White Paper

SGI and Intel on the Grid
Unique Capabilities for Grid Computing

Table of Contents

Introduction

What is Grid Computing?
A "Grid" is a model for organizing networked computing to help achieve flexible, streamlined, protected, shared, user–controlled access to resources. A dispersed user community accesses virtualized Grid resources without regard for their physical location. A key goal of a Grid is to automate resource access and management to avoid the need for a big technical staff to setup, route and monitor jobs. Grid computing has evolved to allow scientists and engineers to share and gain quick and transparent access to critical or powerful computing resources and advanced instrumentation. A technical Grid user directly accesses computers, software, data, and other critical resources such as visualization with little concern for their physical location. Such Grids help enhance the ability to solve more difficult problems, enable greater collaboration across organizations, and increase the utilization of and access to expensive, specialized resources.

For example, the COSMOS Project at Cambridge University, under the direction of Professor Stephen Hawking, is studying the formation of galaxies after the Big Bang, using a dedicated cosmology Grid. Grid users have access to a variety of computing resources located at universities across the UK without regard for their location. Specialized software helps them identify and access the best resource for a particular job. Similar grass-roots small Grids are being organized all over the world.

This first wave of Grid computing users encompasses national, government, and academic HPC arenas, driven primarily by an insatiable appetite for compute power and the need for large–scale, geographically independent collaboration. Participants in this wave are the true pioneers of Grids—thought leaders and early adopters of the latest technologies. They have little dependence on commercial Grid software, meaning that much of the development is either open source or done in–house. A few examples include the CosmosGrid, Welsh Visualization Grid, WestGrid, the UK e–Science initiative, and the European DataGrid. Some of these projects have pioneered the transfer of technology for industrial Grid development and become foundational for the use of Grid in other sectors.

The second wave of Grid adopters, gaining momentum today, is made up of primarily commercial enterprises, especially in the financial services, pharmaceutical, and manufacturing segments. The participants in this wave see Grids as essential for remaining competitive, and use Grids to increase the efficiency of their infrastructure and business processes. They introduce new use models to perform mostly computationally–intensive tasks such as financial risk analysis, drug discovery, car crash simulations, while integrating these capabilities with existing product data management systems. The ultimate objective is to go straight from design to production in a manufacturing environment, while enabling effective cooperation with suppliers and contractors over the Grid. Second–wave Grid users have a greater dependency on

Fig.1. Grid computing creates an infrastructure of virtualized, shared resources that can be accessed

commercial Grid software than first–wave users, and they may require either commonly accepted standards and multiple software vendors or open source to help protect their investment. They often employ application developers to help fill the gaps in the commercial Grid offerings available now or adapt them to the use scenario.

While distributed computation—using multiple computer systems to solve a single large problem—is an important aspect of first and second wave Grid computing, distributed access to critical data (such as product data, test results, simulations, etc.) and unique capabilities are becoming ever more important. Particularly in Europe, a great deal of effort has gone into creating Grids for sharing huge data volumes and for resource sharing. Connection to a Grid helps increase system utilization and return on investment for advanced computing resources.

The third market segment wave views Grid as a critical infrastructure component that can span from scientific and industrial pursuits to encompass an entire organization. This wave helps address the sharing and management of many of an organization's compute, information, and communication assets: e.g. handheld devices, laptops, appliances, applications, databases, visual displays, edge servers, mid–tier servers, and back–end servers etc. Many companies are exploring the deployment of Intra–,Extra–, and Inter–Grids to cost–effectively meet their computational needs. Intra–Grids—sometimes referred to as Enterprise Grids—are designed to help meet computational needs within an organization. Extra–Grids, or collaborative Grids, allow collaborators and suppliers to enhance collaboration, while Inter–Grids might span several commercial organizations to facilitate broader sharing of resources. Companies typically follow a progression from Intra–Grids to Inter–Grids for business to business application, and only then consider Grids to accommodate customers and end–users.

Grid Processing Strategies

The following are four unique perspectives to consider when discussing strategies for Grid deployment:

1) End–users focus on ease of use and access to greater computational power.
2) IT professionals focus on ease of management, speed of deployment and the ability to economically deliver services to more end–users.
3) Executives focus on ROI, operational efficiency, and adding new value for their users and customers.
4) Vendors focus on growing the market segment and enhancing ease and speed of solution deployment.

