

Sponsored by:

NetworkWorldFusion

This story appeared on Network World Fusion at
<http://www.nwfusion.com/research/2003/0106grid.html>

Feature /

Grid-dy determination

Grid computing systems gird to move out of the lab and into the enterprise.

By Curtis Franklin Jr.
Network World, 01/06/03

Stu Jackson needs CPU cycles - lots of CPU cycles. As IT architect for Incyte Genomics, Jackson designs systems that use computing resources the way a blast furnace uses iron ore. The Palo Alto firm's genomic applications burn up every available CPU resource.

Jackson doesn't need supercomputers, however. He builds his applications for pharmaceutical and biotech firms on computing grids. "For businesses that consume CPU cycles as a raw material, grids make sense in almost every case," he says.

Organizations have spent large sums of money building their computing infrastructures, which primarily consist of computers that spend a lot of time doing nothing. Harnessing those unused CPU cycles to power compute-intensive applications is the driving idea behind grid computing.

A [grid computing](#) system is a distributed parallel collection of computers that enables the sharing, selection and aggregation of resources. This sharing is based on the resources' availability, capability, performance, cost and ability to meet quality-of-service requirements.

Grids come in various sizes, from cluster grids that pull workgroup computers into a single system, to those that link clustered computers, to enterprise grids that tie computers in a single organization, to global grids that tie computers from multiple organizations into massively parallel high-performance computing engines.

There also are several types of grids, from the traditional grids that focus on aggregating CPU horsepower, to data grids that move terabytes of data between sites for analysis, to access grids that provide high-performance video conferencing and application sharing between multiple sites. Each grid, no matter the size or type, is tied together with job scheduling and management software.

[Avaki](#), [DataSynapse](#), [Entropia](#) and [Platform Computing](#) are four companies specializing in grid management and scheduling software. Entropia specializes in linking PCs into parallel-computing grids. The other three focus on high-performance servers and midrange computers. All are building products based on the Open Grid Services Architecture ([OGSA](#)), a standard developed by the [Global Grid Forum](#), a trade group seeking

to create a common basis for grid computing. In addition to the commercial offerings, the [Globus Project](#) has developed an open source grid framework based on OGSA standards.

[Hewlett-Packard](#), [IBM](#) and [Sun](#) each have developed grid initiatives based on their own hardware. While each has unique elements, all claim allegiance to the OGSA standard. Dan Powers, vice president of grid computing strategy and business development at IBM, says rallying around a standard is a must for the growing grid market. "We didn't need eight different ways to build networks, so we ended up with TCP/IP. We don't need eight different ways to build grids," Powers says.

Grid to go

Grid computing's first moves out of the academic and research arenas have been into compute-intensive applications. Bioinformatics, oil and gas exploration, automotive and aerospace engineering, and financial services industries were among the early corporate adopters.

Financial services firms are using grid computing to prepare complex models of individual currencies or complete portfolios, and get the results quickly enough to trade based on the model's predictions.

Frank Cicio, COO of grid computing vendor DataSynapse, says there's no mystery behind the move. "On Wall Street, turning information around in real time that normally takes hours could mean billions of dollars. Everybody is watching the dollar, and everybody wants more for less."

Grid computing has been cost-effective for Incyte Genomics. The company moved from a 32-processor Sun E10000 to an Intel-based grid running Platform Computing's software, and Jackson's price/performance calculations show the grid is about 10 times less expensive for the same computer power.

Incyte has used some form of grid system for nearly five years. "Five years ago we were clustering 50 to 100 Alpha processors, where today we tend to use Linux on an Intel platform," Jackson says.

Another reason for the grid deployment is ease of upgrading. "Our first grid was 125 processors, and we've used as many as 1,000 processors for the same application," he says.

Information for life-sciences applications also is the focus of the North Carolina Bioinformatics Grid Project in Research Triangle Park, N.C. Phil Emer, chief architect for the project, says the organization has built a grid incorporating hardware and software from Avaki, Platform Computing, IBM and Sun.

Emer says the project didn't start with the goal of building a grid; the grid architecture grew out of the needs of several organizations. "By the time we looked at our requirements - high-performance computing, high scalability, user interface transparency - we had described a grid," he says.

The grid spans computers at three universities, several commercial and government research facilities and the North Carolina Supercomputing Center. Emer found the organizational and accounting challenges were at least as great as the technology hurdles. "The human policies are significant issues. You have to put in place enough monitoring applications to prove to institutions that, by cooperating with the grid, they'll get out more resources than they put in," he says.

Grid computing was a good choice for getting start-up Butterfly.net to take flight. Developers of a framework for multiplayer online games, Butterfly.net sees the demand for computing resources vary in a short time, says [Butterfly.net](#) CEO David Levine. The company built its infrastructure on Globus because of its ability to run on a Linux platform and uses IBM's global grid to provide resources for game developers and players around the world.

While the largest game hosted so far has about 50,000 concurrent users, Levine says they have to prepare for more. "Some games being ported over from China already have millions of players, so when we first put the infrastructure in place we had to have resources for a million players," he says. The company's contract with IBM let it underbuy in the early stages of the company's life, but scale to accommodate more users as needed.

