

Preface

Grid computing is currently an area of intense activity involving a large community of researchers, developers, and users. Hundreds of experimental and production grids are in operation throughout the world. The development of grid computing has followed a path similar to the development of the World Wide Web, which started as a technology for scientific collaboration but was later adopted for use by a multitude of industries and businesses. Recent standards activities driven by the Global Grid Forum (GGF) have led to the definition of the Open Grid Services Architecture (OGSA), which has met with wide acceptance and is being implemented in commercial grid products.

This issue of the *IBM Systems Journal* starts with three papers of general interest: a tutorial on grid computing, a paper introducing intraGrid, an experimental grid deployed on the IBM intranet, and a paper describing the vision of grid computing that drives the current and future development of grid-oriented features in IBM products. The next four papers deal with various aspects of the grid infrastructure: security, namespaces, logging, and aggregation of services. Special thanks are due to Marcos Novaes for his help with the planning and the production of papers devoted to grid computing. A nontopical paper on MyMED, a database system for biomedical research, concludes this issue.

In their paper “Evolution of grid computing architecture and grid adoption models,” Joseph, Ernest, and Fellenstein discuss the motivation for standardization in grid computing and the relevance of standards to the IBM on demand approach. They also discuss the recent merging of Web services and grid technology and its impact on grid architecture and grid adoption models.

The *intraGrid* is an experimental grid based on the Globus Toolkit** and deployed on the IBM intranet. In “Design and implementation of an enterprise grid,” Meliksetian et al. introduce intraGrid and describe the major challenges encountered during its design and implementation. They review the use of intraGrid by various teams of IBM researchers to date and describe their plans for future work, including migrating to an OGSA-based architecture.

Bourbonnais et al. present their vision of an information infrastructure for grid computing in their paper “Towards an information infrastructure for the grid.” They describe an infrastructure that supports virtualization of resources, achieves quality-of-service goals through policy-driven mechanisms, and is based on industry standards. The authors show how this vision of grid computing is reflected in the development work on current and future products.

Carpenter and Janson tackle the difficult and important problem of dealing with heterogeneous security systems when building grids that span multiple organizations. In particular, there is a need to specify security policies that allow certain users in one organization to access certain resources in another organization. In their paper “Abstract interdomain security assertions: A basis for extra-grid organizations,” they propose the use of an abstract form of interdomain security assertion that relies on appropriate universal identifiers to refer to users, resources, and attributes.

A uniform namespace with global scope and hierarchical ownership allows the sharing of file data among multiple organizations without compromising security or autonomy. Anderson et al. propose such a name service in their paper “Global

namespace for files.” This name service enables federated access to various file systems and can be extended to include databases and live data feeds.

In their paper “A Logger System based on Web services,” Horn et al. propose a virtual logging system for a distributed heterogeneous environment. The interfaces to such a logging system are expressed through the conventions of the recently released Web Services Resource Framework and provide an example of how this framework might be used to define meta-operating systems for grid computing.

According to Tan et al. a service domain maps a collection of comparable or related services to a single logical service. In large computing grids in which tens or hundreds of service instances are offered to customers, the use of the service domain could help cope with the complexity of that environment. In their paper “Service domains,” the authors describe an architecture for service domains, a reference implementation, and the application of these tools to the implementation of a large grid.

MEDLINE**, a citation database from the National Library of Medicine, is a frequently used resource in bioinformatics research. Access to this database is free but limited in terms of the number of daily requests allowed from the same computer. In their paper “MyMED: A database system for biomedical research on MEDLINE data,” Lewis et al. describe their experiences in implementing and using MyMED, a database system based on the IBM DB2 Universal Database* and the DB2* Extender products.

The next issue of the *Journal* is dedicated to business innovation.

Alex Birman, Associate Editor
John J. Ritsko, Editor-in-Chief

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of the University of Chicago or the United States National Library of Medicine.