Any Grid deployment must prioritize and take into account these different perspectives to succeed. With data volumes continuing to grow rapidly in virtually all market segments, mechanisms for efficient data storage, access, and management are essential to a successful Grid. End–users require high speed data access to take full advantage of Grid resources. IT professionals need effective, flexible, and economical ways to manage large volumes of data on which Grid users depend, while executives are concerned with the overall economy and benefits of Grid storage systems as well as data security.

In some applications, a Grid that enables access to previously inaccessible information provides huge benefits, regardless of the performance delivered. However, in the majority of Grid uses a data infrastructure is required that can deliver data with high bandwidth and minimum latency to feed computing and visualization resources.

Most organizations have a variety of computing resources deployed for different purposes. Often, similar resources are duplicated in different groups because they cannot easily be shared. By facilitating wider access, a Grid can help decrease this duplication and make the benefits of expensive resources available to a wider base of users.

The processing problems tackled by Grid computing can be broken into two classes:

- **Capacity computing**, in which a large number of small jobs are run simultaneously. (Large problems that are easily partitioned into many smaller jobs also fall into this category.) Each job requires a small number of processors or even a single processor. Peak loads may require use of a large number of individual systems. The self contained nature of capacity jobs typically requires computing ability but places fewer demands on internal system bandwidth, system I/O resources, and Grid data bandwidth.
- **Capability computing**, in which a huge number of processors, large shared memory, and multiple I/O channels are required to tackle the most difficult computing problems. Capability problems often stretch the limits of the processor interconnect as much as they do the compute ability and require exceptional bandwidth to Grid data storage to help avoid bottlenecks.

A Grid can make a diversity of resources available to tackle both classes of problems. Grids can by nature be very heterogeneous, offering a range of resources from multiple vendors such as large commodity clusters, large single system image (SSI) machines, smaller servers, and advanced visualization systems.

First and second wave Grid adopters typically have requirements for both classes of processing to tackle scientific and technical problems, making data bandwidth an essential element. Third wave adopters may focus resources on capacity computing, deploying a large number of relatively economical compute nodes and centralized network storage that can be flexibly re–deployed to help meet rapidly changing computing needs while simplifying the overall infrastructure and helping to reduce capital and management costs. Third wave Grids often benefit from including capability computing resources within their corporate Grids as well. When not in use for scientific and engineering tasks, these resources are utilized for large scale financial analyses, to power graphics pipes for advanced visualization, and other purposes.

Enabling Middleware for Grid Computing

Because Grids emphasize distributed resource sharing—often crossing organizational boundaries—issues such as scheduling, accounting, access control, and security are critical. To realize the true vision of Grids, robust and ubiquitous standards must be in place, and organizations need access to a wide range of industry–standard building blocks and solutions that can be mixed and matched as needed.

Standards are critical to enabling volume economics, and for moving technologies from early adopters to widespread acceptance. The move to more distributed Grid–like environments has been enabled by Java, J3E, and XML, among others. The evolution of Web Services has had its own track, including standards such as Simple Object Access Protocol (SOAP) and Web Services Description Language (WSDL), but is now intersecting with Grid–specific standards such as Open Grid Services Architecture (OGSA) via the Web Services Resource Framework family of OASIS standards.

The Global Grid Forum (GGF) is an established body to discuss novel Grid concepts, foster the uptake of Grids and define Grid community standards. It has evolved from a largely academic organization into a clearinghouse for scientific and commercial Grid users and developers of building blocks and solutions alike. The GGF has defined principles of Grid architecture in the OGSA (Open Grid Services Architecture), and it has been instrumental in bringing about the convergence of Web Services and Grids as expressed by the Web Services Resource Framework (WSRF) family of OASIS standards. It continues to provide leadership to the Grid community, and maintains close relationships with other Grid special interest groups (like EGA) and standard bodies (like OASIS and W3C).

