

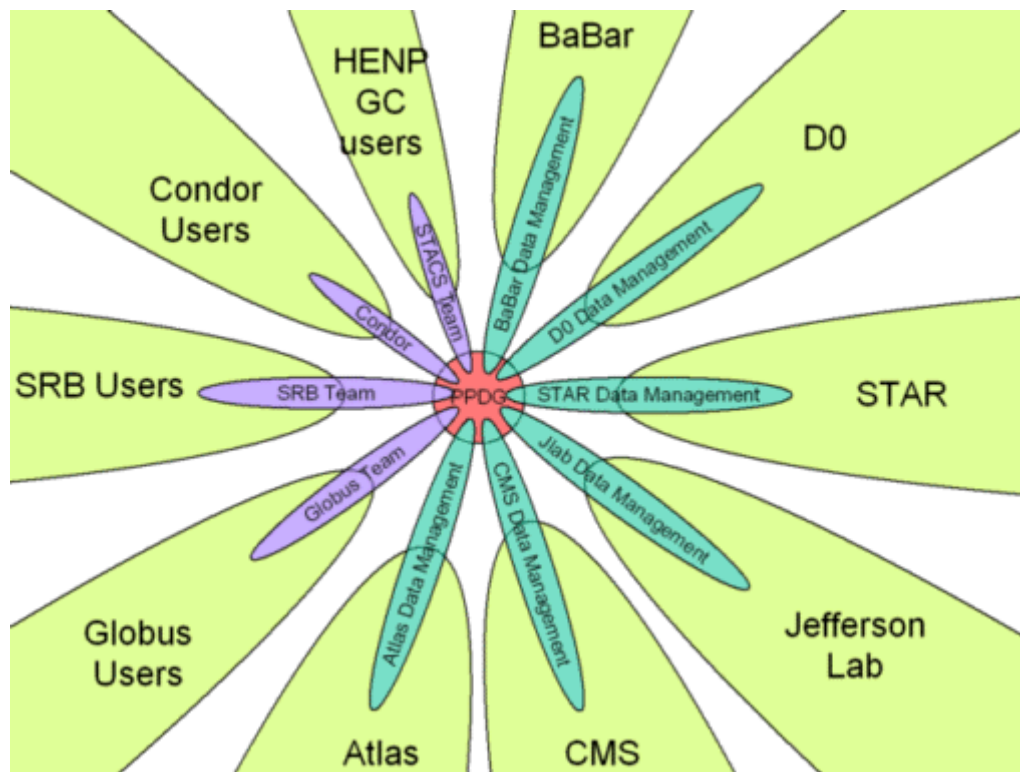
# US Grid Projects: PPDG and iVDGL

Richard P. Mount, SLAC

On behalf of the PPDG and iVDGL Collaborations

## The Grid Coordination Challenge

The HENP vision of the Grid<sup>1</sup> includes fascinating computer science challenges, many of which are evoked by the “virtual data” concept, as well as the challenge of bringing the easy use of these new paradigms to the thousands of experimental scientists. Astrophysics, whose data handling needs are now approaching those of HENP experiments, are joining the HENP experiments in Grid projects. Arguably, the coordination of the ensemble of Grid developments relevant to these experiments is an even bigger challenge and is totally vital to success. The coordination challenge also exists within individual projects, as exemplified by the figure below that attempts to relate the project-funded PPDG<sup>2</sup> effort to the larger teams of developers and users that must be coordinated for total project success.



Extending the picture of the coordination challenge to include, EU DataGrid<sup>3</sup>, GriPhyN<sup>4</sup>, GridPP<sup>5</sup> iVDGL<sup>6</sup> etc. is beyond my graphical capabilities. Below I give separate accounts of the now-funded PPDG project and the iVDGL proposal. In the future separate accounts of Grid-project success or plans will be of less and less value.

