THE CLIPPER GROUP
Navigator™

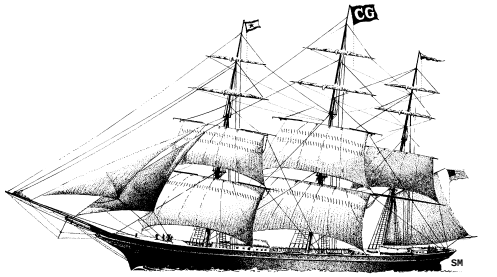

Navigating Information Technology Horizons℠

Published Since 1993 September 24, 2002

SGI Grids Accelerate Business Process

Analyst: Anne MacFarland

Management Summary

Providing computing for scientific, technical and creative workers has particular pains. Hardware is expensive, and buying more-of-cheaper-products is not an option. When the files are too bulky to quickly replicate, file sharing options shrink. Many of the "divide and conquer" solution approaches of the past few years just don't apply. Meanwhile, **accelerating workflow is incredibly important to these businesses, because the workflow is the business.** Together, these factors require a new solution. Silicon Graphics, Inc. (SGI) proposes to solve them with grids.

These grids are very different from the grids of PCs that search for extraterrestrial life. SGI focuses on the use of graphics and visualization in diverse industries including energy, manufacturing (design), medicine and government. Enterprises in these markets choose their hardware for integrated capability, looking for lower cost of ownership rather than the lowest priced system. With computing key to their business processes, they more than make up the money spent through the revenues enhanced by operational efficiencies and greater throughput. These customers are looking to grids not to aggregate processing over distances, but to share files and computing resources more effectively, and to make their applications real-time and pervasively accessible, despite the bulk of the data and processing involved. They are interested in developing after-markets for surplus cycles of their potent computing nodes, which is something the aggregative and extensive architecture of grids can facilitate. They also look to the grid concept to provide solutions for the following new challenges:

1. **The emergence of integrated sciences and the proliferation of monitoring and measurements has increased not just the bulk of data, but the relationships between data sets.** This complexity has made communication between processors, and memory structures, increasingly important, and it is important to use the right-sized engine (server or workstation) for the job. The resource aggregation of grids can help use these potent resources more efficiently.
2. **Better bandwidth has made remote access an inherent part of business functionality.** In SGI's customer base, remote workers need access to rich graphics presentations, often on ordinary laptops, ruggedized tablets or operating room monitors far from the servers that render the images. The distribution capabilities of grids are also important.
3. **Timely resolution of complex business decisions is becoming both more difficult and more important.** SGI's "immersive visualization" has proven a valuable tool for collaborative brainstorming where spreadsheets won't do. Both the aggregation (#1) and the distribution (#2) capabilities of grids can leverage immersive visualization to enable the acceleration of business decision making that underlies enterprise profitability.

SGI's take on grid computing is focused on limited markets, but involves a panoply of resources, including its hardware architecture, a new distributed file system, software from Platform Computing, and the capabilities of *Trusted IRIX*. **It is a story of building grids, not to do what you couldn't do, but to do it faster, better, and more pervasively.** Walter Stewart, SGI Spokesperson for Grids, calls grids "the convergence of high end computing and bandwidth." For more details, read on.

IN THIS ISSUE

SGI's Focus

For years, **SGI has focused on those scientific, technical and creative markets** for which its products provide critical functionality, such as:

- The exploration needs of the **oil and gas industry**, and the similar research needs of the scientific community,
- The parameter-tweaking design needs of **manufacturing and pharmacology**,
- The rich visual effects needs of the **entertainment industry**, and
- The complex analytics and modeling required by **government organizations** from NASA to the military.

These diverse markets share a need to deal with computatios involving massive quantities of data.

The current plague of intractable complexity may lead other organizations to see if visualization techniques can help them, where decomposition of the problem doesn't work. SGI will sell to, but not market to, such organizations. SGI has partnered with Hitachi Limited to take *CXFS*, one of the components of SGI grids, into other markets as Hitachi *Net Century*. SGI will sell IA-64 hardware, running Linux. *IRIX* will continue to run on MIPS processors, for which it is optimized.

What SGI brings to the Grid model

The grid model comes in many forms. There are enterprise grids, partner grids, and maybe – eventually - Internet grids. Common factor of all these grids are the holes (the allowance for distance somewhere in the equation) and the intersections (the resource aggregation). **The grid needs of any particular customer will depend on what scale of capability the customer needs – and where in the IT network the distance scalability is needed.**

SGI Nodes

SGI's nodes are optimized for high performance. SGI's *Origin* and *Onyx* servers are internally integrated with all processors sharing all the memory[1] through a crossbar switch, rather than a backplane.[2] SGI's workstations, *Octane* and the new *Fuel*, feature substantial memory and high memory bandwidth. *Trusted IRIX* allows multiple workloads to run on these processors with a B1 level of security safely isolating the workloads. SGI also supports clustering for failover (with *FailSafe* software) and for parallel processing (*Capability Clusters*), and a massively-parallel interface model for Grid implementations.