To help address Grid standards specifically at the enterprise level, the Enterprise Grid Alliance (EGA) was recently formed to develop Grid solutions standards in servers, databases, storage, and other enterprise–level applications. The EGA, a consortium of leading technology vendors, aims to help move Grids from their historical strengths in scientific and technical market segments to enterprises in the full range of market segments.

Among other standards–related groups important to the development of Grids is the Globus Alliance, which develops the Globus Toolkit, a reference implementation of OGSA and WSRF that supplies the building blocks (software and tools) for many current Grid deployments.

In addition, the open source project UNICORE (Uniform Interface to Computing Resources) provides middleware and client software that combines the resources of supercomputer centers with science and engineering Grids. UNICORE enables access to many different systems while hiding platform differences from the user, creating a seamless portal for accessing clusters and HPC systems, compiling and running applications, and transferring input/output data. Strong authentication is performed at all times in a consistent and transparent manner. Funded by German and European research projects since 1997, UNICORE is currently used at several European HPC centers as well as in major Grid projects in Japan and China.

The OGSA architecture specification and the WSRF interface specifications allow users to implement middleware services based on the earlier Grid projects in a uniform and interoperable way. Globus Toolkit version 4 is one of these new implementations. While standards and special interest groups are moving quickly to address Grid middleware standards and specifications, other challenges remain.

SGI, in cooperation with Intel, is working hard to address many of these challenges to help:

- Increase system and application scalability. The largest problems consistently push the limits of available technologies. Different problems may require different programming models for the best solution
- Enable high–speed data sharing between heterogeneous systems
- Improve Grid security
- Allow remote visualization with compressed output from high–value, real–time visualization systems delivered to any display device on the grid

SGI on the Grid

As a recognized leader in high–performance computing and advanced visualization, SGI has been involved in Grid computing from the earliest days and has a long history of achievement in Grid computing:

- The first public demonstration of Grid technology happened at the Supercomputing ’97 conference and was powered exclusively by SGI® systems. The demonstration was led by Argonne and USC
- Initial development of the Globus Toolkit was done on SGI systems using the SGI® IRIX® operating system and associated development tools
- SGI systems are used in almost all of the major technical Grid installations in Europe, North America, Japan, and Australia
- SGI is a Platinum sponsor of the Global Grid Forum and is actively involved in its Working Groups and Research Groups

SGI faces and solves many of the important challenges of Grid computing every day—including scheduling, bandwidth, data distribution, and coherency. SGI computer systems, visualization, storage systems, networking technologies, and software provide unique capabilities for Grid computing:

- With the SGI Altix family of servers and superclusters, SGI combines the benefits of the Linux operating system and Intel Itanium® 2 processors with global shared memory and high performance clustering, yielding a platform ideally suited for a wide variety of capability and capacity problems.
- SGI’s InfiniteStorage Shared File System CXFS provides high–performance, shared access to data over storage area networks (SANs).This data can be accessed by all major operating systems(32–bit Linux®, 64–bit Linux®, SGI IRIX®, Sun™ Solaris™, IBM® AIX®, Windows®, Mac OS® X, and other versions of Unix®.)
- In partnership with Yotta Yotta, SGI offers a solution that extends high performance file sharing to the wide area. Data rates of 700MB/sec to shared storage over a distance of 27,000 kilometers have been demonstrated.
- SGI is pursuing a variety of security certifications for Novell/SUSE Linux running on the Altix platform.
- With Visual Area Networking (VAN) technology, SGI brings the power of collaborative, remote visualization to the Grid. Grid users can view the output of advanced visualization systems anywhere on the Grid. Multiple Grid users can collaborate by viewing and interacting with the same image at the same time in different locations.

This paper discusses these technologies and their applicability to Grid computing.

The SGI Altix Family

Grid computing installations rely increasingly on industry–standard computers running Linux to tackle both capability and capacity problems. Commercial off–the–shelf (COTS) based systems combined with Open Source Linux applications and tools create a flexible and cost–effective Grid computing infrastructure.

The SGI Altix family extends the benefits of COTS and Linux through the addition of unique SGI technology that allows individual Altix systems to scale beyond the capabilities of other Linux systems. SGI utilizes Intel Itanium 2 processors for its Altix systems. The use of Intel technology allows SGI to utilize the breadth and depth of the Intel R&D effort which includes not only a the Intel Itanium 2 processor line, but also Intel's software tools. These include compilers, libraries, tools for application organization and tuning on a single node, tools for developing applications across HPC clusters, and is expected to include tools for distributed computing environments.

