The European DataGrid Project

Robert J. Wilson*
Institut de Fisica d'Altes Energies, Barcelona, Spain
and Colorado State University, Fort Collins, Colorado, USA
(Dated: October 22, 2001)

We give a brief report on a European initiative to produce a production capable *data grid* within three years. The goal is to support the data access and computation needs of demonstration projects in High Energy Physics, Earth Observation Data, and the Biosciences. It is hoped that it will also provide a functional framework for the next generation of information technologies beyond the limited scope of these demonstration projects.

I. WHY GRIDS FOR HIGH ENERGY PHYSICS?

The 500 lb gorilla of particle physics computing requirements in the next decade will be the Large Hadron Collider at CERN. The physicists envision accumulating terabytes of data per day - petabytes per year. The resources needed to process and store this vast amount of information will far outstrip those that the CERN budget could provide. Even if the financing was available, it is not clear that the traditional centralization of computing resources would provide adequate access for the hundreds or even thousands of users spread over a score of countries.

A potential solution has emerged from the change in orientation of US meta-computing activity from interconnected super-computers, towards a more general concept of a computational grid. This is essentially a shift from high performance to high throughput computing. Following years of development among computer scientists, there has been a flurry of computational grid activity in high energy physics. Two large projects have been funded in the US: the Particle Physics Data Grid (PPDG), supported by the US Department of Energy, and the Grid Physics Network (GriPhyN) funded by the National Science Foundation. A similar initiative, the Information Processing Grid, is being pursued by NASA. There are also several large national projects in Europe. For example, the Grid Technology Evaluation project sponsored by the Italian particle physics institutes (INFN), and from the UK government there is a commitment to provide £100 million over three years for the implementation of a Prototype Grid infrastructure.

Each of these efforts is being shepherded by HEP groups. An example of the high-level of interest in Europe, and the perceived potential impact, is an article last October in the Spanish national newspaper, El Pais; they declared: "Los fisicos preparan la nueva 'web'", which translates to "Physicists build the new 'Web'".

II. THE GRID CONCEPT

The name *Grid* comes from that holy grail of democratic computing - data and computation access as easy and available as using electricity from the power grid. You plug into a socket in the wall and you have instant, effortless access to power – so long as you can pay for it.

The *de facto* bible on the subject is "The Grid: Blueprint for a New Computing Infrastructure," edited by Ian Foster (Argonne National Laboratory) and Carl Kesselman (University of Southern California's Information Sciences Institute). They begin the book by summarizing some of the key requirements for such a system: services on demand, high reliability, dynamic data/computing distribution, transparent access to multi-petabyte databases, and, a key part of the analogy to the power grid, hidden complexity of the infrastructure.

One can imagine innumerable uses for such a grid, some more obvious than others. One grand scenario envisions a governmental "national grid", as a kind of strategic computing reserve or for national collaborative projects (dubbed a "collaboratory"). Private grids in industry will provide seamless secure data exchange across national borders. It is also inevitable that the entertainment industry will find ways to make profitable use of virtual "shared spaces" – though it should be noted that proponents of this concept hope that educational uses of the technology won't get completely lost in the rush to profits.

Implementing these goals will require a substantial extension to current national and international network and computing systems, both hardware and software. For example, in addition to reliable high-speed access to

*Email: wilson@lamar.colostate.edu; supported on this project by a grant from the Spanish Ministry of Science and Technology

TABLE I: Grid Components Hierarchy

Applications					
Chemistry Biology Cosmology Environment High Energy Physics					
Application Toolkits					
Distributed Computing	Data Intensive Applications	Collaborative Applications	Remote Visualization Applications	Problem Solving Applications	Remote Instrumentation Applications
Grid Services (Middleware)					
Resource-independent and application-independent services: *authentification, authorisation, resource location, resource allocation, events, accounting, remote data access, information, policy, fault detection*					
Grid Fabric (Resources)					
Resource-specific implementation of basic services: *transport protocols, name servers, differentiated services, CPU schedulers, public key infrastructure, site accounting, directory services, OS bypass etc.*					

remote data, the most effective operation will require multiple copies of large databases dynamically generated by the grid itself. This data "replication" will need sophisticated management, such as tools for automated discovery of the "best" (closest) copy of data by grid clients. Optimal system performance also will need co-scheduling of distributed computing, storage and network resources. A task quite beyond current systems. A coarse hierarchy of grid services is summarized in Table I.

Among the most pressing of the middleware tasks are *authentification* and *authorization*. These are coupled but functionally independent aspects. The first addresses the basic security question of verifying the identity of the resource requester. A priority issue even for today's e-commerce infrastructure. Authorization is the act of determining which resources this requester is allowed to use. The challenge here is to do this on a global scale, such that these tasks can be performed once, then propagated securely through a globally distributed system of resources. This same statement applies to many of the other services essential for an effective data grid.

III. THE EUROPEAN DATAGRID PROJECT

The DataGrid project was commissioned by the European Union, in part, to give a central focus to several national initiatives. This three year, 10 million euro, project officially started on the auspicious date of 01/01/01 after an unusually rapid 9 month review and approval process.

