GRID COMPUTING: THE NEXT (REALLY, REALLY) BIG THING

Interconnecting millions of distributed-computing devices

CASIMER DECUSATIS

Imagine a time in the not-too-distant future, when every computer-like device in the world—from mainframes to cell phones—is connected by a global system that dwarfs the current Internet in scale. A network so pervasive, so ingrained in all aspects of our daily lives, that nobody even thinks about buying a computer any more than you think about whether your air conditioner is getting its electricity from a coal-, nuclear-, or solar-powered generating plant. When you need to solve a particular problem, you just pick up a receiver and listen for a data tone to indicate that you're connected to the rest of the world—then draw on as much or as little processing power as you need, depending on whether you want to find the perfect anniversary gift or simulate gene splicing on a newly discovered chromosome. This network will be far too complicated for any single person to manage, but it will always appear to be available when you need it because the data servers will be able to manage themselves, automatically correcting problems, responding to unpredictable events, and fetching resources as needed. In fact, more than anything, this infrastructure will not only be intelligent, but it will also likely resemble a human nervous system in many respects. At the end of the month, you'll receive a bill for the amount of resources you've used and, since the resource is a commodity, you can switch to another service at any time.

While this seems like the stuff of science fiction, companies like IBM, Hewlett-Packard, and Sun Microsystems believe that this vision of the future is not only strategic, but inevitable, and they're betting billions of dollars on making it into reality. The concept is an extension of today's distributed-computing approach, in which a network of computing devices taps into a main server where important software and data reside. By contrast, so-called "grid computing" involves many thousands or millions of small, distributed-computing devices, linked together in local or regional networks, which are in turn interconnected on a global level. The grid functions as if it were one giant virtual computer, with very close integration of servers, storage, and other resources; everyone connected to the grid is able to share these resources in a very fast, efficient manner. Grids can be used to address so-called "Grand Challenge" problems, including advanced genetics research, modeling

---

Casimer is a senior engineer for IBM and can be contacted at decusat@us.ibm.com.

Are you rebuilding your engine every time you have to parse?

Speed. Power. Flexibility. That's what you want in a parsing solution. Whereas other tools require that you painstakingly generate and compile source code whenever your parsing task changes, with the **ClearParse™** engine you simply load a grammar and start parsing. Grammars can even be changed and loaded on the fly during run time. Weighing in at only 70 KB, is there any reason to parse without it?

ClearParse. Horsepower without the hassle.

- Small, efficient 70 KB parsing engine
- Practically unlimited backtracking capability
- Dynamic grammar handling
- Intelligent error recovery
- Parses multiple sources simultaneously
- Modular grammar language
- Source code available (ClearParse Professional)
- Run time royalty free

For more information, a free demo and documentation click ➡ www.clearjump.com/a/byte/clearparse

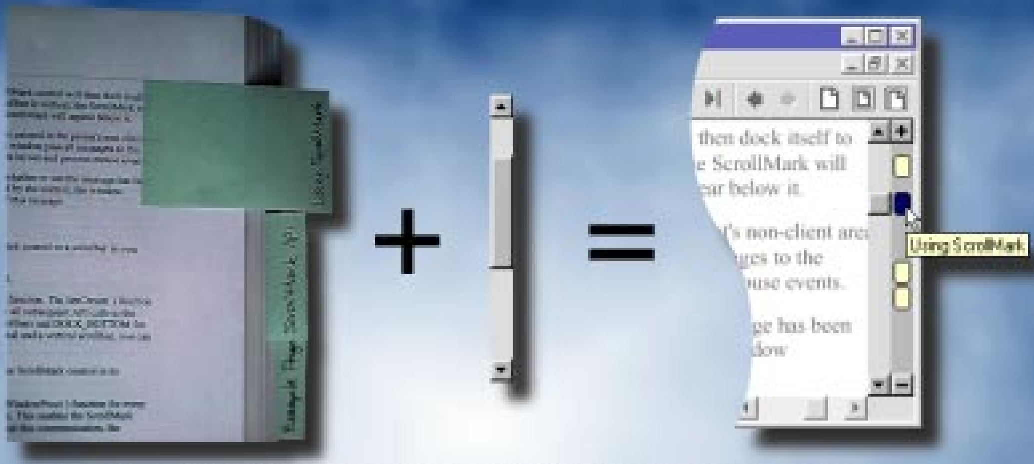

then dock itself to
e ScrollMark will
ear below it.
's non-client area
ges to the
ouse events.
ge has been
dow

Put bookmarks where they belong. On the interface.

