Enterprise Grid Computing: State-of-the-Art

By Rajkumar Buyya, Ph.D., and Krishna Nadiminti

The term "grid" as used today means many different things to different people. It's often used to refer to various forms of distributed systems, such as cluster-based systems, Point-to-Point (P2P) networks, wide-area distributed storage solutions, and the like. Numerous large companies add to the confusion by liberally using the term grid while describing their products and services. So, it has now become such a common industry >

buzzword that the actual meaning needs to be inferred from the context.

Let's provide yet another definition, one that encompasses several existing definitions and describes some basic attributes of a grid:

"Grid is a type of parallel and distributed system that enables the sharing, selection, and aggregation of geographically distributed, 'autonomous' resources dynamically at run-time depending on their availability, capability, performance, cost, and users' quality-of-service requirements." (Source: Grid Computing Info Centre: Frequently Asked Questions, accessible at www.gridcomputing.com/gridfaq.html)

Given this definition, today's distributed systems have a varying degree of grid-like characteristics. There are many systems developed and deployed for various purposes and myriad names have emerged to describe these: compute grid, data/storage grid, campus grid, enterprise grid, global grid, knowledge grid, sensor grid, cluster grid, PC grid, commodity/utility grid, and so on. Ian Foster has written an interesting article titled "What Is the Grid? A Three Point Checklist," that describes some characteristics of a grid system. (You can access this article at www.gridtoday.com/02/0722/100136.html.)

This article lists the benefits of grid computing for the enterprise and describes a grid-oriented open source project with a compelling service-oriented framework.

Grid Benefits and Challenges

In a typical Small or Medium-Size Enterprise (SME), there are many resources that are generally underutilized for long periods. A resource in this context means any entity that could be used to fulfill any user requirement; this includes compute power (CPU), data storage, applications, and services. An enterprise grid can be loosely defined as a distributed system that aims to dynamically aggregate and coordinate various resources across the enterprise and improve their utilization for greater productivity.

Grid computing technology provides enterprises an effective solution for aggregating distributed resources and prioritizing allocation of resources to different users, projects, and applications based on their Quality of Service (QoS) requirements. These benefits ultimately result in huge cost savings for the business.

There are various applications of grid computing. Many commercial compute-intensive applications, such as drug discovery, clinical modeling, simulation, investment and credit-risk analysis, large-scale document processing, and data-intensive applications that involve aggregation and management of distributed data storage centers, can vastly benefit from the performance enhancements and resource aggregation capabilities that can accrue from the use of grid technologies.

However, grids today don't address all the issues important to the enterprise. That's because they were born in academic communities where such issues aren't a high priority. There are some important distinctions between the types of grids used in research communities and those that can be used in an enterprise or commercial environment. Figure 1 outlines the characteristics that differentiate an enterprise grid from a research-oriented grid. The stars reflect the importance/desirability of the attribute to each type of grid.

	ENTERPRISE GRID SYSTEMS	NON-COMMERCIAL GRIDS
Criticality of efficient and optimal resource usage	★★★★★	★★★★★
Sharing of inter-organizational resources	★★★	★★★★★
Authentication and authorization	★★★★★	★★★
Security of stored data and programs	★★★★★	★★★
Secure communication	★★★★★	★★★
Centralized / semi-centralized control	★★★	
License Management issues	★★★★★	★★★
Auditing	★★★★★	★★★
Quality of Service (QoS) and Service Level Agreements (SLAs)	★★★★★	★★★
Economy-based & Service-Oriented Architecture (to support QoS)	★★★★★	★★★
Interoperability of different grids (and, hence, the basis of open standards)	★★★★★	★★★★★
Support for transactions	★★★★★	★

Note: A maximum of five stars means the particular feature/attribute is of utmost importance. Absence of a star means the attribute isn't required/desirable.

