SUMMER 2001

Campus Advisor

Sun microsystems

News, Views and Trends for the 21st Century Campus

www.sun.com/edu

IN THIS ISSUE

GRID COMPUTING SPOTLIGHT

GRID COMPUTING IN EDUCATION AND RESEARCH

This issue of *Campus Advisor* illustrates several applications of "The Grid." The term, which was coined in the 1990s, refers to one of the fastest growing trends in high-end scientific and engineering computing. The Grid is a distributed computing architecture in which computing and data resources are delivered over the Internet in much the same way that electricity is delivered over the power grid.

Today, we at Sun Microsystems see at least three different "tiers" of grid computing within the academic community. These are illustrated in the chart at right. The first tier, which refers to global access and data-grids, is what brought the term "grid computing" into widespread use among academic IT professionals. In this case a number of world-class high energy physics sites created a virtual network of research centers throughout the world to analyze jointly vast amounts of data provided by big collider experiments such as the European Organization for Nuclear Research (CERN). Creating that virtual network posed a number of IT challenges, beginning with how to build a coherent data infrastructure onto heterogeneous platforms and how to transfer large data sets for authentication. IT experts turned to Sun technologies like Java™ and Jiro™ to address some of the critical requirements like open source and multi-platform support.

CIOs of large universities face a different set of issues, which Sun has responded to by establishing the second and third tiers of grid computing. These issues are largely related to compute cycles and capacity.

It is well understood that there are plenty of idle compute cycles within the campus environment. For example, a library server usually is not very active at night or on weekends, while ERP servers generally are challenged only by payroll runs, not by day-to-day operations. At the same time, universities are often hard-pressed to find enough computing capacity to meet the demands of the campus scientific community. In Sun's vision this is a second tier grid, also called the Campus Grid. This solution, Sun™ Grid Engine (formerly known as Codine from Gridware), allows the university IT manager to automate the scheduling and resource management across the campus. Creating such an infrastructure is a huge benefit for the entire campus.

The third tier of grid computing refers to technical compute farms, which are now known as "computational grids" or "Tier 3 grids." These have different technical requirements than Tier

(continued on page 2)

For more information and news, visit www.sun.com/edu/hpc.

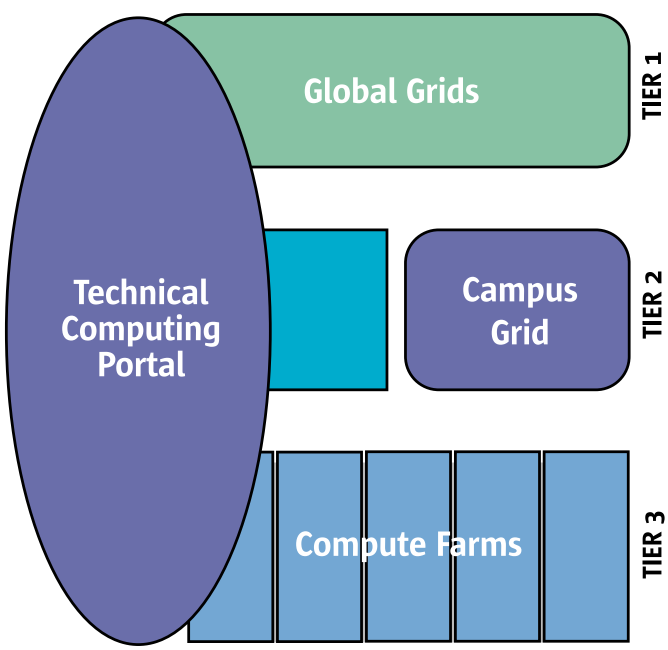

VISION FOR COMPUTATIONAL BIOLOGY

A fundamental transformation has occurred in the field of computational biology during the past few years. The field has become an explosively growing mixture of regular and medical bioinformatics, molecular modeling/computational chemistry applied to peptides/proteins, biological simulation and networks, and bioengineering. Developments like robotics and "parallelized" thinking have given rise to truly massive amounts of data, which is accumulating at fantastic rates

As a result, computers are now absolutely essential in the process ranging from data to information to knowledge and, finally, to action. But a lack of standards has hampered the smooth flow of data to action because legacy databases simply do not "talk" with each other. Also, the complexity of modern research requires that labs work collaboratively, relying on external data such as public databases like GenBank, Swiss-Prot, and PDB. These changes underscore the pressing need for open data standards and interoperability.