## **The Particle Physics Data Grid (PPDG)**

The Particle Physics Data Grid collaboration brings together as collaborating peers 6 experiments at different phases of their lifecycle and the recognized Grid middleware teams of Globus, Condor, SRB and LBL-STACS. PPDG will develop, evaluate and deliver vitally needed Grid-enabled tools for data-intensive collaboration in particle and nuclear physics. Novel mechanisms and policies will be vertically integrated with Grid middleware and experiment specific applications and computing resources to form effective end-to-end capabilities. A three-year program is now underway, taking advantage of the strong driving force provided by now-running physics experiments together with recent advances in Grid middleware. Our goals and plans are guided by the immediate, medium-term and longer-term needs and perspectives of the LHC experiments ATLAS<sup>7</sup> and CMS<sup>8</sup> that will run for at least a decade from late 2005 and by the research and development agenda of other Grid-oriented efforts. We exploit the immediate needs of running experiments – BaBar<sup>9</sup>, D0<sup>10</sup>, STAR<sup>11</sup> and Jlab<sup>12</sup> experiments – to stress-test both concepts and software in return for significant medium-term benefits. For these “mid-life” experiments the new Grid services must be introduced and deployed without destabilizing the existing data handling systems. While this imposes constraints on our developments, it also ensures rapid programmatic testing under real production conditions.

PPDG is actively involved in establishing the necessary coordination between potentially complementary data-grid initiatives in the US, Europe and beyond.

The BaBar experiment faces the challenge of data volumes and analysis needs planned to grow by more than a factor 20 by 2005. During 2001, the CNRS-funded computer center at CCIN2P3 Lyon, France will join SLAC in contributing data analysis facilities to the fabric of the collaboration. The STAR experiment at RHIC has already acquired its first data and has identified Grid services as the most effective way to couple the facilities at Brookhaven with its second major center for data analysis at LBNL. An important component of the D0 fabric is the SAM<sup>13</sup> distributed data management system at Fermilab, already deployed in several major US and international sites. The D0 SAM data grid moves the Monte Carlo data to the central Fermilab mass storage systems as well as replicates acquired data for analysis. The LHC collaborations have identified data-intensive laboratories as a vital component of their plan to analyze tens of petabytes of data in the second half of this decade. Both US Atlas and CMS are actively developing distributed simulation data production systems which will be used extensively in the next few years for detector and physics studies

### **PPDG methodology**

The PPDG Collaboration intends to take advantage of the existing requirements analyses performed by the physics collaborations and to leverage existing software developed by members of the collaboration and others. Our research will build on the frameworks and technology that have been developed by the PPDG team over the last two years. This effort was funded by DOE Mathematical Information and Computer Sciences and by High Energy and Nuclear Physics. In addition, the HENP experiments participating in PPDG necessarily have access to very significant national and international computational, storage and network facilities that will be exploited as the Grid fabric for PPDG with the full support of the physics collaborations.

The challenge of creating the vertically integrated technology and software needed to drive a data-intensive laboratory for particle and nuclear physics is daunting. Success is rendered likely by the existence of several efforts with complementary focus: in the US the NSF-funded GriPhyN collaboration is focusing on the longer-term computer science issues, notably the support of “virtual data”, and the complexities of a multi-tiered computing and data system. In Europe, the European Union-funded DataGrid will address the medium to long term goals of its collaborators including LHC-focused particle physics, Earth observation sciences and biology. The PPDG team will focus on

providing a practical set of Grid-enabled tools that meet the deployment schedule of the HENP experiments. It will make use of existing technologies and tools, including those from GriPhyN and the European DataGrid as they become available, to the maximum extent, developing, on the CS side, those technologies needed to deliver vertically integrated services to the end user. Areas of concentration for PPDG will be the sharing of analysis activities, the standardization of emerging Grid software components, status monitoring, distributed data management among major computing facilities and Web-based user tools for large-scale distributed data exploration and analysis. Many key personnel in these projects have a long history of collaboration on physics, computing, data management and network development projects, and all are committed to ensure overall coordination.

Explicit management and coordination structures are in place to ensure that the work plan addresses the goals and meets the needs of both the physics and CS teams. While, as expected, each of the two teams has a different perspective on the challenges we face, they agree on the goals and on what it will take to meet them.

The PPDG work plan will focus on several distinct areas as follows:

1. Deployment, and where necessary enhancement or development of distributed data management tools:
  - Distributed file catalog and web browser-based file and database exploration toolset
  - Data transfer tools and services
  - Storage management tools
  - Resource discovery and management utilities
2. Instrumentation needed to diagnose and correct performance and reliability problems
3. Deployment of distributed data services (based on the above components) for a limited number of key sites per physics collaboration:
  - Near-production services between already established centers over 'normal' networks (currently OC12 or less);
  - Close collaboration with projects developing "Envelope-pushing" services over NTON<sup>14</sup> and other high-speed research testbeds (currently OC48 or more).
4. Exploratory work with limited deployment for advanced (i.e. difficult) services:
  - Data signature definition (information necessary to re-create derived data) and catalog
  - Transparent (location and medium independent) file access
  - Distributed authorization in environments with varied local requirements and policies
  - Cost estimation for replication and transfer
  - Automated resource management and optimization.