Figure 1: Comparison of Features/Characteristics of a Commercial and Non-Commercial Distributed Systems

The Current State of the Enterprise Grid

Grid technology is evolving to provide solutions that more fully address enterprise requirements. The technology is rapidly moving from academia and scientific research and applications toward mainstream enterprise applications with a special emphasis on Service-Oriented Architectures (SOAs) and utility computing. The enterprise grid currently includes a range of applications that use data centers and application clusters to distribute workloads of applications such as:

- Accounts receivable
- Investment portfolio risk analysis
- Pricing securities in the finance and insurance sector
- Finding solutions to bottlenecks in product design and development cycles in the manufacturing sector
- Drug discovery in the pharmaceutical sector
- Digital media creation, rendering, and distribution management.

While most early-adopters are still running batch-oriented applications, the concepts of SOA and virtual organizations are already being used to explore the possibilities of running transactional and interactive applications on enterprise grids where the QoS is expected to be reliable, especially when bound by Service Level Agreements (SLAs).

Investment in enterprise grids is expected to grow manifold in the next five years as more companies come up with value-added

services, according to a *Network World* article titled "Grid Taking Shape in Enterprise Nets" (accessible at www.networkworld.com/news/2005/101005-grid.html). Various major companies are already offering a range of services such as:

- IBM's "Grid and Grow," (described at www-1.ibm.com/grid/) includes IBM's grid hardware, operating systems, schedulers, servic-