In addition, because the academic community is beginning to offer compute services around the clock, they are starting to recognize the importance of Sun's emphasis on "uptime" (reliability, availability and scalability).

Sun has formed a new Special Interest Group in Computational Biology (CB-SIG). The CB-SIG will provide Sun education and research customers with relevant information on Sun technology and open standards initiatives. It will also offer support for the user community and a forum to discuss recent scientific developments of interest. In addition, the CB-SIG will collaborate with Sun's Informatics Advisory Council and will promulgate the use of standards by the community. For more information on the CB-SIG, visit http://www.sun.com/edu/hpc/compbio-sig.

These are the reasons why Sun's Global Education and Research organization based its vision for computational biology on three interdependent factors: Sun's strong hardware and software platforms, strong support for open data standards, and support for the computational biology community.

Grid Computing *(continued from page 1)*

1 grids and latency, the time it takes a message to travel between nodes, is a very important consideration. Large Symmetric Multi Processing (SMP) systems, such as the Sun Fire™ series of servers, have extremely good latency between compute nodes, which allows programs to ignore latency while taking advantage of the memory address space.

For some applications either the scalability or the cost of compute cycles requires a different approach. In electrical engineering, for example, the most efficient approach is to schedule multiple simulations on a cluster of smaller, inexpensive workstations or workgroup servers. Sun has addressed this requirement by offering a pre-packaged solution, the Sun Technical Compute Farm, or "SunTCF." SunTCF comes as a pre-configured, loosely coupled, efficient cluster that is ready to run.

To make grid computing more accessible, technical computing portals are emerging as easy-to-use interfaces for job submission and job control. Instead of selecting the computing queue of a specific system on the Internet, all a user must do is submit his or her job using a technical computing portal. Sun Grid Engine will then find the appropriate resource and monitor execution.

Technical computing portals are quickly becoming popular. In fact, many institutions have plans to use a technical computing portal as the single interface for many kinds of grids. They can help resource managers, for example, who often struggle with basic questions like whether it is more efficient to run a job on campus or somewhere else. To answer such questions today the resource manager must be aware of data size, bandwidth and transfer times, transaction costs, computational execution time, latency and depth of queues. In the future, after a computing job has been submitted, the resource manager may be able to answer basic questions simply by using Sun Grid Engine to allocate resources on campus as well as in other locations. The result will be very cost efficient technical and scientific computing.

THE GRID ENABLES GEOSCIENCE TO TAKE A BROADER VIEW

One of the most exciting applications for grid computing is in the area of geoscience. At the start of the 21st century geoscience is taking on a broader view, which now includes several traditional sciences. With the help of The Grid, earth scientists, atmospheric scientists and oceanographers are now working together to determine how each of their disciplines affects the entire global change process.

Because of enhanced computational capabilities brought about by The Grid, on an international basis major funding agencies are now collaborating as well. There are now joint projects among the National Science Foundation (U.S.), the European Network for Research on Global Change, the Asia-Pacific Network for Global Change Research, and the Inter-American Institute for Global Change Research.

These agencies are collaborating to help leverage our increasing knowledge to benefit mankind in ways that were previously never thought possible. Joint projects are now improving our ability to predict catastrophes such as earthquakes and severe weather conditions. Accurate and timely prediction will minimize the effect of dangerous weather and other catastrophes on human beings and business.