The above work breakdown reflects the viewpoint of physicists. From a CS viewpoint, the research and development agenda of this effort will map principally on to issues related to the Grid fabric layer and within or close to the application layer. We intend to leverage existing intermediate Grid layers provided and supported by other projects to interface the application layer with the fabric. We are working in close collaboration with the teams that have created these software systems and seeking their advice and support in building fully integrated Grid-enabled capabilities. The principal CS work areas, forming an integral part of the above breakdown are:

1. Obtaining, collecting and managing status information on resources and applications, (managing these data will be closely linked to work on the replica catalog)
2. Storage management services in a Grid environment
3. Reliable, efficient and fault-tolerant data movement
4. Job description languages and reliable job control infrastructure for Grid resources.

Tools provided by the CS team are being adapted to meet local/specific requirements and will be deployed by members of the Physics team. Each experiment is responsible for its applications and resources and will operate a largely independent, vertically integrated Grid, using as far as possible standardized components and often sharing network infrastructures. The schedule and deliverables of the CS team are being coordinated with the "milestones" of the experiments. The "Grid Reference Architecture" developed by GriPhyN will guide the work. As GriPhyN starts to deliver Virtual Data technology that is based on its long-term research agenda, it will be included in our vertically integrated Grid services.

PPDG sub-projects are organized as collaborations between a single experiment and computer science group. The goal is to stimulate implementation and production employment to bring Grid Services early on to provide demonstrable benefit to the many physicists doing analysis and data processing work. Experience from the production and performance issues will be fed into subsequent sub-projects and versions of software as it is adapted and reused. The goal is then to sponsor general services and software packages that are can be trusted to perform.

The PPDG emphasis on end-to-end applications makes it a natural forum for collaborating with other Grid projects on national and international test beds - such as the iVDGL proposal below.

### **PPDG Anticipated results**

Results and deliverables will be produced in three areas:

- a. Data-intensive collaborative tools and services of lasting value to particle and nuclear physics experiments. Support responsibilities for this technology will be transferred to the experiments and to a dedicated US support team for which funding has been requested within the DOE High-Energy Physics program.
- b. Advances in computer science and software technology specifically needed to meet the demanding needs of a data-intensive collaborative. The validation and hardening of ideas currently embodied in early Grid services and proof-of-concept prototypes is considered a most important component of these advances.
- c. Advances in the understanding of the infrastructure and architectural options for long-term development of data-intensive Grid and collaborative services. The involvement of key scientists from long-term Grid projects will ensure that practical experience gained from this collaborative pilot can become an integral part of forward-looking architectural planning.

## iVDGL

The proposed international Virtual-Data Grid Laboratory (iVDGL) will have unprecedented scale and scope, comprising heterogeneous computing and storage resources in the U.S., Europe—and ultimately other regions—linked by high-speed networks, and operated as a single system for the purposes of interdisciplinary experimentation in Grid-enabled data-intensive scientific computing.

Our goal in establishing this laboratory is to drive the development, and transition to every day production use, of Petabyte-scale virtual data applications required by frontier computationally oriented science. In so doing, we seize the opportunity presented by a convergence of rapid advances in networking, information technology, Data Grid software tools, and application sciences, as well as substantial investments in data-intensive science now underway in the U.S., Europe, and Asia. We expect experiments conducted in this unique international laboratory to influence the future of scientific investigation by bringing into practice new modes of transparent access to information in a wide range of disciplines, including high-energy and nuclear physics, gravitational wave research, astronomy, astrophysics, earth observations, and bioinformatics. iVDGL experiments will also provide computer scientists developing data grid technology with invaluable experience and insight, therefore influencing the future of data grids themselves. A significant additional benefit of this facility is that it will empower a set of universities who normally have little access to top tier facilities and state of the art software systems, hence bringing the methods and results of international scientific enterprises to a diverse, world-wide audience.