SGI Delivers Visual Area Networking

Solving the challenges of remoteness in the distributed enterprise cannot be done by potent hardware nodes alone. There is a matter of getting the rich graphic information feed to the end user's usually less sophisticated device. SGI's *Visual Area Networking (VAN)* is a thin-client alternative to the traditional client-server network where a server "serves" data to clients for local manipulation. Instead, with *VAN*, the data is centrally manipulated and pre-rendered. The results are compressed and streamed to clients.

Customers who routinely access large files know the pain of getting the data in and out of the computational system. Since, even with SANs, data has to go out over the LAN, moving data often takes more time than processing it. *VAN* capabilities like Dynamic Pipe Allocation and multiple compression modes[3] address this issue. APIs for third party authentication beyond the UNIX standard and a per-user usage log to enable chargeback give this solution enterprise functionality. The *VAN* works on a reservation basis. Massive rendering is not often a matter of whim.

With *VAN*, the engineer on an oilrig can explore the structure of the earth beneath the drill that is to be deployed. With *VAN* and *OP30*, a special application for surgical situations, a doctor can explore the three-dimensional rendering of a patient's MRI images in the operating room.

CXFS Enhances the Solution

Large files strain the scope and file-size limits of traditional, operating system or application-bound file systems**. Large file system users need a single name space for all shared files, addressable by all the servers directly**, not through *FTP*, *NFS* or *CIFS*. **They need support for a range of file types, and high file transfer rates, as well as the locking and failover of a**

[1] In other architectures, each processor has its own allocated subset of memory, with some provision of shared memory and over-provisioning.

[2] That's the NUMA of SGI's *NUMAflex* architecture.

[3] Plus an API for importing other compression algorithms.

traditional file system.

Where such shared access is critical and LANs don't supply enough bandwidth, SGI's global file system *CXFS* can help, either as software or productized with hardware as the *SAN Server 1000*[4]. It is targeted at large data sets and supports files as large as 9 EBs[5]. The tested average transfer rate is 15 GB/sec. CXFS's locking process is token based.[6] The token includes the metadata, which allows the data transfer to go directly from storage to the requesting server, not through the out-of-band metadata server, once the session has been negotiated. It can handle both 32 and 64-bit addressing and supports memory-mapped files.

CXFS is a journaled file system, which uses "I/O fencing" to segregate failing nodes. File system block sizes, selected on file system deployment, run from 512 bytes to 64 KB in normal deployments, or 1 MB in large-file deployments[7]. CXFS supports *IRIX*, *Solaris* and *Windows NT 4* clients. It will be ported to the rest of Windows later this year, and to Linux when McKinley servers are available. Hitachi will port CXFS to *AIX* and *HP-UX*. Non-CXFS servers can access files through *NFS*, *CIFS* or *FTP*. Authentication, replication, and HSM[8] are afforded by other SGI products. CXFS costs $1800 for dual processors on NT, and twice that on a UNIX server.

Grids Support Complex Business Relationships

While the processing in SGI grids is relatively local, its network connections are global, and ownership can be complex. An example is the grid sited at the University of Manchester in the U.K. SGI owns the hardware. Computer Sciences Corp. runs the grid. The University uses the grid academically, and CSC sells the surplus capacity to a large corporation overseas. Such **complex arrangements are increasingly common in this era of limited resources and collaborative projects**.

Grids need the ability to expense by use. SGI provides that ability today.

Grids and Computing as a Utility

SGI's Stewart feels that computing – as a utility – is a natural progression in our search for greater productivity. "Every major productivity increase has been based on our capability to harness physical power," he says. The result is more gain with less sweat. Enterprises must similarly harness computational power to optimize use of expensive computing resources. The aggregation and distribution capabilities of grid architecture provide the automatable infrastructure. SGI adds the graphics, auditing and charge-back capabilities.

Conclusion

For those enterprises whose information flow is inherently graphical or visual in nature, where a picture is worth a thousand pivot tables, **SGI's grids may enable a better and more streamlined business process to achieve your objectives.** These grids marry the tightly integrated processes of high performance computing, the graphic strength of SGI's visualization capabilities, and the bandwidth of today's networks to outflank enterprise complexity with visualization, served up through grids to wherever it is needed. **While other vendors are talking about grids, SGI is making them work for their targeted customers.** If you are working in any segment similar to those described, look at SGI's grid solutions.

---

4 *SANServer 1000*, a turnkey solution, has 32-client support. The base system, at $173,000, scales to 29 TB. Larger implementations are possible, as the file system scales much larger (get number)

5 9 Exabytes (Ebs) is 9 million terabytes (TBs).

6 *CXFS* supports *POSIX*, *BSD* and *SVR4* file locks.

7 *CXFS* allows "individual contiguous extents" of up to 1 TB.

8 Hierarchical Storage Management (HSM) is a way of moving files to the most appropriate tier of storage, based on their use.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland **is Director of Enterprise Systems Research with The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at AnneM@clipper.com or at 781-235-0085 Ext. 28.***

Regarding Trademarks and Service Marks

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.