Today's increased use of wireless technologies allows for Geographic Information Systems (GIS) information to enable new applications. Existing and upcoming generations of wireless handheld devices are making location services the fastest growing application based on GIS data. At the same time the expanding use of database technology is making more data than ever before available to applications developers on an anytime, anywhere basis.

"GEO Sciences Beyond 2000," a comprehensive summary of the future of geosciences published by the National Science Foundation with the help of a broad group of researchers, tells more about these exciting developments. The summary is available by visiting **www.nsf.gov.**

COMPUTATIONAL FINANCE DREAMS BECOMING REALITY

To analyze thousands of securities transactions instantly requires compute power previously unheard of. Likewise, visualizing the results of that analysis in real time used to be unthinkable. Today, however, Sun is working with major universities and other partners like Numerical Algorithms Group (NAG) and Informix Software to turn these computational finance dreams into reality.

In one exciting project, data is being loaded from actual securities transactions almost instantaneously as each transaction is completed. The data is being loaded at a rate that is three to five times faster than what was previously possible. Using sophisticated software algorithms, the data is then analyzed using compute farm technology. The results of the analysis can then be used to make intelligent business decisions on the spot. The next step in the project will be to utilize high resolution visualization to succinctly summarize the large volume of results being generated by the new techniques.

Mathematicians are using the same strategies and techniques to explore the use of the new technology for other trading platforms. For example, in the commodities markets, researchers are looking at methods to take advantage of similar simulations for the benefit of traders. Researchers will include weather data as the next step in developing models that are even more realistic and yield the most accurate data for commodities traders.

Sun will announce new developments in the field of computational finance in coming months. Future issues of *Campus Advisor* and **www.sun.com/edu** will report these developments.

GRID AND PORTAL COMPUTING: THE NET EFFECT IN SCIENTIFIC AND ENGINEERING COMPUTING

By Brian Hammond, PhD

To solve the growing number and complexity of scientific, engineering, and business problems, a new model of computing has developed.

In the distant past, data was created, stored, processed, and analyzed all at a single site on a campus. Later, with workstations becoming ubiquitous within departments, small amounts of data and computing power began to be distributed around the campus. At the time, the workstations were connected by a slow network and data was distributed either via FTP or NFS, which left each user to obtain what data they needed and compute locally on their desktop machine or send off jobs to a central computer.

That model of distributing data and computing was extremely wasteful because desktop computers are generally underutilized and because important computing and data resources were distributed all over the world. A new model was needed to harness those resources.

Evolution of "The Grid"

Today, the wide deployment of high-speed networks is putting us well on the way to solving the many problems involved in creating coordinated resource sharing in dynamic, multi-institutional confederations. This model of computing is loosely termed "grid computing" or simply "The Grid". Through emerging grid technologies we are able to create persistent ad hoc conglomerations of resources that allow distributed sources of data and computing to be assembled, solve a problem, and then be dissolved.

