body's central nervous system ('The Body's Internet'). The Sun Grid Engine Master ('the brain') with its sensors in every computer (comparable to the sensations of touch, sound, smell, taste, and sight) dynamically acts and reacts, according to set policies (comparable to move, eat, drink, sleep, ...) to allow for full control and achieve optimum utilization and efficiency of the available resources. Sun Grid Engine has been developed as an enhancement of Codine from the former Gridware Inc, according to well defined requirements from the Army Research Lab in Aberdeen, BMW in Munich, and hundreds of other customers. Today, Sun Grid Engine manages hundreds of powerful compute servers in each of these local grids. Average usage increased from well under 50% to over 90%, in these managed environments.

### **Grid Users, Markets, Business**

Grids are especially suitable in science and engineering. Biochemists, for example, can exploit thousands of computers to screen hundreds of thousands of compounds in an hour. Hundreds of physicists worldwide pool computing resources to analyze petabytes ( $10^{15}$  bytes) of data from CERN's Large Hadron Collider experiment. Climate Scientists visualize, annotate, and analyze terabytes of computer simulation datasets. An emergency response team couples real-time data, the actual weather model, and population data. Engineers in the automotive industry combine their multidisciplinary results with analysis data from their suppliers. A community group pools members' PCs to analyze alternative designs for a local road. And many more.

"The Grid" itself is not a piece of software, not a product which one can sell into specific technical or commercial markets. It is very much like the Web – we cannot sell the Web. The Grid is simply the evolving next generation of the Advanced Web. And like the Web, The Grid will be ubiquitous. It will simply become the basic IT infrastructure for all markets, where it makes sense. What hardware/software companies can and will provide is the appropriate hardware and software stack – the gridware – which helps organizations meet their IT requirements on top of these grids.

We are ready to build such grids today. But there is no commonly agreed standard, there is no common architecture yet. Currently, there is still a lot of manual work involved to set up and run a prototype grid for computing, collaboration, and communication.

Since the Grid is still a highly technology-driven software infrastructure, there are mainly three user groups: the early adopters, the early majority, and then everybody else. Currently, we are in the early-adopter phase: Users are highly technology oriented, like in research and education and in some large engineering companies. Currently, grids are being built together with these users.

To have a closer look at what kind of grid activities some companies already provide, let's take Sun Microsystems as an example.

Sun's origins are in distributed computing. Originally Sun stood for "Stanford University Network,". Scott McNealy's vision was "the network is the computer." Today, one might say that the grid is the computer. Since its beginning, Sun contributed to network computing and grid technologies, for example with Java/tm, Jini/tm, Jxta/tm, and with Sun Grid Engine and Sun HPC ClusterTools/tm technology, which both have become open source projects recently. In addition, there is the iPlanet/tm Grid Access Portal, the Solaris Resource and Bandwidth Manager, the Sun Management Center, and the LSC Distributed Storage Manager. Sun Grid Engine is integrated with the major grid technologies currently developed and deployed in the Grid community, such as Globus, Legion, Punch and Cactus. Sun is currently forming these building blocks into one integratable stack.

Sun is actively involved in building a large number of grids, Departmental Grids, Campus Grids, Research Grids, and Enterprise Grids, mainly within a customer program called "The Sun Center of Excellence". One example is the Edinburgh Parallel Computing Center (EPCC), a Sun Center of Excellence in Grid Computing. It is the location of the UK National eScience Center (along with eight regional centers). Basically, the eScience program is to build a UK-wide Grid which interconnects all these distributed computing centers to aid scientists in their research. Edinburgh will evaluate software building blocks like Sun Grid Engine, the iPlanet/tm Portal Server, Sun Management Center, Sun HPC ClusterTools, and Forte/tm Workshop to build the next generation Grid infrastructure. Through this infrastructure, EPCC will deliver compute power to and exchange expertise with its partners in research and industry all over the UK.

Another example is the Ohio Supercomputer Center (OSC), which became a Sun Center of Excellence for High Performance Computing, earlier this year. Together with Sun, OSC is building the grid infrastructure which enables distributed computing, collaboration, and communication with other partners, e.g. Ohio State University, Universities of Akron and Cincinnati, Nationwide Insurance, and Exodus.

One more example is the Technical University of Aachen, Germany, which is a Sun Center of Excellence for Computational Fluid Dynamics. Among other objectives, the Center will be providing remote access to its large Sun system (which will grow to over 2 Teraflops) for researchers on the university campus, much like an ASP Application Service Provider. Therefore, one of their grid contributions is the enhancement of open source Grid Engine toward a Grid Broker, using the software code available in the Grid Engine open source project.

## **Beyond the Web**

Soon, grids will break out into commercial and industrial computing. Grid computing is already implemented in many industrial settings, on a departmental or enterprise level. Up to now it has gone by many different names, such as compute farms, cluster computing, etc.. For example, Caprion Pharmaceuticals has recently worked with Sun engineers to implement Sun Grid Engine for proteomics, on a large server farm. Grids are also coming into the commercial setting through third parties, since many ASPs and ISPs are implementing grid-enabled access to applications. So in the end, Grid Technology will become the glue which will unite both technical and commercial computing. Just as we have "The Web" today, we then will have "The Grid".



The Grid forces us to reinvent the network itself: the data centers, the clients, the applications, and the services. Everything is changing. Everything will be seen in the context of grids – Campus Grids, Enterprise Grids, Research Grids, Entertainment Grids, Community Grids, and many, many more. The network will be service driven, the clients will be light-weight appliances with Internet or wireless access to any kind of resources. Data centers will be extremely safe, reliable, virtually always available, from anywhere. Applications will be part of a wide spectrum of services delivered over the network, such as compute cycles, tools for data processing, accounting and monitoring, with customized consulting, with additional information and communication tools, and with software which allow you to sell or trade your results.

The good news is that there will be no disruptive change in the technology or in the way we use it. For many good reasons: We currently have so many new technologies on the table that we will be busy enough to efficiently implement them. This takes time. Then, after we have interconnected all kind of computing devices through grid infrastructures, the next step will be to embed and hide these devices into any kind of thing that serves our daily business and private needs, in our houses, cars, airplanes, clothes, the environment, maybe some even in our body.

Wolfgang Gentsch is Director of Engineering and Grid Computing at Sun Microsystems, Inc. He joined Sun in July 2000, with Sun's acquisition of Gridware, a US/Germany based software company exclusively focussing on distributed computing. He was cofounder of Gridware in 1999, and founder of its predecessor Genias Software, in 1990. From 1985 to 2000, he was a professor of mathematics and computer science at the engineering college in Regensburg, Germany, and at the same time a consultant for many vector, parallel and distributed computing companies.

© 2001 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, Solaris, Starfire, iPlanet, the iPlanet logo, Java, Jini, Sun HPC ClusterTools, and Forte, are trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Compant, Ltd.