The use of SGI's patented NUMAflex architecture and NUMAlink interconnect make it possible to scale the number of processors in Altix systems well beyond the level that has been possible in other shared–memory designs. Instead of the traditional backplane design, SGI uses crossbar switches and high–speed cabling, allowing each node direct access to the memory in other nodes with a relatively slight increase in latency versus accesses to local memory.

The success of any computer system depends largely on the quality of its system and application software. Within the open–source community, SGI is committed to provide ongoing development and improvement of the performance and reliability of the Linux operating system. With its focus on high–productivity computing environments, SGI has already contributed to Linux scalability, scheduling, memory usage, I/O, and other efforts critical to high–demand application performance. The result is the world's first Linux operating system that supports global shared memory across cluster nodes, and an advanced development environment optimized for the technical applications common on the Grid. By supporting all the commonly available tools for Linux, the Altix family ensures immediate productivity, creating a "Grid powerhouse" for the open source community

The SGI Altix 3700: Supercomputing Capability for Linux

The SGI Altix 3700 supercluster supports individual cluster nodes with up to 256 processors and 3TB of memory in a single system image (SSI). Each node is itself a formidable computing system capable of tackling a wide variety of capability problems. Multiple Altix 3700 nodes can be joined into superclusters with global shared memory across nodes and thousands of processors. Altix 3700 superclusters can be configured to tackle the most difficult capability problems or to provide the capacity necessary to handle vast numbers of smaller jobs. The Altix 3700 is ideal as a central shared compute platform for Grid deployments that accommodate a high variety of applications and job sizes with many users.

The SGI® Altix 350 Mid–Range Technical Server for Capacity Computing

SGI Altix 350 offers the same capabilities as the Altix 3700 in a smaller package at a breakthrough price point. Scaling from 1 to 16 processors and 192GB of global shared memory, the Altix 350 is an ideal platform for tackling capacity problems or smaller capability problems. The Altix 350 can be clustered using standard cluster interconnects to create clusters in which the individual nodes offer substantially greater computing power than typical Linux nodes, allowing each node to handle more individual jobs or more complicated jobs. Reduced node count ultimately simplifies cluster management.

Fig 2. SGI Altix Family

Powering Heterogeneous SANs

For large capability problems, data I/O is just as important as computing capabilities. SGI offers storage area network (SAN) solutions and industry–leading software to meet the most demanding storage needs. SGI has more experience deploying high–performance SAN technology than any other system vendor. SGI was the first system vendor to ship Fibre Channel storage and the first to ship a complete SAN fabric. SGI also developed the InfiniteStorage Shared File System CXFS, the first shared filesystem for the SAN, and was the first system vendor to begin shipping 2Gb–per–second Fibre Channel technology. Using SAN technology, SGI has demonstrated I/O throughput in excess of 7GB per second to a single system.

In addition to enabling shared access to computing resources, Grids are, in many cases, intended to provide shared access to important data. Storage solutions from SGI can help make that goal a reality.

The SGI InfiniteStorage Shared Filesystem CXFS: High–Performance Shared Data Access

Shared, fast, and reliable data access is a critical component of successful Grid computing. Without it, data either has to be copied over the Grid to local storage each time it is needed—creating potential issues with data security and integrity—or slow network filesharing technologies like NFS must be used. In either case, valuable time is wasted waiting for data.

The SGI InfiniteStorage Shared File System CXFS provides high–performance shared data access to the same files and filesystems for any system on a SAN. Since a SAN can span tens of kilometers, CXFS is an ideal complement to Grid computing when used in local Grids such as campus environments, and SGI is actively working to extend the technology to better meet the needs of the Grid. (See the following section.)

CXFS provides the ability to make data accessible where it's needed for capability problems without copying. Capacity problems that require shared access to the same data set also benefit immediately from the high–performance data sharing provided by CXFS. For example, Cardiff University installed a SAN with 2TB of RAID storage and CXFS for its Grid environment. This storage system provides direct high–speed data access and data sharing for its 32–processor capability system and four 8–processor capacity systems.