The principal goals of the project are to develop the software to provide basic grid functionality and associated management tools for a large scale testbed for demonstration projects in three specific areas of science. The primary project milestones and associated deliverables are shown in Table II.

The driving force of the project is the need for production quality mock data, simulation and analysis capability for each of the Large Hadron Collider experiments (ATLAS, CMS and LHCb). The other demonstration areas, which began receiving attention somewhat after HEP, are in the fields of Earth observation data (e.g. satellite images) and the biosciences, principally genome data access and analysis. In addition to these demonstration projects, an explicit component of the project is to encourage academic/industry synergy through such venues as the *GRID Forum* and the *Industry and Research Forum*.

The main partners in the HEP part of the project are CERN, INFN(Italy), CNRS(France), PPARC(UK), NIKHEF(Netherlands). The European Space Agency has taken the lead in the Earth Observation task and KNMI(Netherlands) is leading the biology and medicine tasks. In addition to the major partners, there are associated partners from the Czech Republic, Finland, Germany, Hungary, Spain and Sweden. A relatively recent important development is the establishment of formal collaboration with some of the US grid projects.

Formal industrial participation is limited to a few corporations such as CS SI in France, DataMat in Italy, and IBM in the UK. However, the Industry and Research Forum includes representatives from Denmark, Greece, Israel, Japan, Norway, Poland, Portugal, Russia, and Switzerland, demonstrating at least a healthy curiosity about the project across a wider arena. Even with the existing breadth of participation, the EU has decided to encourage another multinational collaboration to generate a second related proposal (called "Cross-Action Grid"). The goal there is to encourage the application of the basic grid infrastructure to a broader range of disciplines.

The DataGrid project is divided into twelve work packages described only by concise goals, with few implementation details: Grid Workload Management (WP1), Grid Data Management (WP2), Grid Monitoring services (WP3), Fabric Management (WP4), Mass Storage Management (WP5), Integration Testbed (WP6), Network Services (WP7), HEP Applications (WP8), Earth Observation Science Applications (WP9), Biology

TABLE II: Project Milestones

Milestone	Deliverable	Time (month)
Testbed topology	Widely deployed grid of basic services	6
Tier 1, 2, 3 computing model experimentation	Testbed for traditional applications	12
Grid prototypes	Scheduled client/server applications testbed	18
Scalability tests	Grid fabric for distributed Tier 1	24
	Chaotic client/server applications testbed	30
Validation tests	Grid fabric for distributed all Tier hierarchy	36

Applications (WP10), Dissemination (WP11), Project Management (WP12). The project manager is Fabrizio Gagliardi of CERN.

IV. STATUS

Though many project participants had been working for many months before the official starting data, activity kicked into a noticeably higher gear in January, 2001. In some cases the first challenge was to translate the goals of each work package into specific tasks and deliverables.

The first full collaboration meeting occurred in March, 2001, in Amsterdam, in association with the Global Grid Forum meeting. By that time many new DataGrid specific personnel had been hired and hardware for the first phase tests (informally called *Testbed-0*) had been purchased in many of the participating countries. Of more formal significance was the presentation of the first draft of DataGrid architecture. Though it was not a great surprise, the architecture group presented an initial Testbed implementation, scheduled for month-9 of the project (*Testbed-1*), based on the Globus toolkit. However, they were quite clear that the option of replacing individual components was open, perhaps expected, in later versions. Nevertheless, a subsequent announcement of a cooperative agreement with the Globus team, to include synchronization of the some aspects of the Globus release schedule and the DataGrid milestone requirements, may go a long way towards maintaining a substantial Globus presence in the final product.

Despite a seemingly haphazard first few months, by early summer 2001, the basic Grid software components had been installed at more than 50 sites across Europe. Though inter-operation was limited to just a few regional sites, this key aspect is expected to develop rapidly after the release of a new Testbed installation package early in the fall.

V. CONCLUSION

The Grid is a useful metaphor for large-scale HEP computing needs such as the LHC experiments, and even for smaller scale collaborations. The middleware, fabric and interfaces should be general enough to accommodate a broad range of scientific, industrial and commercial applications. However, it appears that many of the grid infrastructure components still need basic R&D, so a "production capable" testbed is a bold move by the European Union. They hope that by making that step now, they can leap-frog over the ground breaking work done in the US to a usable product in a scant three years.

The world's major funding agencies have shown that they are prepared to fund large testbeds in the USA, Europe and Japan, so this is an excellent opportunity for HEP again to be a leader in a world-changing technology transfer. If successful, it could indeed develop to be the next generation of internet computing.

In the portentous words of the project monitor from the EU, "this project has high priority and visible, it cannot fail." The hundreds of LHC physicists who will depend on it could not agree more.

Acknowledgments

Much of the material for this report came from presentations by the DataGrid project manager Fabrizio Gagliardi. I would like to thank the director and staff of the Institut de Fisica d'Altes Energies for providing me with the opportunity to participate in this project.