

26 February, 2003

What is Grid Computing?

For a number of years, the economics of high performance distributed computing have been changing dramatically. Servers and storage have continued to rapidly improve their "price for performance" by leveraging new innovations and manufacturing efficiencies, and the same trend has finally taken hold for bandwidth. The effect is to transform distributed computing into a competitively priced commodity. At the same time, TCP/IP has become the only networking protocol suite considered, and UNIX (TM) or Linux has become the operating system of choice for scientific computing. In contrast to the commoditization of technology, skilled people have remained scarce, and concerns about security and quality of service have increased. The logical response to these changes is to move from a model based on discrete infrastructure components to a distributed computing model which fully leverages the computing capabilities of the infrastructure. This distributed computing environment operates as a uniform service, which looks after resource management and security independently of individual technology choices. In other words, distributed computing is treated like a *public utility* comparable to the electricity grid.

In order for users from multiple organizations to make use of a computing grid efficiently and securely, they must belong to a *virtual organization (VO)* sharing a common goal and a common set of resources. Assigning resources, users, and applications to VOs is the fundamental Grid technical value proposition.

These ideas were first described in *The Grid: Blueprint for a New Computing Infrastructure* edited by Ian Foster (Argonne National Laboratory, USA) and Carl Kesselman (University of Southern California) [Morgan Kaufmann Publishers, 1999, ISBN 1-55860-475-8]. Today, their ideas flourish in various practical implementations, most notably the open source Globus toolkit developed by Foster, Kesselman and their collaborators (<http://www.globus.org>) and the Unicore project (<http://www.unicore.de/>). Additionally, a number of companies have proprietary Grid computing products.

The current Globus toolkit builds on well-known Internet standards such as SSL/TLS for security and LDAP for service directories. It contains components such as

- Grid Security Infrastructure
- Grid Resource Allocation Management

Expanded Coverage from ISOC

In-depth articles, papers, links and other resources on a variety of topics are available from the ISOC site at: www.isoc.org/internet/issues

Examples in the News

Feb 5, 2003 IBM, *United Devices and Accelrys Aid U.S. Department of Defense in Search for Smallpox Cure* <http://www-916.ibm.com/press/prnews.nsf/jan/A1E7D8AEA826518C85256CC4005E575E> IBM, United Devices and Accelrys today announced a project supporting a global research effort that is focused on the development of new drugs that could potentially combat the smallpox virus post infection. The project will be powered through a massive computing "Grid" that will enable millions of computer owners worldwide to contribute idle computing resources with the goal of developing a wide collection of potential anti-smallpox drugs.

For More Information

www.gridforum.org
www.globus.org/ogsa
www.gridcomputing.org
www.thegridreport.com
www.unicore.de
<http://eu-datagrid.web.cern.ch/eu-datagrid/>

- Interface to resource brokers (e.g., LoadLeveler, Condor Matchmaker)
- Grid Resource Information Service
- GridFTP

Standards are being developed for Grid computing by the Global Grid Forum (GGF) (<http://www.gridforum.org>), using procedures directly modeled on those of the IETF. In addition to many academic and research institutions, at least thirty industrial companies are committed to Globus, or involved in the GGF, or both. These include, for example, Microsoft, Hewlett Packard, IBM, Sun, Dell and BEA.

Industrial interest and the link to Web Services

This raises the question of *why* so many companies are interested in a trend within the research community. Consider the economic factors mentioned above: the commoditization of technology and the expense of skilled people. These apply at least as much to the business community as they do to research institutions. At the same time, the pressures on commercial IT managers are towards cost reduction, more efficient utilization of resources, and of course increased security and quality of service. These are *exactly* the same factors and pressures that brought about the development of Grid computing in the research community, so it is in fact no surprise that industry is so interested.

However, the initial response to these factors in the computing industry was in fact the development of Web Services standards and products. Web Services refers to the set of standards developed in the World Wide Web Consortium (<http://www.w3.org>), OASIS (<http://www.oasis-open.org/>), and elsewhere to permit the implementation of sophisticated distributed services. The basic mechanism of Web Services is the Web Services Definition Language (WSDL), based on XML. WSDL documents are exchanged between service requesters and services, thus avoiding the need for the requester and a service to run the same underlying software - a Web Services run-time system is likely to be implemented in Java (TM), but other solutions are possible as long as they communicate using standard WSDL documents. WSDL is wrapped in the SOAP protocol for transmission, and SOAP can in principle travel over any transport protocol, although HTTP is mainly used today. With Web Services, it is possible to build distributed e-business computing systems that run on heterogeneous, commoditized servers. For further information, see S. Graham, S. Simeonov, T. Boubez, G. Daniels, D. Davis, Y. Nakamura, R. Neyama, *Building Web Services with Java: Making Sense of XML, SOAP, WSDL, and UDDI*, Sams, 2001.