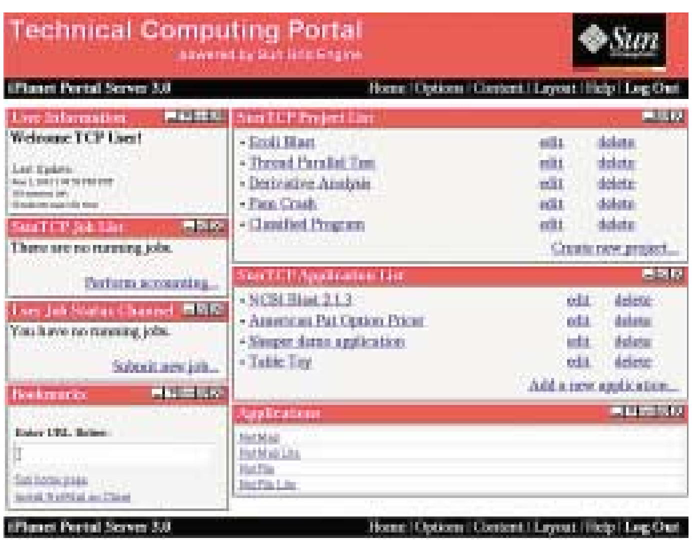

Closely aligned with the concept of grid computing is the idea of "portal computing." A user or application portal is a web-based collection of information presented together on a browser page. In many cases a portal can replace the need for users to log on to a number of different computers to gain access to the various resources that exist. A portal also allows IT managers to focus on maintaining services rather than allocating resources to user issues. In portal computing well-defined services are delivered to the portals instead of giving users direct access to the UNIX prompt. As a result, resource management and security can be exercised at a higher level.

One example of portal computing is the GridPort toolkit developed by Mary Thomas and coworkers at the San Diego Supercomputer Center and the National Partnership for Advanced Computational Infrastucture (NPACI). User portals built on top of GridPort allow users to run codes, manipulate data and files, and otherwise use components of the computational grid through a web interface. The web pages on the client side only require simple HTML and JavaScript. Therefore, any browser can be used to view the pages, while on the server the services are built using Perl and CGI scripts. Behind the scenes GridPort uses technologies such as Globus and PKI to provide secure, interactive services. **(Visit the NPACI HotPage portal at http://hotpage.npaci.edu/)**

Those who prefer to use off-the-shelf software will soon have another option: the Technical Computing Portal. This portal, which is currently under development, will integrate the iPlanet™ portal server with the Sun Grid Engine (explained on page 6). Instead of selecting the computing queue of a specific system on the Internet, all the user will have to do is submit his job using the Technical Computing Portal. The Sun Grid Engine will then find the appropriate resource and monitor execution. A prototype of the Technical Computing Portal is shown in the illustration at left.

The Grid Architecture: Global and Campus Grids

A grid may have multiple levels of hierarchy, including the *Campus Grid* and the *Global Grid*. The campus grid is a relatively tightly coupled set of resources all controlled by a central resource management system such as Sun Grid Engine. This grid may consist of scattered workstations and servers as well as centralized resources (e.g., a Sun Enterprise 10000, Sun Fire™ 6800, or Technical Compute Farm, all of which are scheduled by a single resource manager).

A global grid is a collection of campus grids, all of which have agreed-upon global usage policies and protocols, but not necessarily the same campus grid implementation. The global middleware layer must therefore implement a unique set of capabilities to solve the additional problems presented when integrating disparate campus grids into an accessible, unified compute structure. Global grid middleware leverages existing technologies, including the technologies used to implement the local grid level.

The Global Grid

Although global grids are still in their infancy, pilot projects are receiving a great deal of attention. A global grid solution consists of a common set of rules and protocols required for the sharing of data and resources in a manner acceptable to all the individual campus grids (such as defined in Globus). The global grid framework must fulfill certain basic requirements:

- Security in both the computer center and network
- Flexibility in the sharing of resources
- The ability to manage and account for a wide range of resources
- The sharing of programs, data, and computers
- Agreements on performance and cost

Building a global grid begins when all the resources are "grid ready." For the most part the grid can rely on the basic TCP/IP protocol stack (IP, TCP, UDP, DNS, etc.) for transport, routing and naming requirements; it can rely on FTP for bulk data transfer (which supports third-party transfers). In addition the user should have a single sign on (authentication) to be able to access a wide range of grid resources via one of a number of different security protocols such as Kerberos or Unix security. Finally, the resources must be able to identify themselves using a protocol such as LDAP.