Grinding gears

Grid computing architectures provide advantages in performance and flexibility, but there are still issues keeping many companies from leaping too quickly onto the grid bandwagon. Questions of scheduling and management, security and accounting make grid computing a risky proposition for many IT executives.

Patricia Kovatch is manager of high-performance computing at the San Diego Supercomputer Center (SDSC). She's involved in building the TeraGrid, a large grid connecting systems at four premier high-performance computing centers - SDSC, the National Center for Super-computer Applications (NCSA), Argonne National Laboratory and Cal Tech.

Scheduling and management issues present a big challenge in building the TeraGrid. "You need a metascheduler so each piece of the program can run on different computers at the same time. The tasks have to talk to each other and make sure that data is returned to the central control portion of the application. There are still a lot of problems that aren't solved, and that's part of what this project is about," she says.

When computing resources are aggregated, security can become a significant issue. The basic issues of user authentication and access control suddenly are multiplied by the number of clusters, departments or organizations that link to form the grid. Questions of who can create a job of a particular priority, which resources can be accessed and other questions are part of any grid that multiple users can access.

Standardizing methods for enforcing policies - such as security and financial - is a major thrust of standards efforts such as the OGSA. While major players in the industry are behind the standardization effort, the history of technical standards committees is not filled with standards developed as quickly as the market would like.

There's also the issue of who pays for all those CPU cycles. The accounting applications are made more complex because the task is not simply about traditional IT costs, says Bob Fabbio, CEO of software developer [Vico](#).

"You're setting up a financial marketplace in the computing center, so you're matching supply and demand," he says. "You have to look at application service levels and have a sophisticated understanding of the infrastructure beneath the application." Accounting management for grid systems has not developed at the same rate as application support. Until corporations can adequately account for the use of resources, grids will remain platforms for single applications rather than many applications for a variety of departments.

The combination of security, administration and accounting issues has resulted in most grids being centered around computers from a single vendor or based on a single operating system. Though the promise of grid computing is shared resources regardless of the underlying platform, building a grid based on multiple hardware and operating systems involves massive customization efforts.

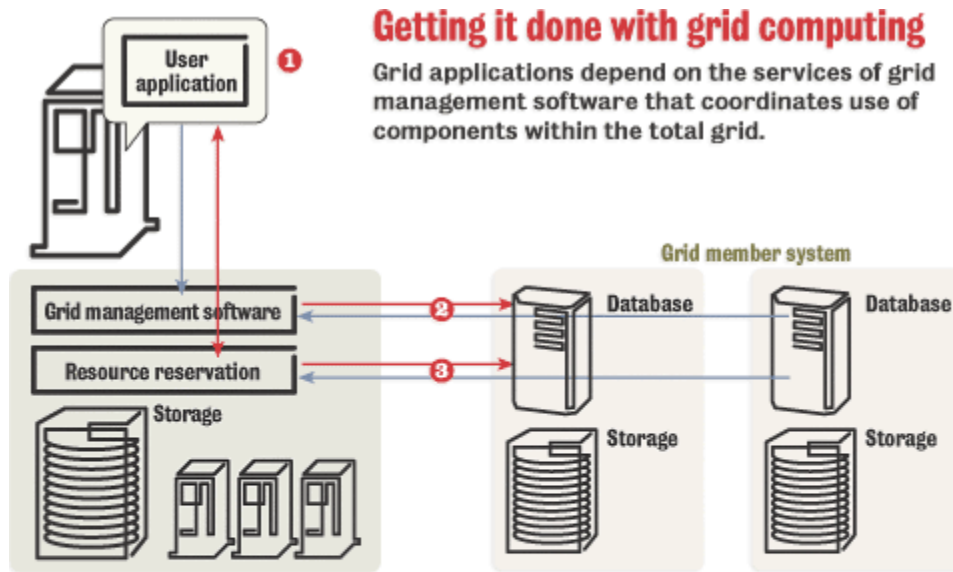
Even single-platform grids are custom efforts today, and few organizations are willing to commit to that level of additional effort.

The future grid

Though commercial grids still are moving through the early adopter stage, Levine is bullish on the

technology. "In five years, I can't imagine a company not using a grid," he says.

Other observers are more cautious about the time scale, but not the ultimate results. "I really do think [grids] will become the way to share resources within and among enterprises," says Jane Clabby, research analyst at Bloor Research. "Within five to 10 years we'll be talking about grids the way we talk about the Internet today."



The major components of a commercial grid system include grid-enabled user application, resource scheduling and management software, and multiple database servers.

- 1 Initial application request for data goes to the grid software, which reserves resources on available systems within the
- 2 Updates take place between grid member systems and the management software, and between the management software and user application.
- 3 When processing is complete, the management software releases grid computers and results are returned to the user application.

RELATED LINKS

Franklin is an editor and writer in Gainesville, Fla. He can be reached at kg4gwa@arrl.net.

[The Global Grid Forum](#)

A community-initiated forum of individual researchers and practitioners working on distributed computing, or "grid" technologies.

[Globus Project](#)

A group developing fundamental technologies needed to build computational grids.

[TeraGrid](#)

A multi-year effort to build and deploy the world's largest, fastest, distributed infrastructure for open scientific research.

All contents copyright 1995-2003 Network World, Inc. <http://www.nwfusion.com>