Bookmarks can be a significant productivity enhancement for applications that access large amounts of text or data such as word processors, document viewers, programming editors and integrated development environments. So don't hide the functionality behind function keys and menu items. ClearJump's **ScrollMark™** control attaches to existing scrollable controls to make bookmarks clearly visible, intuitive to use and always within reach.

ScrollMark. The graphical bookmark control.

- Visual bookmarks
- Attaches to existing scrollable controls
- Variety of bookmark shapes
- Bookmark labels
- Tooltips
- Extensive customization capability
- Source code available (ScrollMark Professional)
- Run time royalty free

For more information, a free demo and documentation click ➡ www.clearjump.com/a/byte/scrollmark

ClearJump™
Get a jump on your development

www.clearjump.com

global weather patterns, or global air traffic control. This class of high-risk/high-reward problems is also known as "Deep Computing," and has commonly been addressed by interconnecting a large number of computer processors in parallel.

Lightweight versus Heavyweight Grids

While they hold the potential to be the next disruptive technology in the computing industry, the concept of grid computing isn't new. It can be traced back to early experiments with distributed, parallel processing, and so-called "lightweight grids" such as SETI@home, a continuation of the former NASA Search for Extra-Terrestrial Intelligence (http://setiathome.ssl.berkeley.edu/). Using free software downloaded over the Internet, your home computer scavenges spare processing cycles to analyze signals from the Arecibo Radio Telescope. There are currently over 1.6 million SETI@home subscribers in 224 countries, averaging 10 trillion operations per second and having contributed the equivalent of over 165,000 years of computing time to this project. This is arguably the world's largest distributed supercomputer, interconnected over the existing Internet.

Small-scale grid models have been under investigation for years by the University of Southern California, Argonne National Laboratory, and NASA, among others. Also, commercial applications have recently started to attract the attention of major corporations. Grids are still in their infancy, however; the ultimate goal is the creation of so-called "heavyweight grids," which are much larger and more powerful systems linked on a national or international scale. Today, heavyweight grids are under development in several countries, including the United Kingdom's National Grid and The Netherlands. Recently, IBM was selected to build the Distributed Terascale Facility (DTF), which would be the world's most powerful computing grid, capable of 13.6 trillion calculations per second and a storage capacity of more than 600 terabytes of data, or the equivalent of 146 million full-length novels. The Terascale Facility, with $53 million funding from the NSF, is a joint undertaking of the National Center for Supercomputing Applications (NCSA), the San Diego Supercomputing Center (SDSC),

Type of WDM	Number of Wavelengths (Channels)	Wavelength Spacing	Notes
Coarse (CWDM)	2 to 4	Very wide (typically 1300 nm and 1550 nm)	Very limited applications.
Wide Spectrum (WWDM) also known as Sparse (SWDM)	Up to 16	No standard defined, typical spacing of 1 to 30 nm have been used	Four-channel systems are under consideration for emerging 10 gigabit per second Ethernet. *
First-Generation DWDM	8 to 10	ITU grid **	Example: IBM 9729.
Second-Generation DWDM	Up to 16	ITU grid	
Third-Generation DWDM	Up to 32	ITU grid	Largest systems commercially available for datacom. Example: IBM 2029
Fourth- Generation DWDM (Ultra Dense)	Up to 64	No standard defined, smaller than ITU grid (0.4 nm proposed) ***	Not commercially available.

Table 1: *Classifying WDM systems. (* See IEEE 802.Az higher speed discussion group [10 gigabit per second Ethernet] at http://grouper.ieee.org/groups/802/3/10G_study/public/. ** ITU grid nominally centered on 1550 nm with minimum wavelength spacing of 0.8 nm [100 GHz] or multiples thereof, anchored to a reference of 193.1 THz per ITU G.MCS Annex A of COM15-R 67-E. *** See, for example, M. Ferries, "Recent Developments In Passive Components and Modules for Future Optical Communication Systems," Proceedings of the OSA Annual Meeting, 1999.)*

Argonne National Laboratory, and the California Institute of Technology. The grid will include the fastest supercomputers and high-resolution visualization environments, toolkits for grid computing, and data storage facilities integrated into an information infrastructure called the "TeraGrid," which will enable thousands of scientists around the country to share computing resources over the world's fastest research network in search of breakthroughs in life sciences, climate modeling, and other critical disciplines.

