Stacey Ann Phillips

03/11/2003 02:18 PM

To: Benjamin Khoo/Singapore/IBM@IBMSG, Michael Monty/Singapore/IBM@IBMSG
cc:
From: Stacey Ann Phillips/Singapore/IBM@IBMSG
Subject: Asia Computer Weekly - Supercomputing on demand

Ben, Michael,

Thanks for doing the interview. The last paragraph is particularly interesting.

Singapore (November 3 - 9) Asia Computer Weekly - Supercomputing on demand

Singapore (November 3 - 9) Asia Computer Weekly - Supercomputing on demand

By Ong Boon Kiat

Like the first wave of e-business, grid computing development is now at a stage where enthusiasm runs high, but so does confusion.

This is because while all the building blocks are in place, a common infrastructure has yet to emerge. And like e-business was five years ago, the term grid computing is still very much a moving target, with different vendors offering different takes on the technology.

Judging from the growing number of commercial deployments around the world, grid computing is definitely moving out of the science and campus networks and into commercial organisations. But like the evolution of the Internet for business use, the move into a more stringent environment means that more changes will happen before the dust settles. So how will these changes look like?

A smaller grid

For one, the definition of grid computing will continue to shift. The most recent shift came last month, when database giant Oracle entered the grid computing fray, joining IBM, Sun Microsystems, Hewlett-Packard, Platform Computing, Avaki, DataSynapse and others.

Oracle announced a grid computing-compatible "10g" database and application server, which lets companies combine database application servers to form a self-managing computing entity.

10g applies grid computing to a database-centric environment. Building grid computing components into both its database and application servers, Oracle's grid computing model is essentially purpose-built to cluster application servers which can in turn work on a virtualised database.

By intent, Oracle's approach is narrower in scale than, say, the multi-party grid computing deployment offered by IBM. 10g's "closeness" have even led some vendors into criticising it as not being a true grid computing platform, merely an extension of Oracle's existing Real Application Clusters (RAC).

Contrast this with a recent IBM grid computing implementation for investment bank JP Morgan, which comprised task scheduler application from DataSynapse, grid domain manager from Avaki, software agents to manage data access between connected devices, and IBM's WebSphere to stitch the various pieces into a coherent platform. IBM's deployment in JP Morgan spans a wide gamut of applications and hardware platforms, and is clearly wider in scope.

But while it is natural to laud a more open platform, a narrower one is not necessarily less desirable in this case. Firstly, the notion of grid

computing remains a political minefield for many organisations. "As far as building a bigger grid goes, the biggest obstacle is people, not technology," pointed out Peter Thomas, senior director of technology, Asia Pacific, Oracle.

It is hard—impossible in some cases—to convince different organisational departments to surrender their precious servers and storage arrays into a common commune, he said. As a result, building smaller grids are the most viable way to start for most organisations. Smaller and more homogeneous grids are also easier to manage.

Another reason is security. A large scale grid computing deployment spanning across public networks can be risky due to increased exposure. "Policy negotiation, especially with regards to security, will be a key hindrance (to large-scale grid computing) adoption," said K Sudershan, director, Infrastructure Solutions, Enterprise Systems Group, HP Asia Pacific.

Which means that many organisations are likely to take to the safer approach of deploying relatively enclosed grid computing within their firewalls in the near future.

Killer apps

"Killer application" is a favourite term for technology pundits and at least one of grid computing's killer applications may have already arrived. Oracle's participation in grid computing offers this technology its most accessible application to date: the database and the rich software market that surrounds it, including ERP, SCM, CRM and data warehousing.

In the longer term, Web services is probably grid computing's best application ally. "Web services, which are essentially connectors that let applications communicate with each other, will be an effective and complimentary interface to grid computing," said Sudershan.

Think of Web services as "sockets" for an electrical power utility, he said. An underlying grid computing infrastructure will make Web services more powerful by efficiently provisioning the hardware resources needed for a Web service application.

Not surprisingly, both technologies are likely to be closely aligned at the standards level. The de facto open standard for grid computing, the Open Grid Computing Architecture (OGSA), already defines grid service in terms of the Web Services Definition Language (WSDL) with extensions. This makes it possible to invoke grid services using existing Web services standards like SOAP, XML, and WS-Security.

According to Benjamin Khoo, grid computing architect, IBM Asean, South Asia and ANZ, IBM Global Services, the standards for both technologies will overlap and eventually merge in future releases. "OGSA and WSDL will come together, and the impact on enterprise computing will be immense," he said.

One such groundbreaking impact is the fact that application binary codes will be able to run across all OS and hardware platforms without modification. "Web services is the universal application operating system, and if people write codes according to WSDL and OGSA, applications can truly run on heterogenous systems in future," he said.

Sounds plausible? It gets better. Khoo believes that when Web services integrate tightly with grid computing, applications will become more hardware-aware, capable of invoking such tasks as computing aggregation or database federation when it encounters the appropriate hardware. "Web

services is designed at the applications communication level today, but grid computing will bring that to an even lower level," he said.

Standards
Another key area in the grid computing technology that should see much work over the next two years is the development of its core and related standards.

The main grid computing standards body today is the Global Grid Forum (GGF), which defined OGSA. Web services-based interfaces for files and databases are beginning to be defined at GGF in the Database Access and Integration Services Working Group (DAIS-WG). Standardisation and compliance by key vendors will be essential for interoperability and widespread adoption.

But standards are very much a vendor concern. For users, what may be more important is understanding the basis of what constitutes a sound grid computing infrastructure. Given its hype, it is easy to forget that grid computing is not about high performance per se, but on getting good hardware utilisation. Which means that an organisation's existing hardware must be good enough to start with.

"If your computing resources have lousy performance, grid computing won't save you," said Michael Monty, business development executive, Grid Computing business, IBM Asean, South Asia and ANZ.

He gave an analogy which drives home the point: "If you have two dying donkeys and one heavy load, spreading that load evenly between the two won't miraculously give you the energy of one healthy donkey." Hee-haw to that.

Regards,
Stacey Phillips
Corporate Communications
IBM Singapore
80 Anson Road IBM Towers Singapore 079907
Tel: (65) 6320 1569 Fax: (65) 6223 1150 Tieline: 919-1569 e-mail: phillipa@sg.ibm.com