

Grid Technologies & Applications: Architecture & Achievements

Ian Foster

Mathematics & Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA
Department of Computer Science, The University of Chicago, Chicago, IL 60637, USA
foster@mcs.anl.gov

Abstract

The 18 months since CHEP'2000 have seen significant advances in Grid computing, both within and outside high energy physics. While in early 2000, Grid computing was a novel concept that most CHEP attendees were being exposed to for the first time, we now see considerable consensus on Grid architecture, a solid and widely adopted technology base, major funding initiatives, a wide variety of projects developing applications and technologies, and major deployment projects aimed at creating robust Grid infrastructures. I provide a summary of major developments and trends, focusing on the Globus open source Grid software project and the GriPhyN data grid project.

1 Introduction

The term “Grid” was coined to denote a proposed distributed “cyberinfrastructure” for advanced science and engineering [7]. The term is now understood to refer to technologies and infrastructure that enable *coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations*. This sharing relates primarily to direct access to computers, software, data, and other resources, as is required by a range of collaborative problem-solving and resource-brokering strategies emerging in industry, science, and engineering. This sharing is, necessarily, highly controlled, with resource providers and consumers defining clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occurs. A set of individuals and/or institutions defined by such sharing rules form a *virtual organization* (VO).

Grid concepts are particularly relevant to high energy physics (HEP) due to the collaborative nature of HEP experiments and the increasing complexity of data analysis tasks, and hence a need for next-generation experiments to exploit large distributed collections of shared resources. We thus see considerable interest in the creation of so-called *HEP Data Grids*, i.e., Grid infrastructures, tools, and applications designed to enable distributed access to, and analysis of, large amounts of HEP data.

The broad significance of Grid concepts, in advanced scientific collaborations and in business, means that HEP is just one of a number of communities (albeit a particularly ambitious and important one) that are developing and/or driving Grid technologies. The resulting interrelationships make it important to understand the state of the art and likely future directions in this field. In this paper, I attempt to provide such an understanding.


In order to focus my presentation—and also to satisfy the CHEP organizers, who proposed the title “Overview of Globus and GriPhyN”—I approach this overview from the perspective of two major projects in which I am involved. These two projects address complementary aspects of the Grid problem, with Globus conducting research and development on foundation Grid issues such as security, policy, data management, and resource management, and GriPhyN tackling issues relating specifically to Data Grids, addressing in particular data grid architectural concepts and issues relating to virtual data. Both projects also produce significant software systems, the Globus Toolkit and GriPhyN Virtual Data Toolkit, respectively. Hence, a summary of Globus and GriPhyN activities and their partnerships provide a broad-brush, if not complete, view of Grid computing.

In the following, I review Grid architecture concepts; describe the current state of, and future plans for, the Globus Toolkit; review major Grid projects that are using Globus; and finally describe GriPhyN goals and initial achievements. Other sources provide a survey of Grid concepts and technologies [7] and a detailed picture of Grid architecture [8].

2 Grid Architecture

Grid technologies comprise protocols, services, and tools that address the challenges that arise when we seek to build scalable VOs. These technologies include security solutions that support management of credentials and policies when computations span multiple institutions; resource management protocols and services that support secure remote access to computing and data resources and the co-allocation of multiple resources; information query protocols and services that provide configuration and status information about resources, organizations, and services; and data management services that locate and transport datasets between storage systems and applications.

The figure illustrates a categorization that we have found useful when explaining the roles played by various Grid technologies. In the *Fabric*, we have the resources that we wish to share: computers, storage systems, data, catalogs, etc. The *Connectivity* layer provides communication and authentication services needed to communicate with these resources. *Resource* protocols (and, as in each layer, associated APIs) negotiate access to individual resources. *Collective* protocols, APIs, and services are concerned with coordinating the use of multiple resources, and finally application toolkits and applications themselves are defined in terms of services of these various kinds. Other papers present views on necessary components of a Grid architecture [8] and the additional services required within a Data Grid architecture [5].



3 The Globus Toolkit

The Globus Toolkit [6, 8] is a community-based, open-architecture, open-source set of services and software libraries that support Grids and Grid applications. The Toolkit includes software for security, information infrastructure, resource management, data management, communication, fault detection, and portability. It is packaged as a set of components that can be used either independently or together to develop useful Grid applications and programming tools. Globus Toolkit components include the Grid Security Infrastructure (GSI), which provides a single-sign-on, run-anywhere authentication service, with support for delegation of credentials to subcomputations, local control over authorization, and mapping from global to local user identities; the Grid Resource Access and Management (GRAM) protocol and service, which provides remote resource allocation and process creation, monitoring, and management services; the Metacomputing Directory Service (MDS) [4], an extensible Grid information service that provides a uniform framework for discovering and accessing system configuration and status information such as compute server configuration, network status, or the locations of replicated datasets. Data Grid-specific technologies include a replica catalog, GridFTP, a high-speed data movement protocol [1], and reliable replica management tools. For each of these components, the Toolkit both defines protocols and APIs and provides open source reference implementations in C and, in most cases, Java. A variety of higher-level services can be, and have been, implemented in terms of these basic components.