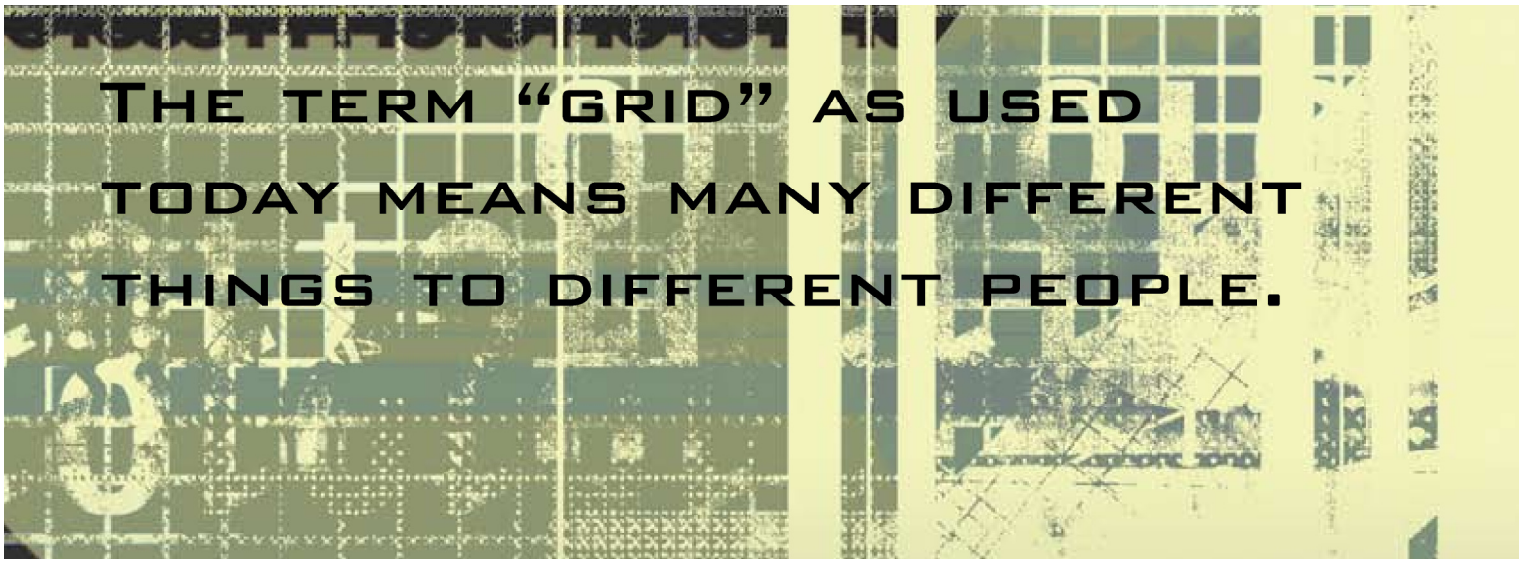

es and client training, and is intended to give businesses a competitive edge by using available resources more efficiently.
- Oracle's grid computing solution (described at www.oracle.com/technology/tech/grid/index.html) lets businesses standardize on modular servers and storage; consolidate servers and storage with Oracle Database (10g) and Real Application Clusters; and automate daily management tasks.
- Sun Microsystems' "Grid Utility Computing" (described at www.sun.com/servers/grid/) is a pay-per-use service that lets users dynamically provision compute power, depending on application requirements. Sun provides access to a standardized grid computing infrastructure that lets you offload your compute-intensive workloads with minimal risk and no capital investment.
- HP is delivering grid-based storage products today that are built according to their "StorageWorks" architecture (described at http://h71028.www7.hp.com/enterprise/cache/125369-0-0-0-121.html). These products either use early versions of smart cell technology or exemplify other design attributes of the architecture, such as single system image management.

The Enterprise Grid Alliance (EGA) is an open, non-profit, vendor-neutral consortium formed to develop enterprise grid solutions and accelerate the deployment of grid computing in enterprises. It consists of more than 30 members, including grid users, vendors and solution providers such as IBM, Oracle, Sun, Intel, HP, DataSynapse, Univa, and Dell. The EGA, as described at www.gridalliance.org, aims to encourage and accelerate movement to an open grid environment through interoperability solutions. It will work on grid computing standards by endorsing and supporting existing specifications, assembling and profiling component specifications, and defining new specifications where needed.

Open Source Software (OSS) is involved in a big way in the development of enterprise grids. Many grid solutions (including IBM's grid service offerings for the enterprise) are currently based on the open source Globus toolkit developed by the Globus Alliance, Argonne National Laboratory, and the University of Chicago and described at www.globus.org. It's a set of software services and libraries for resource monitoring, discovery, and management plus security and file management that facilitate construction of computational grids and grid-based applications, across corporate, institutional and geographic boundaries.

Globus offers grid middleware that mainly runs on the Unix-like platforms. Several open source grid projects have developed user-level middleware that work with the Globus toolkit. One such effort is the Gridbus project at the University of Melbourne, which developed the Grid Service Broker. As described at www.gridbus.org/broker/, the grid service broker supports creation, scheduling, and deployment of computational or data grid applications (including work flows) on enterprise and global grids.

Another open source grid initiative from the Gridbus project is the Alchemi enterprise grid-computing framework that harnesses the power of a network of computers running Windows.

Alchemi: An OS Enterprise Grid Computing Framework

Alchemi, as described at www.alchemi.net, is an open source, .NET-based enterprise grid computing framework developed by researchers at the GRIDS lab, in the Computer Science and Software Engineering Department at the University of Melbourne, Australia. It lets you painlessly aggregate the computing power of networked machines into a virtual supercomputer and develop applications to run on the grid with no additional investment and no discernible impact to users. It's been designed to be easy to use without sacrificing power and flexibility. It supports the Microsoft Windows operating system, which is seen as a key factor in industry adoption of grid computing technology, since more than 90 percent of machines worldwide run variants of Windows.

The main features offered by the Alchemi framework are:

- Virtualization of compute resources across

the LAN/Internet
- Ease of deployment and management
- Object-oriented "grid thread" programming model for grid application development
- File-based "grid job" model for grid-enabling legacy applications
- Web Services interface for interoperability with other grid middleware.

Figure 2 shows the Alchemi architecture, which has three types of components:

- The Manager
- The Executor
- The User Application.

The Manager node is a computer with the Alchemi Manager component installed. Its main function is to service user requests for application distribution. The Manager receives a user request, authenticates it, and distributes the workload across the various Executors that are connected to it. The Executor node is the one that actually performs the computation. Alchemi uses role-based security to authenticate users and authorize execution.