Next, a layer of middleware is needed. The middleware must be able to identify and authenticate users, discover and monitor resources, provide remote data access and data movement, and act as a cross-organization resource broker and job scheduler. This is the layer that is designed to take the user's requirements and match them with available resources anywhere within the global grid using a set of resource and connectivity protocols. These network connections must have a wide range of bandwidths and latencies, operating systems, security software, and other resources, so it is important that the middleware be able to deal with all these possibilities. Specifically, this layer must provide:

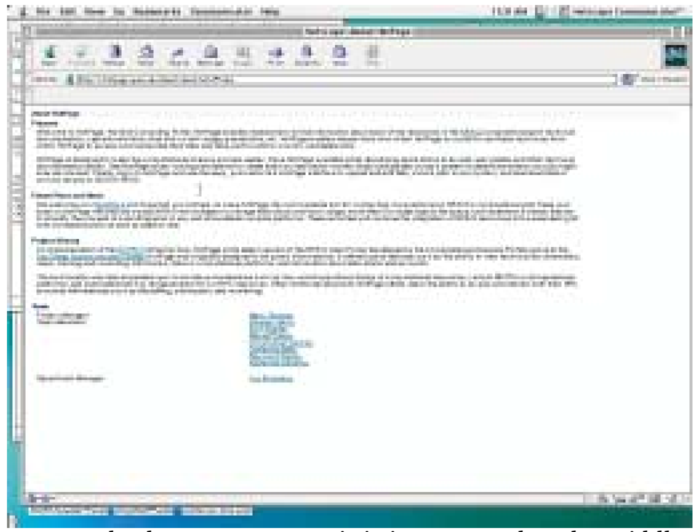

- Directory services
- Allocation, scheduling and brokering services
- Monitoring, error handling and diagnostics
- Transparent data movement
- Software discovery services
- Grid policy enforcement

This type of middleware is prompting a number of research projects. **(Visit Globus at www.globus.org, Legion at www.legion.virginia.edu, Cactus at www.cactuscode.org and Punch at http://punch.ecn.purdue.edu.)**

The third step in a global grid solution deals with a resource scheduling and management system, which must be at each individual campus grid to receive the data and control information. This should be built upon a local campus resource management system such as Sun Grid Engine, PBS, LSF, or NQE. The only requirement for the local resource manager is that they provide interfaces to the global grid scheduler to catch the request from the global grid and translate the requirements into actual job scheduling. Then they must execute the job and return any requested data.

Finally, to make this process seamless to the end user a portal interface (such as the one described above) is needed to present choices and information about all the different phases in a simple and comprehensible manner. This portal can be application specific, or can represent a conglomeration of information about work being performed throughout the grid (such as the NPACI HotPage, shown above).

(continued on page 6)

The Campus Grid and Sun Grid Engine

For most users new to grid computing, setting up a campus grid is the first step. Punch at Purdue University (shown below) is one example of the many fully operational grids that span individual campuses throughout the U.S. Constructing a campus grid is made simpler because all the resources are under the authority of a single university or company. This means that in most cases global middleware is not required. Instead, all the resources can be managed by a single resource manager such as Sun Grid Engine.

Sun Grid Engine is the product name for an advanced batch scheduler formerly known as Codine. In 2000 Sun acquired Gridware, a grid software company that featured two products, Codine and GRD (short for Global Resource Director), a system to implement and monitor enterprise-wide resource usage policies. The capabilities of GRD will be integrated into the Sun Grid Engine product line at a later date.

Sun Grid Engine is a powerful distributed resource manager using work queues on each system it controls and a centralized management and scheduling master server. While Sun Grid Engine does not currently have the ability to discover, broker and manage resources not under its control, it is able to effectively manage all resources within a campus that are explicitly under its control. It can then interface with packages such as Globus to access global resources. When Sun Grid Engine software is installed on the workstations and servers that belong to the grid pool, users submit resource requests (both interactive and batch) to the Sun Grid Engine system rather than to a particular resource. Sun Grid Engine software then distributes the requests to the most suitable workstation or server, handles error conditions, and can account for and regulate usage of individual and groups of users.