During 2001, partly stimulated by discussions at a Global Grid Forum meeting held in Amsterdam, the Web Services community and the Grid community started to think about the obvious synergy between their objectives, and whether common solutions could be found. In fact, the use of LDAP for resource and service descriptions in Globus was the

About the Author



Former Chair of the ISOC Board of Trustees, Brian Carpenter is an IBM Distinguished Engineer, active in Internet Standards and Technology. Carpenter coordinates IBM relations with the IETF and works on related technical strategy. He is also a former Chair of the Internet Architecture Board (IAB), and currently serves as co-chair of the IETF Differentiated Services working group.

Acknowledgments

The ISOC Member Briefing series is made possible through the generous assistance of ISOC's Platinum Program Sponsors: APNIC, ARIN, Microsoft, Qualys, and Ripe NCC. *More information on the Platinum Sponsorship Program* : <http://www.isoc.org/isoc/membership/platinum.shtml>

clue to exploiting this synergy. Both Grids and Web Services need such descriptions, but LDAP does not truly offer the sophistication and flexibility needed for a description language. In contrast, WSDL was designed precisely for that purpose. Hence the idea arose of using the Web Services mechanisms systematically for the interactions between requesters and services within a Grid - thus allowing both service and resource descriptions, and actual interactions with services, to be expressed in a system-independent language, WSDL. This proved to be a very powerful idea and led rapidly to an outline design known as the Open Grid Services Architecture (OGSA) (<http://www.globus.org/ogsa/>). There is a good general introduction in I. Foster, C. Kesselman, J. Nick, S. Tuecke, *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration* (available from the Globus web site). At the present time, the Global Grid Forum is developing details of OGSA standards in three working groups

- The main OGSA working group concentrates on the architecture and roadmap.
- The Open Grid Services Infrastructure (OGSI) working group is defining very precisely the required WSDL structures required in an OGSA environment.
- The OGSA Security working group is developing a security architecture expected to be based on the recently published Web Services security proposals (see <http://www.ibm.com/developerworks/library/ws-secroad/>). Exemplary security is unavoidable in a world where users from multiple real organizations collaborate to share resources in a single virtual organization.

The Globus team announced some time ago that the next version of the Globus toolkit will be based on OGSA, and in fact they released "alpha" OGSA code in January 2003. A number of industrial companies, notably Avaki, Entropia, IBM, and Platform Computing, have indicated plans to support OGSA. Thus, and almost uniquely in the history of the distributed computing, we see a strong convergence between leading edge trends in scientific and business computing. As open source and proprietary implementations of OGSA appear, we will see a rare opportunity to construct robust, secure heterogeneous distributed computing systems.

Implications

If we assume that Grid computing and OGSA standards lead within a year or two to a standardized way of building distributed systems supporting virtual organizations, how will this transform the way we use the network?

Remembering the underlying economic factors, it is no surprise that the IT departments of commercial companies are anxious to control costs and avoid installing excess capacity, at the same time as they insist on increased security and quality of service. This readily leads to

About the Background Paper Series

Published by:

The Internet Society
1775 Wiehle Avenue, Suite
102
Reston, Virginia 20190 USA
Tel: +1 703 326 9880
Fax: +1 703 326 9881

4, rue des Falaises
CH-1205 Geneva
Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445

Email: info@isoc.org

Web: <http://www.isoc.org/>

Series Editor: Martin Kupres

Copyright © Internet Society
2003.

All rights reserved.

a vision of computing and storage, not as a capital asset, but as a service purchased on demand over the network from trusted service providers. The combination of Grid and Web Services outlined above is a powerful mechanism for realizing this vision. An increasing number of people believe that this will be the biggest change in the nature of the IT industry for many years.

However, Grid computing will only be adopted by commercial users if they are confident that their data and privacy can be adequately protected and that the Grid will be at least as scaleable, robust and reliable as their own in-house IT systems. Thus, new Internet technologies and standards such as IPv6 and IPsec take on even greater importance. Needless to say, users of the Grid want easy, affordable, ubiquitous, broadband access to the Internet.

Grids raise a number of other issues of societal importance: research investment; usage for education, government, and regional development; liability and antitrust questions; intellectual property; universal access; and taxes and tariffs. Lack of space prevents detailed discussion of these issues in this briefing. Despite these complexities, which the community is starting to address, the various open-source and proprietary implementations of Grid computing allow commercial enterprises, like the scientific community, to adopt the technology for their in-house applications today.

ISOC Position

ISOC welcomes the emergence of this major new trend in network utilization, originating like so much else from the research community. However, if successful, Open Grid Services will further increase the dependence of large sections of the economy on the Internet and its technology and standards. It is essential that we strive to improve the quality and security of standards (whether developed in the IETF, the Global Grid Forum, or elsewhere). In the international context, it is essential that Grid technology can be deployed at large scale and across frontiers. Thus, three of ISOC's constant concerns - the general deployment of IPv6 and of better security, and the avoidance of restrictive regulation - are only reinforced by the rise of Grid computing. Since the widespread use of Grids depends on widespread access to broadband services, ISOC should redouble its efforts to encourage government policies that spur investment and competition in the telecommunications sector, particularly in the provision of wireless and broadband Internet services.

Acknowledgment

Michael R. Nelson made important contributions to this briefing.