A simple grid is created by installing Executors on each machine that's to be part of the grid and linking them to a central Manager component. The Windows installer setup that comes with the Alchemi distribution and minimal configuration makes it easy to set up a grid.

An Executor can be configured to be dedicated (meaning the Manager initiates application execution directly) or non-dedicated (meaning that the execution is initiated by the Executor). Non-dedicated Executors can work through firewalls and Network Address Translation (NAT) servers since there's only one-way communication between the Executor and Manager. Dedicated Executors are more suited to an intranet environment and non-dedicated Executors are more suited to the Internet environment.

Users can develop, execute and monitor grid applications using the .NET Application Program Interface (API) and tools that are part of the Alchemi Software Developer's Kit (SDK). Alchemi offers a powerful grid thread programming model that makes it easy to develop grid applications and a grid job model for grid-enabling legacy or non-.NET applications.

An optional component is the Cross Platform Manager Web Service that offers interoperability with custom non-.NET grid middleware. Alchemi also comes with a Java API that can be used to develop Java-based clients that need to harness the computing power of an Alchemi grid.

Alchemi is widely used for a variety of applications. It has been used for teaching and setting up test grids and also some serious applications in the commercial world. Some Alchemi-based industrial applications and projects include:

- Large-scale document processing (Tier Technologies, U.S.)
- Natural resource modeling (CSIRO, Australia)
- Asynchronous Excel tasks using Managed XML Linking Language (XLL) (stochastix GmbH, Germany)
- Detection of patterns of transcription factors in mammalian genes (The Friedrich Miescher Institute (FMI) for Biomedical Research, Switzerland)
- Finding the location of a high-frequency radio transmitter using Secure Sockets Layer (SSL) technology (Correlation Systems Ltd., Israel).

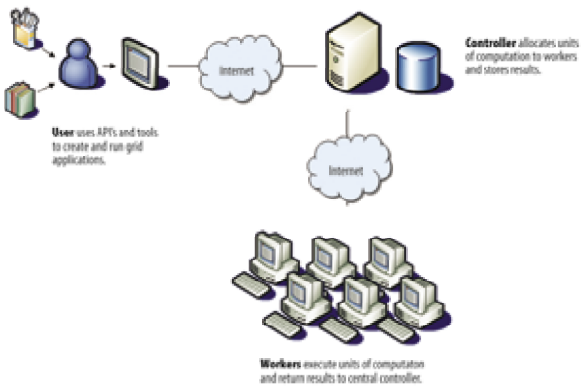

Figure 2: A Simple Alchemi Grid

A Peek Into the Future

The enterprise grid is still in its nascent stages in terms of development and industrywide adoption, but is poised for rapid growth. However, there are issues that preclude the big revolution that the grid promises to bring to IT. Some problems to solve include security, development and wide adoption of standards for representing and executing applications and workflows, resource description, monitoring and management, dynamic service composition and aggregation. There are also issues relating to managing data, intellectual property, developing new software licensing models and their enforcement, representing QoS and formulating and enforcing SLAs that are especially important in a commercial environment. Considerable research is under way in these areas and the standards are constantly evolving. Finally, before grid computing becomes ubiquitous, a sustainable business model has to be developed so all parties obtain value from adopting grid technologies.

It's widely believed that the grid of the future will be based on SOA and software and hardware will be available as a utility with demand and supply regulated by the concept of an economic market—just like it works for any other utility such as electricity, telephone, and water.

Rajkumar Buyya, Ph.D., is director of the Grid Computing and Distributed Systems (GRIDS) Laboratory within the Department of Computer Science and Software Engineering at the University of Melbourne, Australia. He has pioneered the economic paradigm for service-oriented grid computing and demonstrated its utility through his contribution to conceptualization, design and development of cluster and grid technologies, such as Gridbus, that powers many emerging eScience and eBusiness applications.
e-Mail: raj@cs.mu.OZ.AU

Krishna Nadiminti is a research programmer for the GRIDS laboratory, Department of Computer Science and Software Engineering, the University of Melbourne, Australia.
e-Mail: kna@cs.mu.OZ.AU