The Globus System

Systems such as these are based mainly on protocols, standards, and software tools under development by Globus, an open-source community led by Carl Kesselman at the University of Southern California Information Sciences Institute and Ian Foster of Argonne National Lab and the University of Chicago (http://www.globus.org/). In the same way that the Linux community has become a major part of open standards, Globus is working with the grid movement to help various standards and technologies to reach maturity; this includes developing tools to enable remote sharing of massive computing and storage resources, sophisticated resource management, scheduling and scalability routines, privacy and security tools, and other software functionality that is largely unavailable in today's peer-to-peer networking environments. Over the past five years, Globus toolkits have been deployed at over 20 multimillion dollar eScience projects around the world. IBM is a major supporter of grid computing, having used Globus technology to establish its own grid—a geographically distributed supercomputer linking research labs in the United States, Israel, Switzerland, and Japan. IBM has also established a state-of-the-art e-Utility Lab and testbed in Texas that is already working with a number of forward-thinking e-businesses on ways to extend grids beyond government labs and research environments into mainstream business of the Fortune 500 companies, and to grid-enable key systems and technologies. Emerging software standards such as Universal Description Discovery and Integration (UDDI) are important to this effort, as well as natural languages to analyze unstructured data such as web pages, HTML, video, and MP3 audio files.

Utility pricing is one of the features that makes grid computing attractive to large businesses; the increased flexibility this model affords could save billions of dollars alone, without even considering the other benefits of grids. Utility pricing, or pay-as-you-go models, are already offered by several vendors including IBM; both Compaq and Hewlett-Packard offers such models, which are likely to persist in some form following their recent merger. For example, IBM has announced its Dynamic E-Business initiative, which addresses business relationships between potential grid users, suppliers, and service providers. This approach is different from capacity-on-demand models, in which companies own or lease servers and pay an additional, fixed charge when they activate excess processing capacity. By contrast, in the utility model, the company doesn't own anything—it simply pays for its use of servers, storage, and disk I/O capacity. These resources are metered through software measurement tools and can be billed in various ways, usually either on a per-user or per-processor basis; either way, the cost per user goes down as the server utilization rises. This approach avoids one of the major pitfalls of today's information technology systems, namely capacity planning. Many corporations have difficulty accurately forecasting their demand for information technology (IT) resources. This has become more obvious during the recent economic slowdown, when anticipated revenues fail to materialize and overbuilt IT capacity becomes a burden most companies would prefer not to bear. Since utility pricing directly ties cost to revenue, it should allow companies to weather fluctuations in the market better by amortizing costs over time. The metering information can be used to determine exactly how much capacity is required under different system loads, ensuring that the end user always has adequate resources to run whatever their application demands. It's also envisioned that recurring monthly utility costs will prove to be lower than current purchased server models, although it is possible that a utility plan could cost more than a conventional approach if utilization runs above expectations for a long

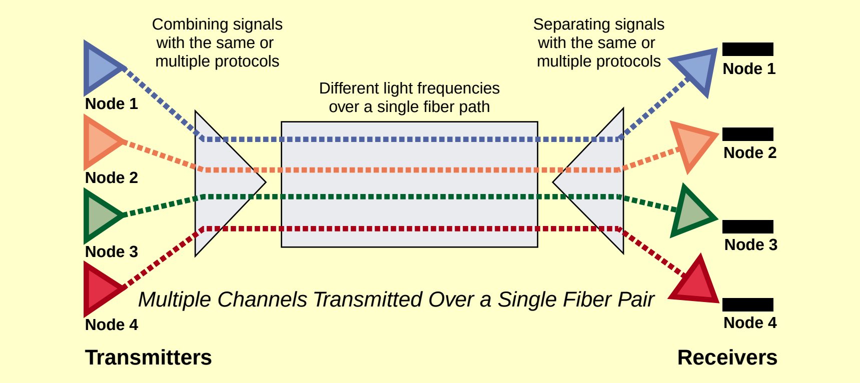

Figure 1: *Wavelength division multiplexing (WDM).*

period of time. Companies should also be able to deploy e-business initiatives faster and more easily respond to changing business conditions, since server costs would become part of a monthly operations budget, rather than an annual capital budget that must continually justify its return on investment. Leasing IT resources is also a good way to guard against technology obsolescence in a world that continues to experience Moore's Law-type growth in server capacity.