SGI designed CXFS for applications where shared data access is critical and local area networks (LANs) simply cannot provide the necessary bandwidth. CXFS gives all SAN–connected systems simultaneous high–speed access to the same filesystems and files. A single system can have multiple connections, making it possible to achieve data rates of multiple gigabytes per second.

One system on the SAN acts as a metadata server, controlling file permissions and mediating shared access. Unlike network file sharing, where all data goes through the file server (which often becomes a bottleneck), once the metadata server grants access, systems with CXFS read and write data directly over the SAN to and from disk.

Should a metadata server fail, a designated backup metadata server automatically takes over management of the CXFS filesystem. This feature—in combination with fully redundant SAN configurations and RAID storage—delivers extremely high availability along with exceptional performance. Even if failures occur, CXFS ensures that a path to access data is always available.

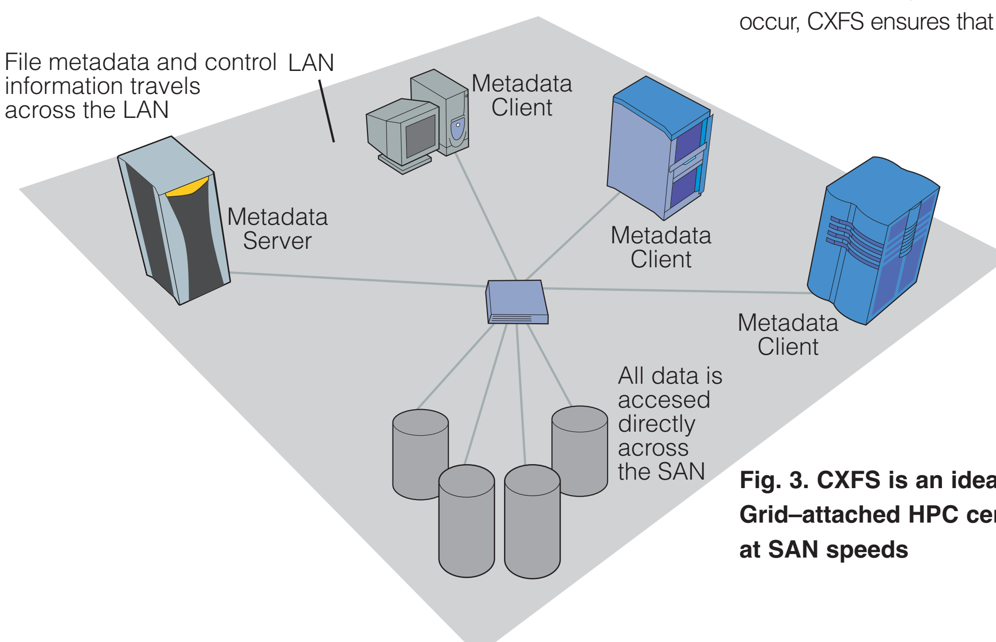

Fig. 3. CXFS is an ideal storage solution for intra–Grids or Grid–attached HPC centers, since it enables data sharing at SAN speeds

Extending CXFS to Meet the Needs of the Grid

CXFS is ideal for data sharing and management needs in localized areas. SGI has joined efforts with Yotta Yotta to bring the capabilities of CXFS to wider geographic areas and widely distributed Grids. With the ability to intelligently share data across thousands of miles, the SGI and YottaYotta solution enables Grid users to share data without the need to waste time and valuable storage capacity making local copies of large data sets. The combination of the YottaYotta® NetStorager® System and CXFS provide shared data access across great distances at near local performance rates. SGI and YottaYotta have demonstrated a CXFS cluster reading and writing to a shared file across 2,900 miles at approximately 700 megabytes per second.

SGI InfiniteStorage Data Migration Facility (DMF)

Existing Grids are already managing huge quantities of data. Since Grids maximize the utilization of computing resources, their potential to generate new data and consume storage is very high, making storage capacity and data lifecycle management (DLM) critical issues. By targeting data to appropriate storage media (primary disk storage, secondary serial ATA storage, tape, etc.) DLM solutions can significantly reduce the cost of Grid storage infrastructures.