Data Grid technologies embody entirely new approaches to the analysis of large data collections, in which the resources of an entire scientific community are brought to bear on the analysis and discovery process, and data products are made available to all community members, regardless of location. Large interdisciplinary efforts such as the NSF-funded GriPhyN and European Union (EU) DataGrid projects are engaged in the research and development of the basic technologies required to create working data grids. What is missing is (1) the deployment, evaluation, and optimization of these technologies on a production scale, and (2) the integration of these technologies into production applications. These two missing pieces are hindering the development of large-scale data-grid applications application design methodologies, thereby slowing the transition of data grid technology from proof of concept to full adoption by the scientific community. In this project we aim to establish a laboratory that will enable us to overcome these obstacles to progress.

Laboratory users will include international scientific collaborations such as the Laser Interferometer Gravitational-wave Observatory (LIGO<sup>15</sup>), the ATLAS and CMS detectors at the Large Hadron Collider (LHC) at CERN, the Sloan Digital Sky Survey (SDSS<sup>16</sup>), and the proposed National Virtual Observatory (NVO<sup>17</sup>); application groups affiliated with the NSF PACIs<sup>18</sup> and EU projects; outreach activities; and Grid technology research efforts. The laboratory itself will be created by deploying a carefully crafted data grid technology base across an international set of sites, each of which provides substantial computing and storage capability accessible via iVDGL software. The 20+ sites, of varying sizes, will include U.S. sites put in place specifically for the laboratory; sites contributed by EU, Japanese, Australian, and potentially other international collaborators; existing facilities that are owned and managed by the scientific collaborations; and facilities placed at outreach institutions. These sites will be connected by national and transoceanic networks ranging in speed from hundreds of Megabits/s to tens of Gigabit/s. An international Grid Operations Center (iGOC) will provide the essential management and coordination elements required to ensure overall functionality and to reduce operational overhead on resource centers.

Specific tasks to be undertaken in this project include the following. (1) Construct the international laboratory, including development of new techniques for low-overhead operation of a large, internationally distributed facility; (2) adapt current data grid applications and other large-scale production data analysis applications that can benefit from Data Grid technology to exploit iVDGL

features; (3) conduct ongoing and comprehensive evaluations of both data grid technologies and the Data Grid applications in the iVDGL, using various (including agent-based) software information gathering and dissemination systems to study performance at all levels from network to application in a coordinated fashion, and (4) based on these evaluations, formulate system models that can be used to guide the design and optimization of Data Grid systems and applications, and at a later stage to guide the operation of the iVDGL itself. The experience gained with information systems of this size and complexity, providing transparent managed access to massive distributed data collections, will be applicable to large-scale data-intensive problems in a wide spectrum of scientific and engineering disciplines, and eventually in industry and commerce. Such systems will be needed in the coming decades as a central element of our information-based society.

## References

---

- <sup>1</sup> Foster, I. and Kesselman, C. (eds.). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
- <sup>2</sup> The Particle Physics Data Grid, <http://www.ppdg.net>
- <sup>3</sup> European Union DataGrid Project home page, <http://grid.web.cern.ch/grid/>
- <sup>4</sup> GriPhyN Project home page, <http://www.griphyn.org/>
- <sup>5</sup> GridPP: A Grid for UK Particle Physics, <http://www.gridpp.ac.uk/>
- <sup>6</sup> iVDGL Proposal to NSF ITR2001 Program, [www.phys.ufl.edu/~avery/griphyn/itr2001](http://www.phys.ufl.edu/~avery/griphyn/itr2001)
- <sup>7</sup> The ATLAS Experiment, A Toroidal LHC Apparatus, <http://atlasinfo.cern.ch/Atlas/Welcome.html>
- <sup>8</sup> The CMS Experiment, A Compact Muon Solenoid, <http://cmsinfo.cern.ch/Welcome.html>
- <sup>9</sup> BaBar, <http://www.slac.stanford.edu/BFROOT>
- <sup>10</sup> D0, <http://www-d0.fnal.gov/>
- <sup>11</sup> STAR, <http://www.star.bnl.gov/>
- <sup>12</sup> Jlab experiments, <http://www.jlab.org/>
- <sup>13</sup> SAM, <http://d0db.fnal.gov/sam/>
- <sup>14</sup> NTON, <http://www.ntonc.org/>
- <sup>15</sup> The LIGO Experiment, <http://www.ligo.caltech.edu/>
- <sup>16</sup> The Sloan Digital Sky Survey : <http://www.sdss.org/>
- <sup>17</sup> National Virtual Observatory, <http://www.srl.caltech.edu/nvo/>
- <sup>18</sup> National Partnership for Advanced Computational Infrastructure, <http://www.npaci.edu/>