Dense Wavelength Division Multiplexing

In addition to software and management, another distinguishing feature of grids is an intelligent, broadband, low-latency network infrastructure, which provides much of the value in a grid computer. In fact, a high-bandwidth network is one of the most critical resources involved in enabling a grid, surpassing even the processor speed in its importance to overall performance. This is not a unique requirement of grids; recent trends toward clustered, parallel computer architectures such as the Parallel Sysplex (http://www-1.ibm.com/servers/eserver/zseries/pso/) and PowerParallel (http://www.rs6000.ibm.com/resource/technology/sp_sw2/spswp2_1.html) systems have also driven the need for high-bandwidth fiber optic coupling links between computers. Future grid computer nodes will likely be more tightly coupled in their data-sharing ability than the SETI@home systems. Instead, they might resemble a scaled-up version of today's largest clustered supercomputers. Current state-of-the-art parallel supercomputers consist of a few hundred nodes, each capable of executing perhaps a few hundred million instructions per second. A hypothetical grid computer may consist of thousands or millions of nodes, each capable of executing billions of instructions per second. A small grid network is likely to require multi-terabyte or petabyte bandwidth, and could easily grow several times larger than this. Such networks require either higher data rates per link or many low-speed links operating in parallel. This presents an opportunity for effective use of multiplexing techniques to handle the grid's increasing demand for bandwidth. Conventional SONET-based networks use time division multiplexing (TDM) to share many data channels across a single physical path; however, this approach does not scale well to the huge bandwidths required in a grid. Furthermore, grid computer nodes are distributed over tens to hundreds of kilometers or more. The combined distance and bandwidth requirements far exceed the capabilities of copper-based networks. Fortunately, optical networks offer a viable alternative, with theoretical available bandwidths of over 100 terabits per second on a single optical fiber (about 20 billion one-page e-mail messages or 2 billion phone calls). The data rate of individual channels in a grid will probably be in the range of 1–10 gigabits per second, and there is potential to combine hundreds of these channels over a common network using wavelength multiplexing approaches.

Wavelength division multiplexing (WDM) is emerging as the preferred technology for grid networks. It takes advantage of the fact that different wavelengths or colors of light will not interfere with each other when they are carried over the same optical fiber, as in Figure 1. The concept is similar to frequency multiplexing used by FM radio, except that the carrier "frequencies" are in the optical portion of the spectrum (around 1550 nanometers wavelength or 210^{14} Hertz). Thus, by placing each data channel on a different wavelength (frequency) of light, it is possible to send many channels of data over a common optical fiber. Data from nodes in the grid would be remodulated using distributed feedback laser diodes with a tightly controlled wavelength spacing or wavelength locking approach. The optical signals would then be combined into a single fiber using a diffraction grating, prism, or similar mechanism. Each node or group of nodes in the grid could be designed to receive only selected wavelength channels of data, and the network could be reconfigured in real time to add or drop any combination of wavelengths at a given node. More data channels can be carried per fiber if the wavelengths are spaced closer together. In this manner, WDM systems may be classified as either coarse, wide spectrum, or dense WDM (DWDM); see Table 1. Successive generations of DWDM equipment have supported more wavelengths, thus more channels, over a common fiber link. Current industry standards ratified by the International Telecommunications Union (ITU) establish a minimum spacing of 0.8 nm (100 GHz) and accommodate about 32 wavelengths per fiber. Future systems are expected to handle hundreds or even thousands of wavelengths at less than half the current wavelength spacing. Using optical amplifiers, DWDM networks can be extended to the distances required for grid networking.