3.1 Recent Developments

2001 has seen significant enhancements to the Globus Toolkit in terms of functionality, robustness, flexibility, and performance. I summarize some major developments here. Some are already part of the main Globus release; the rest are slated for release in 2001.

- A major *packaging* thrust, with NCSA, is transforming the (conceptually) modular but (in practice) monolithic Globus software into a modular collection of distinct packages, hence enabling the creation of source and/or binary distributions for different purposes and architectures.

- Significant enhancements to the *Globus information service*, MDS, produced MDS-2.1, which in addition to a highly scalable distributed design features integrated GSI security and more efficient information source interface and indexing mechanisms.
- Work at NCSA produced the “MyProxy” *online credential repository*, an important contribution to improved usability of public key infrastructure authentication mechanisms
- A collaborative effort with the Condor [11] group at U.Wisconsin significantly enhanced the *GRAM protocol* and implementation to improve reliability and ability to recover from failures.
- The Condor group also produced *Condor-G* [9], a job management system for Grid computations that combines Condor job management with Globus security and remote job submission.
- Further work on *commodity Grid toolkits* (“CoG Kits”) at ANL, LBNL, and Rutgers produced improved Java interfaces to Grid services and early versions of Python and CORBA interfaces.
- Alpha releases of our *GridFTP* protocol implementation saw extensive use. GridFTP extends FTP with additional features aimed at high-performance transport (e.g., parallel data sources), flexibility (e.g., partial file transfers), and robustness (e.g., extended restart capabilities).
- A new *replica catalog* service supports registration and discovery of physical copies of logical files and collections. A new *GDMP* [12] data replication software version uses Globus services.

We also evolved the Globus development process, introducing an external review process for specifications and an alpha-testing program for new software. These changes are part of an ongoing process aimed at creating an effective, scalable open source platform for Grid software.

3.2 Current Developments

An aggressive R&D program at ANL, UChicago, ISI, and other partner institutions is continuing to enhance the capabilities of the Globus Toolkit. Here I mention a few developments underway.

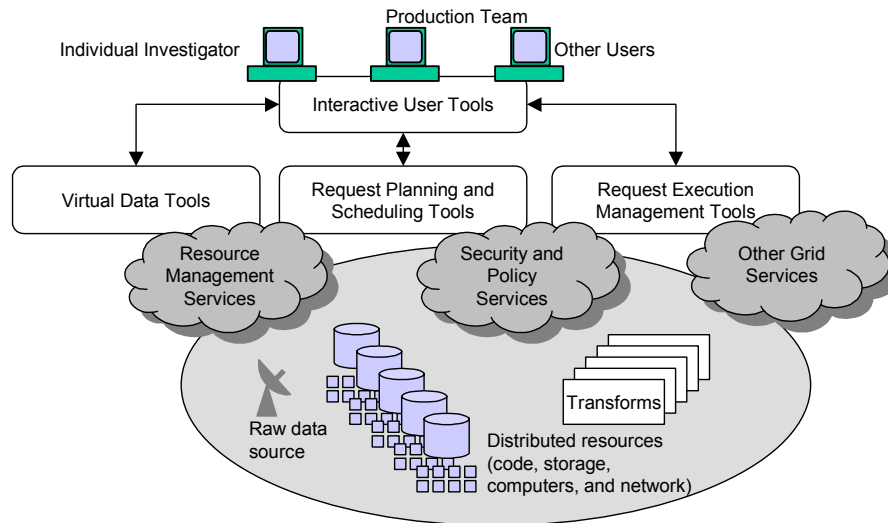
- A *Community Authorization Service* (CAS) will support authorization decisions within large VOs, in which it is desirable to be able to delegate policy decisions concerning who can do what from individual resources and users to community representatives.
- Other work in *security* aims at enhancing significantly the ease of use of security technologies, by for example further enhancements to the online credential repository concept.
- Continued refinements to the *MDS-2 architecture and implementation* will allow its use as an integrating framework for Grid monitoring.
- A next-generation *GRAM-2 protocol* will provide significant enhancements to the robustness and capabilities of GRAM, extending its range of application to other resource types.
- We continue to investigate application of Globus technologies within commodity computing, and expect to propose enhancements to the *Web Services framework* to support Grid applications.