The SGI® InfiniteStorage Data Migration Facility (DMF) is a hierarchical storage management (HSM) system that is in use as an adjunct to online storage for a large number of existing Grids to provide automated data lifecycle management. With DMF, a storage pool of nearly unlimited capacity can be created, making it an ideal solution for storage–hungry Grids. On a Grid, multiple systems can use DMF to transparently migrate unused data to secondary storage or tape. That data can be recalled for access by local systems or for transfer to other Grid–attached locations.

DMF automatically migrates data from primary storage to secondary disk or tape–based storage according to user–defined criteria. Files are automatically recalled to primary storage as they are accessed without user or system administrator intervention. DMF allows access to a nearly unlimited pool of data without concern for the medium on which that data is stored. By comparison, manual data archiving solutions eat up valuable time determining which data to archive and then copying that data to tape. Restoring data from tape for later use is similarly time consuming. DMF frees users to focus on real work by eliminating the need for manual data management.

The COSMOS Project at Cambridge University uses 1.6TB of SAN–based RAID storage with CXFS for shared data access. A 4TB tape library on the back end is managed by DMF, creating a large virtual storage pool that significantly reduces data management tasks, allowing cosmologists to spend more time thinking about the universe and less time worrying about data storage.

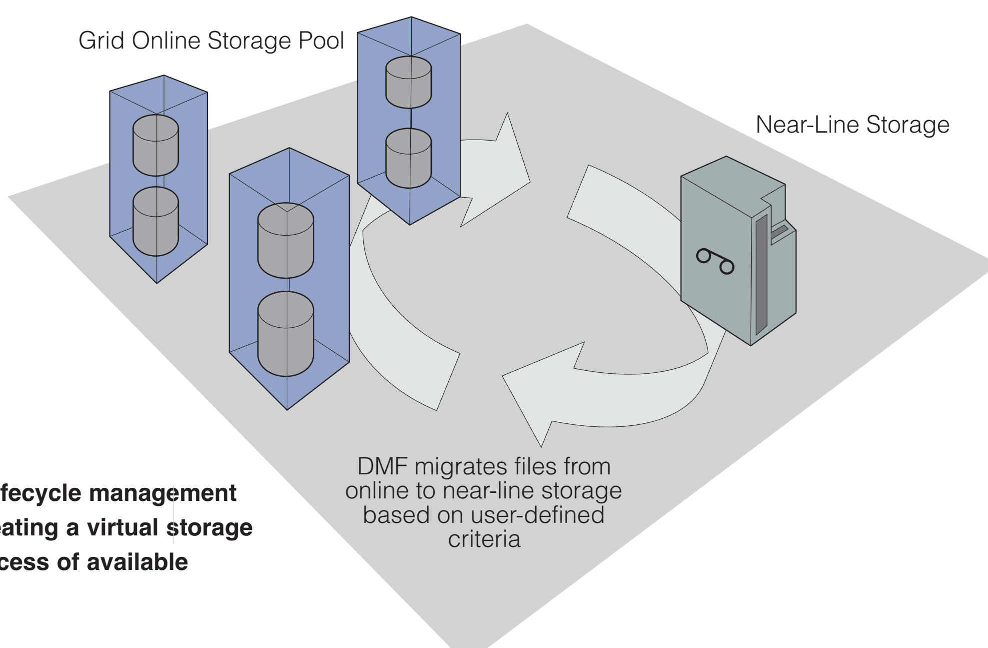

Fig. 4. DMF provides data lifecycle management for Grid–based storage, creating a virtual storage pool with capacity far in excess of available online storage

Security for the Grid

Because Grid computing may involve the sharing of critical data between systems in different organizations, security is extremely important. As Grid technology becomes more widely adopted by private industry, the need for security will only increase. The security of a Grid is only as good as its weakest link. Secure middleware does little good if the systems on the Grid are themselves insecure. Therefore, operating system security is a key element in overall Grid security.