Free Software Download and Training

Sun Grid Engine is currently available and supported by Sun directly for the Solaris and Linux operating systems. This is downloadable now for free at http://www.sunsource.net.

Visit the Sun Grid Engine web site at http://www.sun.com/gridware for updates on availability, partner support, and a free Web-based training, "Administering and Supporting Sun™ Grid Engine."

Sun is supporting all other platforms for Sun Grid Engine (e.g. SGI, Cray, IBM) through third-party partners. The source code to Sun Grid Engine will soon be available at http://www.sunsource.net.

The Sun Grid Engine package has the ability to:

- Manage batch, interactive, and parallel jobs
- Manage homogeneous and heterogeneous clusters running many operating systems
- Specify job options for priority, hardware and license requirements, dependencies, and time windows
- Monitor and control pending and executing jobs
- Define and control user access to resources
- Provide precise records by job, machine, user, group, account, and for licenses provide critical usage information for decision management.

Users and administrators have access to a wide range of capabilities, including :

- Graphical user interfaces
- Defining job limits and priority levels for each user or set of users
- Configuration of queues with specific resource limits, queue owners, managers, and users
- Suspension and enabling of queues
- Secure, reliable performance

The capabilities of Sun Grid Engine are comparable to other resource scheduling systems such as LSF, PBS, and NQE.

For an excellent in-depth discussion of grid architecture, read, *The Grid: Blueprint for a New Computing Infrastructure*, edited by Ian Foster and Carl Kesselman (Morgan Kaufmann Publishers).

Brian Hammond, PhD, is the Business Development Manager responsible for High Performance Computing Environments, Chemistry and Physics in Sun's Global Education and Research Organization. Contact him at Brian.Hammond@sun.com.

THE OHIO SUPERCOMPUTER CENTER LEADS IN HPC, DISTRIBUTED COMPUTING AND SCIENTIFIC PORTALS

Sun has named the Ohio Supercomputer Center (OSC) the Sun Center of Excellence for High Performance Computing (HPC) Environments. The designation, which was announced in April 2001, is the result of OSC's leadership in HPC, distributed computing, and scientific portals.

"This partnership will help establish a new computing paradigm for the State of Ohio, using portal-based interfaces so that a researcher can access his work from anywhere, anytime, over virtually any device, based on Sun's Grid Engine technology," said Kim Jones, VP of the Global Education and Research Group at Sun.

The Sun Center of Excellence in HPC Environments is a collaborative project between OSC, The Ohio State University (OSU), the University of Cincinnati/Cincinnati Children's Hospital, and the University of Akron. The center will also allow a variety of industry and business partners from around the U.S. to experiment with distributed management and portal technologies. Business applications will be driven by the interests of the business partners, such as Nationwide Insurance, LeadScope, Exodus, LabBook, and by the expected development of a Center for Global Business Information Management and Computational Modeling at the University of Akron.

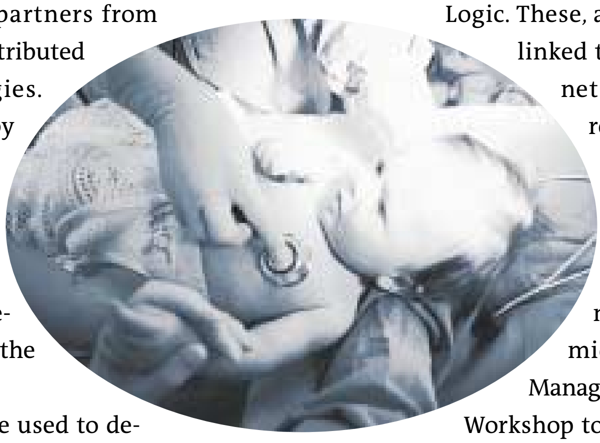

The Sun Center of Excellence will be used to develop and integrate science and business portals, focusing on life science applications. A collaborative testbed infrastructure and distributed storage from Sun will be made available to researchers working on a variety of applications. For example, Ohio researchers, including OSU's Medical Center, will be able to use the computing and storage power of the Sun Center to develop better tools to deal with disease.