4 The GriPhyN Project

GriPhyN is a large collaboration of information technology (IT) researchers and experimental physicists who aim to provide the IT advances required to enable Petabyte-scale data intensive science in the 21st century. Driving the project are unprecedented requirements for geographically dispersed extraction of complex scientific information from large collections of measured data. To meet these requirements, which arise initially from four partner physics experiments (ATLAS, CMS, LIGO and SDSS) but will also be fundamental to science and commerce in the 21st century, the GriPhyN team is pursuing IT advances centered on the creation of *Petascale Virtual Data Grids* (PVDG) that meet the data-intensive computational needs of a diverse community of thousands of scientists spread across the globe.

GriPhyN has adopted the concept of virtual data as a unifying theme for its investigations of Data Grid concepts and technologies. This term is used to refer to two related concepts: transparency with

respect to location as a means of improving access performance, with respect to speed and/or reliability, and transparency with respect to materialization, as a means of facilitating the definition, sharing, and use of data derivation mechanisms. In order to realize these concepts, GriPhyN conducts research into virtual data cataloging, execution planning, execution management, and performance analysis issues (see the figure). The results of this research, and other relevant technologies, are developed and integrated to form a Virtual Data Toolkit (VDT). Successive VDT releases are applied and evaluated in the context of the four partner experiments.



GriPhyN has been underway just over half a year and after an initial rampup/definition phase has established strong research, development, and application projects. Accomplishments to date include:

- The definition of a Data Grid Reference Architecture, defining requirements and proposing interfaces and implementations for the core components required in a Data Grid system.
- The development, in collaboration with experiments, of requirements documents for virtual data services in each of the partner experiments. (E.g., see [2, 10].)
- Successful demonstration of distributed analysis of CMS simulation data using systems at Caltech, U.Wisconsin, and NCSA, with Condor-G used as a computation management system.

Version 1 of the GriPhyN Virtual Data Toolkit, slated for October 2001, will provide packaged, documented, integrated, tested, and supported versions of a core set of Data Grid software, including GridFTP data transfer, replica catalog, Condor-G computation management, and Globus Toolkit security, information, and resource management services. Subsequent versions will provide expanded functionality (e.g., community authorization, scheduling, automated replica selection, request planning), scalability (e.g., number of data, number of resources), and performance.

5 Grid Deployments and Applications

The last year has seen major Grid projects mature and others emerge as a result of new funding initiatives within the U.S. and E.U. Momentum is also building within Asia-Pacific. The table provides a *far-from-complete* list of projects, selected to emphasize the breadth and depth of the endeavor. Some projects address the creation of Grid technology, some the application of Grid technology, and yet others the deployment of this technology to create “production” Grids. Some are still in the proposal review stage, most are underway. Some URLs were not yet operational at the time of writing. All have adopted Globus software as a significant technology component.

Name	URL + Funding Agency	Focus
Access Grid	www.mcs.anl.gov/FL/accessgrid ; DOE & NSF	Create and deploy group collaboration systems using commodity display technologies
DISCOM	www.cs.sandia.gov/discom DOE Defense Programs	Create an operational Grid providing access to resources at three U.S. DOE weapons laboratories
DOE Science Grid	www.sciencegrid.org DOE Office of Science	Create an operational Grid providing access to resources and applications at U.S. DOE science laboratories and partner universities
ESG	www.earthsystemgrid.org DOE Office of Science	Delivery and analysis of large climate model datasets for the climate research community.
EU DataGrid	www.eu-datagrid.org European Union	Create and apply an operational grid for applications in high energy physics, environmental science, bioinformatics.
EuroGrid & GRIP	www.eurogrid.org European Union	Create and apply technologies for remote access to supercomputer resources and simulation codes; in GRIP, integrate with Globus technologies
Fusion Collaboratory	DOE Office of Science	Create a national computational collaboratory for fusion research.
GridPP	www.gridpp.ac.uk U.K. eScience Program	Create and apply an operational grid for particle physics research
GRIDS Center	www.grids-center.org NSF	Integration, deployment, support of an NSF Middleware Infrastructure for research & education
Grid Physics Network	www.griphyn.org NSF	Technology R&D for data analysis in physics experiments: ATLAS, CMS, LIGO, SDSS
Information Power Grid	www.ipg.nasa.gov NASA	Create and apply a production Grid for aerosciences and other NASA missions
Intl Virtual Data Grid Laboratory	www.ivdgl.org NSF	Enable large-scale, international experimentation on Grid technologies & applications
NEESgrid	www.neesgrid.org NSF	Create and apply a production Grid for earthquake engineering
PPDG	www.ppdg.net DOE Office of Science	Create and apply production Grids for data analysis in high energy and nuclear physics experiments
TeraGrid	www.teragrid.org NSF	U.S. science infrastructure linking four major resource sites at 40 Gb/s
UK Grid Support Center	www.grid-support.ac.uk U.K. eScience Program	Support center for Grid projects within the U.K.