The next generation of DWDM involves more than simply increasing the number of channels, however. New grid network topologies are also being enabled by the capability of third-generation and higher DWDM systems to support meshed or nested rings with dual redundant paths, self healing at the physical layer, data regeneration, and more advanced survivability or path protection than currently available. By contrast, conventional telecommunication networks have deployed ATM over SONET rings with separate overlay networks to accommodate IP data and other kinds of traffic as in Figure 2. The overlay networks are optically transparent, but remain service specific. Since grid traffic has different characteristics than voice traffic, overlay networks cannot make efficient use of the available bandwidth as they scale to multiple, concatenated rings. Consequently, optically transparent overlay networks are being replaced by a service-transparent DWDM core capable of allocating bandwidth on demand. This offers the advantages of a highly scalable, low cost, protocol-independent infrastructure, and may be the first step towards all-optical grid networks. Unlike first-generation DWDM, more recent technology allows wavelength multiplexing networks to be cascaded together because they act as a complete "3R" repeater that can:

- Retime the signal to remove jitter and improve clock/data recovery.
- Reshape the signal to remove pulse distortion caused by dispersion.
- Regenerate the signal to ensure that there is sufficient optical power to reach its destination.

Full protection switching at the physical or transport layer is available on a per-channel basis to restore service in the event of either a fiber break or failure of a hardware component in the system; a properly designed DWDM grid network would have no single points of failure.

The historical trend of growth in aggregate system bandwidth is likely to continue for the near future. However, this approach requires the service layer to upgrade more than twice as fast as the transport layer or roughly double capacity every six months (some estimates have shown the service layer growing over 70 times by 2003). A more realistic approach is to have the transport and service layers evolve together, although this still requires service layer capacity to double on a yearly basis. To keep pace with bandwidth growth, future DWDM systems may employ ultra-dense fourth-generation DWDM systems. The desire for efficient bandwidth management is also likely to drive the use of subrate multiplexing, or the combination of multiple TDM data channels within a single wavelength on a DWDM network. This hybrid approach may prove to be the most cost effective way to increase the total number of channels in a grid network. Future grid networks may also allow reuse of wavelengths for different purposes as the data hops between nodes on the grid. It has also become apparent that future DWDM technology will require some level of electronic signal processing in order to ensure the necessary data integrity, quality of service, reliability, security, and manageability of the network infrastructure. Previously, these features have been provided over the legacy telecommunications infrastructure. Indeed, SONET-based traffic is already carried over a physical layer that uses DWDM optical fiber interfaces. However, the growth of IP traffic has led to a complicated arrangement with up to four separate transport layers (IP over ATM over SONET over DWDM). In an effort to simplify this approach and streamline the data flow in future grids, there is a clearly emerging trend toward elimination of the ATM and SONET layers, and toward the direct transmission of IP or grid data over DWDM. This new model requires the DWDM layer to assume many of the traditional functions associated with ATM over SONET, such as protection switching. In particular, two standards efforts are currently under way to link data packets directly to DWDM optical wavelengths, so that the optical network can take some advantage of the intelligence embedded in IP traffic. This so-called "optical IP" effort could eventually let grid users dynamically request portions of a fiber cable's bandwidth for a particular time or service. One such standards effort is the Optical Domain Service Interconnect (ODSI) coalition (http://www.odsi-coalition.com/), a loose connection of vendors providing optical transmission equipment, access services, terabit routers, switches, and network provisioning software. ODSI seeks to define common control interfaces between optical or electrical physical layers and IP media access layers of the Open Systems Interconnect (OSI) model. This work may ultimately rely on derivatives of the Internet Engineering Task Force (IETF) Multi-Protocol Label Switching (MPLS), a means of defining IP flows that is already widely used in the electrical signaling domain. ODSI will propose low-level (below Layer 3) control plane standards that must be met by vendors of both optical transmission equipment and broadband IP switches/routers. A separate

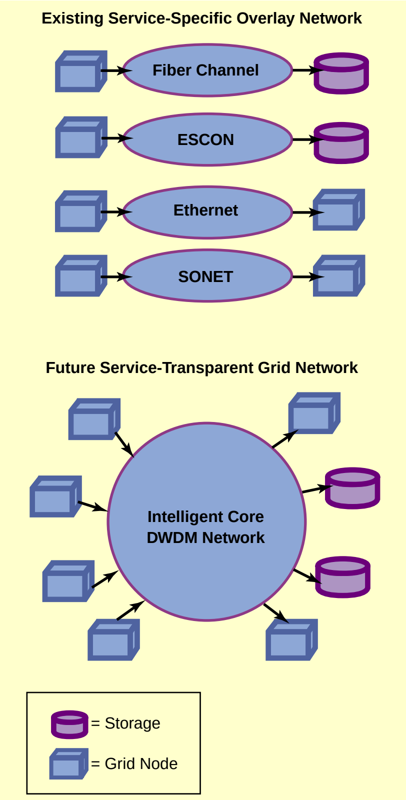

Figure 2: *Deployment of SONET overlay networks versus protocol-independent DWDM core networks.*