Cincinnati Children's Hospital is already exploring high performance computing to deal with pediatrics and disease in their Pediatric Informatics division and is looking forward to using this expanded capacity to explore new research questions.

"Being able to say that OSC is a Sun Center of Excellence builds national recognition for us and our collaborators," said Al Stutz, Acting Executive Director of OSC. "This provides us the opportunity to reach new constituencies in order to accelerate the use of information technologies and to strengthen Ohio's attractiveness and global competitiveness."

The primary environment will include a cluster of four Sun Fire™ 6800 midframe servers with a total of 72 UltraSPARC™ III processors with a peak floating point speed of over 100 gigaFLOPS, 96 gigabytes of memory and 6.5 terabytes of disk storage. The cluster will be located and centrally managed at OSC to allow academic research collaborators to utilize the system and participate in research efforts. A distributed center will be located at the University of Cincinnati/Children's Hospital, which has an existing Sun HPC system. Two Sun Fire™ 6800 servers and a Technical Compute Farm with a capacity of 86 gigaFLOPS, 64 gigabytes of memory, and one terrabyte of disk will be added. To support the work in bioinformatics, both sites will also feature DeCypher bioinformatics accelerator boards from TimeLogic. These, as well as future installations, will be linked throughout the state over OSC's Internet2 "OARnet" network, and will allow researchers and industrial partners to experiment with distributed computing capabilities and high performance computing applications.

The distributed computing environment will be based on Sun's Grid Engine middleware, iPlanet™ Portal Server, Sun Management Center, HPC Cluster Tools, Forte™ Workshop tools, and storage technology from the newly acquired LSC, Inc. OSC will deploy LSC's SAM-QFS hierarchical storage management software and high performance file system to provide production and experimental storage access. Storage Area Network (SAN) and Networked Attached Storage (NAS) will be evaluated in LAN, WAN, and various distributed Sun computational domains. OSC will also test the capabilities of the distributed storage management products from LSC by accessing them from current OSC systems, such as the Cray SV1, the SGI Origin 2000, and the current and future LINUX cluster systems.

One of the center's priorities will be to jointly develop and deploy a scaleable, portable, and extensible technical computing portal, drawing upon projects from both OSC and Sun. This portal will integrate browser-based portal software with the Sun Grid Engine product. Development will draw on OSC's extensive knowledge of creating portals such as the SciPortal project.

GRID SUCCESS DEPENDS ON INFRASTRUCTURE

By John C. Fowler

The Grid offers many benefits and a solid infrastructure is necessary for maximum advantage. The core elements of that infrastructure are security, access, transport, authentication and authorization. When these elements are closely integrated with unified directory services (such as LDAP) in an open standards environment, it will result in a tight, secure framework for building a campus grid.

Security (Authentication and Authorization)

The extent to which compute resources can be used for more than scientific computing depends on how tight the security model is. It is important to secure the data at the transport, authorization and access control layers. The protocols involved are LDAP, NFSv3, IPsec, and SCP.

Kerberos is best for authenticating all user names from a central point using a secure transport so that no username or password is ever transported across the network. iPlanet Directory Services is an LDAP implementation that can use Kerberos as the authentication mechanism. For more information visit http://www.internet2.edu or http://middleware.internet2.edu/ and click on the DoDHE, eduPerson, LDAP recipe links.

Standardizing on NFSv3 will reduce the need for transferring data sets to local temporary storage because it will support the secure mounting of data from a common file service. It also reduces the demand for special code to insure that the data is available and accurate. If transport security is still a concern, consider activating IPsec protocol, which allows secure paths to be defined between remote machines and the data storage environments within the data center. If a batch transfer is necessary, use a product like Secure Copy (SCP) within job streams.