6 Project Coordination and iVDGL

Coordination among various Data Grid projects is vital if we are to avoid replicated effort, profit from hard-won experiences, and achieve interoperable international Grid infrastructures. Coordination is being pursued on several fronts. An international coordination committee involving the various HEP Data Grid projects (EU DataGrid, GriPhyN, PPDG) and representatives of associated physics experiments meets regularly to discuss approaches and opportunities for joint work. The Global Grid Forum (www.gridforum.org) supports information exchange and standards setting. The Globus community also works to encourage communication and common approaches.

Another proposed initiative would create and operate an International Virtual Data Grid Laboratory (iVDGL) [3] comprising computing and storage resources in the U.S., Europe, Japan, and Brazil—and ultimately other regions—linked by high-speed networks, and operated as a single system for interdisciplinary experimentation in Grid-enabled data-intensive scientific computing. We anticipate that the creation of this laboratory will not only drive the development, and transition to every day production use, of Petabyte-scale virtual data applications, but will help encourage the adoption of the protocol and policy standards needed for interoperability and large-scale operation.

Initiatives such as iVDGL are possible in part because of advances in national and international networking. National initiatives such as IWIRE, TeraGrid, and Geant will soon provide 10 Gb/s or more. Internationally, STAR-TAP (www.startap.net) in Chicago plays a valuable role as an international interconnection point for research networks. Now, StarLight (www.startap.net/starlight) is providing an interconnection point for optical networks, with already two groups (SURFnet from the Netherlands, and DATATAG from CERN) planning to connect to the U.S. at 2.5 Gb/s.

Acknowledgements

This work was performed by, and/or in collaboration with, many colleagues. I acknowledge, in particular, [Carl Kesselman](#), Globus Co-PI and GriPhyN co-investigator; [Paul Avery](#), GriPhyN Co-PI; and [Miron Livny](#), PPDG PI and GriPhyN co-investigator. My research is supported by the U.S. Department of Energy, the National Science Foundation's "GriPhyN" project, DARPA, and NASA.

References

We note that GriPhyN reports are accessible at www.griphyn.org.

1. Allcock, W., Bester, J., Bresnahan, J., Chervenak, A., Foster, I., Kesselman, C., Meder, S., Nefedova, V., Quesnel, D. and Tuecke, S., Secure, Efficient Data Transport and Replica Management for High-Performance Data-Intensive Computing, *Mass Storage Conference*, 2001.
2. Allen, B., Deelman, E., Kesselman, C., Lazzarini, A., Prince, T., Romano, J. and Williams, R. LIGO's Virtual Data Requirements. GriPhyN, 2000.
3. Avery, P., Foster, I., Gardner, R., Newman, H. and Szalay, A. An International Virtual-Data Grid Laboratory for Data Intensive Science. GriPhyN, 2001.
4. Czajkowski, K., Fitzgerald, S., Foster, I. and Kesselman, C., Grid Information Services for Distributed Resource Sharing. In *IEEE International Symposium on High Performance Distributed Computing*, (2001), IEEE Press
5. Foster, I. and Kesselman, C. A Data Grid Reference Architecture. GriPhyN, 2001.
6. Foster, I. and Kesselman, C. Globus: A Toolkit-Based Grid Architecture. In *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999, 259-278.
7. Foster, I. and Kesselman, C. (eds.). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1999.
8. Foster, I., Kesselman, C. and Tuecke, S. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *Intl. J. Supercomputer Applications*, (to appear). 2001.
<http://www.globus.org/research/papers/anatomy.pdf>.
9. Frey, J., Tannenbaum, T., Foster, I., Livny, M. and Tuecke, S., Condor-G: A Computation Management Agent for Multi-Institutional Grids. In *10th International Symposium on High Performance Distributed Computing*, (2001), IEEE Press
10. Holtman, K. CMS Data Grid System Overview and Requirements. GriPhyN, 2001.
11. Livny, M. High-Throughput Resource Management. In Foster, I. and Kesselman, C. eds. *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999, 311-337.
12. Stockinger, H., Samar, A., Allcock, W., Foster, I., Holtman, K. and Tierney, B., File and Object Replication in Data Grids. In *IEEE Intl. Symp. on High Performance Distributed Computing*, (2001), IEEE Press.