SGI has an impressive history of trusted systems expertise based on its long involvement with the federal marketplace segment, and is now bringing that expertise to the Open Source Linux community. Security for SGI Altix systems is provided by the underlying Linux operating system. Users can choose either the SGI Advanced Linux Environment, based on Red Hat Enterprise Linux (RHEL), or Novell/SUSE SUSE Linux Enterprise System (SLES).

For network security, the Linux community has incorporated a variety of industry–standard network security interfaces into the Linux kernel. For user and application security, the Linux community has also incorporated industry–standard interfaces. In addition, there are third–party solutions for software licensing and system management.

Trusted Environment for SGI Altix

SGI is pursuing Common Criteria Security Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP) evaluation with Novell/SUSE SLES. Novell/SUSE SLES supports the basic components required for CAPP certification including access control lists (ACLs), capabilities, and auditing via Linux Audit System (LauS).

For the SGI Advanced Linux Environment, SGI supports SNARE for auditing. SNARE is Open Source software from Intersect Alliance (www.intersectalliance.com). SNARE has been ported and extended to support multiprocessor systems and runs on the SGI Advanced Linux Environment. SNARE for Altix is available today to all SGI Altix customers.

Multi–Level Security (MLS) is defined as the ability to compartmentalize user interactions according to specific security labels as defined by the site administrator. For example, a two security level environment can be defined, such as unclassified and classified or many levels of security can be defined. It may be possible to implement Multi–level Security environments today using a guard box configuration. An SGI Origin 350 system running Trusted IRIX or another MLS–certified system can serve as the guard box and manage access to other Grid systems based on their individual classification levels. SGI will have SGI Altix systems capable of supporting MLS once the Common Criteria Security LSPP certification is complete.

Visualization Environments

Visualization has become a critical element in Grid computing as faster supercomputers and higher resolution sensing devices drive an enormous increase in data. Much of this data can only be understood and meaningful decisions reached in a visual context.

For more than 20 years, Silicon Graphics has been creating the world's most powerful visualization systems. These systems have transformed the nature of technical computing and enabled new fields such as 3D CAD, molecular modeling, 3D seismic analysis and the volumetric analysis of 3D medical scans.

However, with the current onslaught of data, faster and faster graphics processors are not enough to keep up. Silicon Graphics is developing a new generation of visualization solutions that have the fastest graphics processors in the world but are also designed to speed the filtering and analysis of data and enable sharing of visual data between colleagues in the same room or at different sites across a Grid.

Silicon Graphics Prism™
The Silicon Graphics Prism family of visualization systems builds on the scalable compute, memory and I/O capability of the SGI Altix family of HPC systems through the addition of scalable visualization capabilities. These systems bring together the economics of industry–standard components with SGI's history of scalable visualization to deliver a quantum leap in visualization performance and price–performance.

Silicon Graphics Prism is designed to enable users to break through desktop computing, memory and visualization limits that separate them from revolutionary results. Modular design allows an organization to deploy a small Grid visualization resource and grow it as their needs and budgets allow.

The unique ability to have tons of graphics processors working in tandem with hundreds of CPUs to analyze terabytes of memory resident data enables researchers to approach problems in completely new ways. The importance of pre– and post–processing of results is reduced as scientists and engineers become an integral part of a real–time simulation and analysis process. The Grid plays a crucial role in a reformulated research endeavor since high–end or even mid–range visualization resources should be available to everyone within an organization and not just a select few with keys to the computer room.

Fig 5: The Silicon Graphics Prism family of visualization systems

Power Systems	Team Systems	Extreme Systems
Designed for single users or small groups with 2-4 graphics pipelines, 4–8 Intel Itanium 2 processors and 96 Gbytes of memory	Designed for teams or small departments with 4–8 graphics pipelines, 8–16 Itanium 2 processors and 192 Gbytes of memory	Designed for departments or enterprise computing centers with 4–16 graphics pipes 16–512 Itanium 2 processors and up to 6Tbytes of memory.

Visual Area Networking

Virtualization is one of the great powers of Grid computing and has historically been one of its great weaknesses. Grid environments empower users to use resources that are located anywhere on Earth, but if they use those resources to generate large amounts of data there has been no easy method for them to analyze and visualize the results.