but complementary effort, also based on MPLS, is being drafted within the IETF itself. Known as Multi-Protocol Lambda Switching (MPLmS; "lambda" refers to switching by native wavelengths), this effort has proposed methods to link optical cross-connects and high-bandwidth routers through a Layer 3 switching methodology. Operating at a higher level than ODSI, this proposal has the advantage of not requiring a new set of protocols for quality-of-service and bandwidth control

We can see how DWDM greatly increases the available bandwidth of installed optical fiber. Today, about 32 wavelength channels are commercially available, each operating at peak data rates between 1 and 2 gigabits per second; in the near future, hundreds of wavelengths will be achievable at data rates exceeding 10 gigabits per second. This enables bandwidth allocation on demand in a grid computer backbone, which may be based on all-optical switching technology or wavelength routers using MPLS. Optical networks enable grid computing and are an important prerequisite to sharing computing resources on a commodity basis, much like the telephone or electric utility infrastructures. IBM and many others see the utility model for delivering computer resources on a pay-as-you-use basis as the future of corporate computing. Today's Internet or application service providers could evolve beyond simple web hosting to offering a full range of services and solutions, including remote backup for disaster recovery and dial-up computing power on demand. Some feel that with the convergence of the telecommunications and data communication worlds, today's phone company central office may grow into tomorrow's Internet Data Center, providing a wide range of content-rich services over a nationwide all-optical DWDM backbone, using emerging technologies such as resilient packet ring (RPR), Optical Ethernet at 10-gigabits-per-second data rates and beyond, and optical MPLS to route data anywhere, at any time.

While the cost of ownership for servers is already being driven sky-high by management and maintenance fees, the sheer size and complexity of grids lends a new meaning to the idea of network management. Indeed, some feel that worldwide grids will be too complex for conventional management approaches to be effective, and that management of ever-growing networks constitutes a looming crisis in the Internet community. For these systems to succeed, both resource and network management must be highly automated. The future of grid computing depends on a core of intelligent middleware that is capable of handling most networking decisions without human intervention. This is the concept of "autonomic computing," named after the body's autonomic nervous system. The autonomic nervous system allows human beings to adapt to any number of situations by unconsciously adjusting heart rate, breathing, body temperature, pupil dilation, etc. The body spends a good deal of time each day fighting off viruses, digesting food, or secreting hormones. But the individual remains totally unaware of this complex activity, how it works, or when it is needed. Similarly, an autonomic computer is self managing and self healing, responsive to changes in its environment, and always accessible. Servers that make up a grid computer will be modular and distributed much like the cells in the central nervous system. This cellular architecture integrates processors, memory, and communications in a way that minimizes the latency considerations that limit current computer architectures. In reality, the nervous system behind a global computing grid would likely be an aggregation of smaller, autonomic networks working together. This middleware needs to be platform agnostic, supporting a heterogeneous mixture of access devices ranging from conventional PCs to network appliances, embedded devices, kiosks, or cell phones. In addition to network management, middleware may also perform functions such as data caching to improve performance, directory and security control, quality of service enforcement, and transcoding (changing data structures to take on the characteristics of the end user's access device).

The eLiza Project

Although the name of IBM's eLiza Project (http://www-1.ibm.com/servers/eserver/introducing/eliza/) may conjure up visions of a certain character from the musical "My Fair Lady," it actually refers to the first artificial intelligence program to permit natural language conversations between a human and a computer, written decades ago by Joseph Wiezenbaum. IBM has adopted a lizard as the symbol for this program; a few years ago, using assumptions in Ray Kurzweil's book *The Age of Spiritual Machines* (Oxford University Press, 1999; ISBN 0195129423), IBM researchers estimated the processing power of the chess-playing Deep Blue supercomputer as roughly equivalent to a lizard's brain. While it may seem odd for such a simple creature to need so much brain power, the lizard actually has a far more complex task than a single-purpose computer. While the computer's environment is completely controlled and comparatively predictable, the lizard faces a constantly changing jungle environment requiring instantaneous responses to new demands and dangers. Grid computers have their own jungles to contend with, complete with totally unpredictable demands from unknown amounts of users, sudden threats from predators, and frequent natural shocks that stress the system. Like the lizard, grid computers require intelligent self-management skills to survive in this environment.

