Particle Physics Data Grid

PI's: Richard Mount, SLAC, Miron Livny, Wisconsin, Harvey Newman, Caltech
Steering Committee: John Huth, Harvard, Tim Adye, RAL, Lothar Bauerdick, FNAL, Lee Lueking, FNAL, Chip Watson, TJNAF, Jerome Lauret, BNL Miron Livny, Wisconsin, Jennifer Schopf, ANL, Ian Foster, ANL, Reagan Moore, SDSC, Arie Shoshani, LBNL (SRM)
Coordinators: Ruth Pordes, FNAL, Doug Olson, LBN
Liaisons: Paul Avery (iVDGL), Larry Price (HICB), Mike Wilde(GriPhyN), Torre Wenaus, Ian Bird (LCG)
(www.ppdg.net, ppdg-exec@ppdg.net)

Summary

The Particle Physics Data Grid, PPDG, has been pioneering large-scale collaboration between computer scientists and physicists who have seemingly unbounded needs for data-intensive computing integrating worldwide resources. Distributed computing for high-energy and nuclear physics experiments has become both a driver for, and an eager consumer of, computational data-grid technology. As a SciDAC project, PPDG has focused on the Grid technology advances, ranging from architecture to robustness, that can be driven by deployment and integration into the demanding end-to-end applications of the physics scientific process.

The Particle Physics Data Grid, PPDG, is collaboration of six physics experiments[1] (ATLAS, BaBar, CMS, D0, STAR, TJNAF) with four leading computer science projects[1] (Condor, Globus, SRB, SRM). Several projects within the US, as well as other countries (especially The European Union), are developing and promoting grid technology. PPDG plays a unique role in stressing and facilitating early adoption of these technologies by the experiments to provide production services to their scientists so that the bugs and missing features that only show their "ugly head" when used "at scale" and "under load" are identified and communicated back to the computer scientists.

PPDG is vertically integrating advances in computer science such as novel mechanisms and policies, grid middleware, experiment-specific applications and computing, storage and network resources to bring effective end-to-end capabilities to the scientist's desk-top. Part of the PPDG mission is to facilitate the integration effort between several groups and projects, each with their own set of timescales and driving forces. We identify a specific piece of middleware to be deployed in each physics collaboration that is likely to bring short-term benefits. Teams from the computer science and physics groups carry out specific development and integration tasks delivering end-to-end capabilities.

Some of PPDG accomplishments during the first year and a half of the project are:

CMS Simulation Production on a Grid:
Members of the CMS physics experiment working together with participants of the Particle Physics Data Grid (PPDG) and the NSF funded grid projects, GriPhyN and iVDGL, succeeded in production run of physics simulations for 1.5M events delivered to the experiment from a test grid of five sites in the U.S. The integration of physics application code with several grid middleware packages running in a production-like fashion led to numerous bug fixes and error conditions being uncovered and fixed. Because the code improvements are cycled back through the Virtual Data Toolkit of iVDGL the benefit is shared across a much broader community than just the CMS experiment. [3]

STAR Robust coast-to-coast terabyte-scale file replication: We have used Storage Resource Manager (SRM) technology to achieve robust file replication for the STAR experiment studying quark-gluon-plasma physics at Brookhaven. [4]. The SRM monitors the staging, transfer, and archiving of files, and recover from transient failures.

In tests with the new grid-enable implementation, rates of up to 8 MB/sec for the wide-area-network stage have been achieved. After resolving some end-point configuration issues rates of 3-4 TB/week are being persistently achieved and making a quantitative difference to the rate at which the experiments can reprocess data.

Integrating D0 SAM with standard Grid middleware: The D0 experiment has had production deployment of their distributed data handling system, SAM, over more than 10 sites over three continents. As part of PPDG, D0 were the first application to use DOE SG certificates between a US and UK site, have integrated GridFTP as one of four data transport mechanisms in SAM, and are starting to deploy planned job scheduling over the Sam-GRID

following a collaboration on extensions to Condor.

US ATLAS interoperable end-to-end file replication: US ATLAS physicists have integrated their data challenge application across a grid based on US middleware in the US and on European middleware in Europe; Deployment of their end-to-end file catalog and replicator, MAGDA, has provided a requirements and evaluation platform for the emerging standard replica management servides.

BaBar throughput data transfers: Babar has established managed data transfer from the SLAC laboratory in California to Lyon France using the Storage Resource Broker, SRB;

TJNAF file replication: TJNAF has established remote data access as a grid web service together end to end automated file replication for their data archives.

In keeping with the goals of achieving end-to-end production services, PPDG recognized a missing connection with site computer-security personnel and so established a working group of the DOE laboratory site security teams to identify issues and recommend improvements to the grid security infrastructure.[6] Collaboration with DOE Science Grid has helped establish the public-key security services for the US physics grid community and trust relationship with European grid resources. Each computer science group in PPDG has made improvements and extensions to its designs to meet needs uncovered by PPDG. These improvements are already benefiting numerous other efforts within DOE's SciDAC program. To quote PPDG computer scientists "Almost every new feature required by Particle Physics turns out to be generic!"

Grid Middleware Benefits from PPDG: PPDG has contributed many improvements to the Globus Toolkit[7] in the form of requirements specifications, testing and hardening of many components including Grid Security Infrastructure, GridFTP, Replica Location Service, Monitoring and Discovery Service, GRAM job submission, and the Community Authorization Service.

PPDG had a similar impact on technologies and middleware developed by the Condor project[8] – most noticeably the grid enabled job manager (Condor-G) and workflow management.

ßCollaboration with PPDG has led to the addition of key features to the Storage Resource Manager[9] (SRM) including space reservation and management needs of disk and tape storage resources. Based on the extensive PPDG experience, the SRM technology has been used by other SciDAC projects, the Earth Science Grid (ESG) and Lattice QCD (LQCD).

Improvements to the Storage Resource Broker[10] (SRB) from work in PPDG includes extensions to the metadata catalog based on requirements of PPDG experiments as well as abstraction of the data transport interfaces to support the diverse data access technologies used across different application communities..

Collaboration with the "Bandwidth Estimation, Measurement Methodologies and Applications" project contributes to PPDG monitoring group and development of future plans for performance monitoring & troubleshooting of end-to-end applications. PPDG collaborates on other joint activities with the NSF-funded Grid Physics Network, GriPhyN and International Virtual Data-Grid Laboratory, iVDGL projects under the umbrella name Trillium.

Future Plans: In the next year, all of the experiment groups in PPDG are planning to roll out grid enabled job scheduling services. Many challenging areas of work require further research and development to meet our goals: troubleshooting end-to-end functional and performance of grid applications; multi-organization authorization and service provision and interactive science analysis services;

---

1 More about PPDG participants at www.ppdg.net.
2 www.ppdg.net/docs/news/news-item-20feb02.pdf
3 www.ppdg.net/docs/news/news-update-cmstestgrid-17may02.pdf
4 www.ppdg.net/docs/news/news-25sep02.pdf
5 www.ppdg.net/docs/news/news-25sep02.pdf
6 www.ppdg.net/pa/ppdg-pa/siteaa/
7 http://www.globus.org/toolkit
8 http://www.cs.wisc.edu/condor
9 http://sdm.lbl.gov/srm-wg
10 http://www.npaci.edu/DICE/SRB