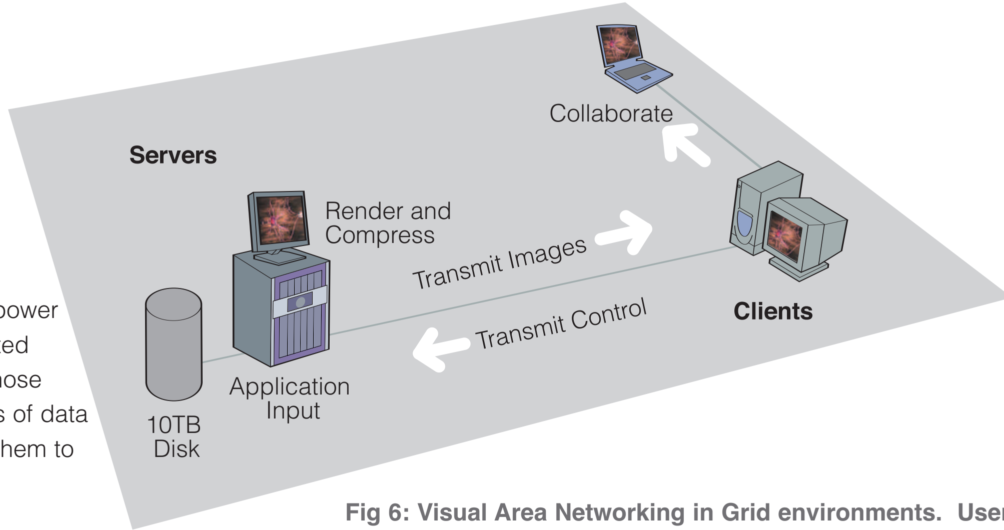

Fig 6: Visual Area Networking in Grid environments. Users anywhere on a Grid can view advanced visualization output acting as a Grid visualization server. No data is required at the user's location.

Visual Area Networking (VAN) helps close the loop between users and data, by enabling them to visualize results of an analysis done at a different location without transferring the data to their local desktop. This provides substantial benefits:

- Users increase productivity by eliminating time consuming data copies from remote data centers. Compressed representations of interactively–generated displays are transmitted across the Grid to the end–user. For large data sets, eliminating data copies may be the difference between finding a cure for a debilitating disease and deleting useful results.
- Users can interactively analyze and visualize 100 gigabyte to terabyte data sets using the full compute and visualization power of a Silicon Graphics Prism family system.
- Teams can collaborate over a distance, with each user seeing the same visual results sharing in the control of the visualization application
- First and second wave organizations can instantly focus their resources on a single problem and deliver tremendous amounts of analysis and visualization power to any group that needs it—increasing the likelihood of breakthrough results
- Third wave enterprises can optimize resource utilization by virtualizing access to visualization resources, deploying only what is needed and then growing the resource as aggregate demand requires

All of this is accomplished using off–the–shelf visualization applications without any recoding.

Silicon Graphics primary VAN solution is OpenGL Vizserver. OpenGL Vizserver is designed as a client–server environment that, in concept, resembles web servers and browsers. All of the heavy computing, visualization and data access is done on the server while a thin client displays results and tells the server what to do next. OpenGL Vizserver has been designed from the start to work in Grid environments, supports standard Internet protocols, and has APIs that enable integration with AccessGrid, today's Globus toolkits, or tomorrows WSRF standards.

For more information on Visual Area Networking and Grid computing see *http://www.sgi.com/products/visualization/van/*.

Conclusion

Grid computing is the next step in the evolution of networked computing. To quote Ian Foster, "globally virtualized Grid computing is a natural extension of today's Internet1." Grid infrastructures are already well established in many scientific and engineering environments, and are becoming increasingly popular in industries such as manufacturing, finance, and pharmaceuticals. Grids are primed for widespread deployment across the enterprise and between enterprises. SGI has been involved with networked computing and Grid computing from the beginning and is already powering many of today's largest Grids. SGI offers advanced technologies for computing, data management, security and advanced visualization—technologies that offer unique capabilities and unique advantages for Grid computing.

[1] Ian Foster, THE GRID: Computing without Bounds, Scientific American, April 2003

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.925.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3723 (11/2004) J14741