Facing the seeming paradox of finding the right people to manage the effort to create self-managed servers, IBM selected a team led by Irving Wladawsky-Berger and Ross Mauri. It is directed toward producing self-managed computer systems hundreds of times more complex than those that exist today. A self-managed computer could, for example, update and maintain its own software at latest release levels, actively protect itself against attacks by computer crackers, monitor its health and perform limited acts of self repair, correct inadvertent human management errors, and guard against unexpected problems or acts of nature. In this way, an eLiza system would ensure its own survival and stability. It would also continually self optimize, eliminating the need for constant tweaking by network administrators, and adjust itself to changes in the server loading. For example, an eLiza server running the stock market could detect sharp upsurges in trading and adjust quickly to meet the higher demand. In addition to providing the means to control systems that are orders of magnitude larger than any existing today, it is hoped that these systems would be significantly lower in cost to own and maintain. While eLiza is still under development, some important related steps have already been taken this past year by IBM, including the release of Intelligent Resource Director (which dynamically allocates server resources for multiple jobs according to demand), eServer clustering (which updates software and manages workloads across massive server clusters), and self-healing technologies such as Chipkill (which eliminates memory failures and allow server components to function for decades without failing). Other aspects of self management are evident in current IBM products, such as the z/OS 64-bit operating system that automatically shifts server capacity between logical partitions in response to high-priority workload spikes.

There are many examples of emerging technologies that will participate in grid computers. One example of technology currently under development is Oceano, a prototype self-managed server farm currently running at IBM Research. Oceano lets a group of servers be automated to handle on-the-fly changes in load requirements, or multicustomer hosting on a collection of virtual hardware. For example, running an Internet application like the World Series online requires a collection of servers to be connected together and loaded with the appropriate web server software; more baseball fans can be accommodated by plugging in more servers. Oceano can automatically sense what the demand is likely to be, bring down servers not being heavily used for other applications, then bring them back online running the World Series web site application, all without manual intervention. This is expected to be a great productivity boost for IT staff, freeing them from the more mundane tasks to focus on other areas. It may also help alleviate projected shortages in trained IT staff (some analysts estimate that in five years, the world will be short at least a million qualified IT administrators; others have noted that a network of 1 billion users would require about 250 million skilled people, almost the entire population of the U.S.). It's widely recognized that there are some things computer systems can't do for themselves, and although their job descriptions may change, systems administrators aren't going to disappear.

However, if eLiza succeeds, the end result could be a global computing network that will be as easy to control as a kitchen appliance. The first self-healing supercomputer may be a mammoth machine, nicknamed "Blue Gene," under development by IBM fellow Monty Denneau, which will be dedicated to unlocking the secrets of protein folding. This is a massive computational problem, and Blue Gene is being designed for the task; when completed, it will be capable of well over a quadrillion (10^{13}) operations per second, 40 times faster than today's top 40 supercomputers combined. Using a total of over 1 million processors, it will occupy the size of a tennis court and require about 10 terabytes of attached storage.

Conclusion

There are many obstacles to overcome before heavyweight grids become a reality, the least of which are the technology challenges of building self-managed, protocol-agnostic systems over high-bandwidth networks. But there are also challenges adapting to new business models; grids allow organizations to think about their infrastructure in new ways. For example, the ability to transparently share resources, data, and applications might allow companies to establish temporary joint ventures—forming a virtual subsidiary that exists only long enough to address a specific business opportunity before being dissolved again. Most analysts agree, however, that grid computing represents a potential multibillion dollar worldwide revenue opportunity, and that it offers real efficiencies that reduce cost, accelerate time to market, and enable the solution of problems that lie beyond our current computational abilities. If current efforts are successful in the coming years, grid computing could literally be the next big thing in distributed computing. ■

Return to Table of Contents
Next Article