Data Access

Within the data center framework is a storage access layer where storage resources available to all the applications can be consolidated within the center. Once authenticated, all applications can access this layer via NFSv3 or SQL calls according to the stated security rules. Gridware nodes can then make standard calls to mount and use the data resources without the use of special procedures or platform specific codes.

Job Control

To make the Gridware product work best, keep the management interface simple by doing the following:

1. Since the Gridware software uses and leverages the common Solaris (UNIX Scripting) or shell commands, consider embedding commands to query the LDAP database for information about users, groups, or programs and resources necessary to support the function or program.
2. Mount and verify access to all the data and programs that will be run.
3. Verify access to all output areas and logs, being sure to take advantage of the accounting features that are built in.
4. Set Check Points within the Job Stream to insure a restart capability.
5. Execute the program.
6. Set a Check Point after each program execution. This will help identify each step in the run time process.

John C. Fowler is Infrastructure Market Segment Manager for Global Education & Research LOB at Sun. He can be reached at John.C.Fowler@sun.com.

Comments/suggestions? Send them to campusadvisor@sun.com

WORLDWIDE HEADQUARTERS SUN MICROSYSTEMS, INC., 901 SAN ANTONIO ROAD, PALO ALTO, CA 94303-4900 USA.
Phone: +1-800-555-9SUN OR +1-650-960-1300 · Internet: www.sun.com

Sun microsystems
We're the dot in .com™

SUN WORLDWIDE SALES OFFICES

Africa (North, West and Central): +9714-3366333 • Argentina: +5411-4317-5600 • Australia: +61-2-9844-5000 • Austria: +43-1-60563-0 • Belgium: +32-2-704-8000 • Brazil: +55-11-5187-2100 • Canada: +905-477-6745 • Chile: +56-2-3724500 • Colombia: +571-629-2323 • Commonwealth of Independent States: +7-502-935-8411 • Czech Republic: +420-2-3300-9311 • Denmark: +45 4556 5000 • Egypt +202-570-9442 • Estonia: +372-6-308-900 • Finland: +358-9-525-561 • France: +33-01-30-67-50-00 • Germany: +49-89-46008-0 • Greece: +30-1-618-8111 • Hungary: +36-1-202-4415 • Iceland: +354-563-3010 • India: +91-80-5599595 • Ireland: +353-1-8055-666 • Israel: +972-9-9710500 • Italy: +39-039-60551 • Japan: +81-3-5717-5000 • Kazakhstan: +7-3272-466774 • Korea: +822-3469-0114 • Latvia: +371-750-3700 • Lithuania: +370-729-8468 • Luxembourg: +352-49 11 33 1 • Malaysia: +603-264-9988 • Mexico: +52-5-258-6100 • The Netherlands: +00-31-33-45-15-000 • New Zealand: +64-4-499-2344 • Norway: +47 23 36 96 00 • People's Republic of China: Beijing: +86-10-6803-5588 · Chengdu: +86-28-619-9333 · Guangzhou: +86-20-8755-5900 · Shanghai: +86-21-6466-1228 · Hong Kong: +852-2202-6688 • Poland: +48-22-8747800 • Portugal: +351-21-4134000 • Russia: +7-502-935-8411 • Singapore: +65-438-1888 • Slovak Republic: +421-7-4342 94 85 • South Africa: +2711-805-4305 • Spain: +34-91-596-9900 • Sweden: +46-8-631-10-00 • Switzerland: German: 41-1-908-90-00 · French: 41-22-999-0444 • Taiwan: +886-2-2514-0567 • Thailand: +662-636-1555 • Turkey: +90-212-335-22-00 • United Arab Emirates: +9714-3366333 • United Kingdom: +44-1-276-20444 • United States: +1-800-555-9SUN OR +1-650-960-1300 • Venezuela: +58-2-905-3800

Printed in USA 6/01 KE334-0 / 2K