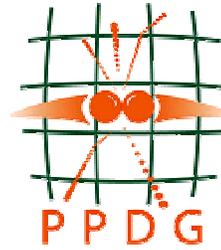


**Particle Physics Data Grid
Collaboratory Pilot**
**Quarterly Status Report of the
Steering Committee,
October - December 2001**

DRAFT 15 Dec 2001



Contents

1. Project Overview	2
1.1. Project Management and Organization	2
1.2. Project Web Pages	2
1.3. End to End Applications and Deployed Systems	2
1.4. Common Services Development and Integration	2
1.5. Meetings and Workshops.....	3
1.6. Architecture	3
1.7. Documents and White Papers	3
1.8. Plans for the next Quarter.....	3
1.9. Interactions with other Projects and Activities	3
1.10. Current Issues and Concerns	3
2. Project Activities.....	3
2.1. GDMP (CMS-DataGrid-Globus).....	3
2.2. D0 Job Management (D0-Condor).....	3
2.3. CMS-MOP (CMS-Condor)	3
2.4. STAR-DDM (STAR-LBNL)	3
2.5. JLAB-Replication (JLAB-SRB).....	3
2.6. ATLAS distributed data manager, MAGDA (ATLAS-Globus).....	3
2.7. BaBar Database Replication (BaBar-SRB)	4
3. Cross-cut Activities	4
3.1. SuperComputing 2001 Demos.....	4
3.2. Certificate/Registration Authority.....	4
3.3. Monitoring.....	4
4. Single Collaborator Efforts and End to End Applications.....	5
4.1. ATLAS.....	5
4.2. BaBar.....	6
4.3. CMS.....	6
4.4. D0.....	6

- 4.5. Jlab..... 6
- 4.6. ANL – Globus..... 6
- 4.7. NERSC – SDM..... 6
- 4.8. SDSC – SRB..... 6
- 5. Appendix..... 7
 - 5.1. SuperComputing 2001 demonstrations related to PPDG..... 7

1. Project Overview

1.1. Project Management and Organization

1.2. Project Web Pages

.

1.3. End to End Applications and Deployed Systems

.

1.4. Common Services Development and Integration

1.4.1. Data Replication and Catalog Services

.

1.4.2. Job Scheduling and Management

.

1.4.3. Storage Resource Management

.

1.4.4. Monitoring and Status Reporting

.

- 1.5. **Meetings and Workshops**
- 1.6. **Architecture**
- 1.7. **Documents and White Papers**
- 1.8. **Plans for the next Quarter**
- 1.9. **Interactions with other Projects and Activities**
- .
- 1.10. **Current Issues and Concerns**
- .

2. Project Activities

- 2.1. **GDMP (CMS-DataGrid-Globus)**
- !
- 2.2. **D0 Job Management (D0-Condor)**
- 2.3. **CMS-MOP (CMS-Condor)**
- 2.4. **STAR-DDM (STAR-LBNL)**
- 2.5. **JLAB-Replication (JLAB-SRB)**
- 2.6. **ATLAS distributed data manager, MAGDA (ATLAS-Globus)**

The principal goal of the Magda project for this period was the completion and deployment of a version capable of production deployment in the ATLAS Data Challenge 0 commencing in December. This was achieved, with a DC0-ready version deployed and announced on December 7. Magda was adopted by international ATLAS as the file cataloging and replication tool for DC0 and by the end of the period was in use cataloging the DC0 data generated to date.

The most important new functionality implemented during the period was the completion and deployment of command-line tools providing a file access interface to production jobs. The `magda_findfile` command searches the catalog on the basis of LFN, LFN substring, location, etc. and reports results in a parsable format. The `magda_getfile` command retrieves files from any accessible location, making them available locally either as a local copy or a soft link to a replica in a managed location. Usage counts of files in managed locations and caches are maintained, with usage decremented when `magda_releasefile` is used,

such that files can be pinned while they are in use. The `magda_ptufile` command archives files in managed store locations and registers them in the catalog. These command line tools provide all the capability currently needed by ATLAS jobs to exploit Magda, so the direct integration of Magda into the Athena framework continues to be deferred until manpower for this more exploratory work is identified.

Integration of GDMP with Magda was identified as the highest priority in further integrating Grid toolkit components with Magda. An integrated deployment of Magda and GDMP is foreseen in the ATLAS Data Challenge 1 commencing in Spring 2002, permitting ATLAS to draw on both PPDG and EDG WP2 data management efforts in a coordinated way. Towards this end, the GDMP design and feature set was reviewed with a view to Magda integration, and an integration plan begun. Problematic issues in the integration were identified and gathered for discussion at a PPDG data management meeting in early Jan.

ATLAS/PPDG has been instrumental within international ATLAS in planning and coordinating a coherent approach to replica and metadata management for the ATLAS Data Challenges, integrating the plans and deliverables of PPDG and EDG.

Magda deployment was completed or initiated at several new sites during the period, including Indiana University (completed), IN2P3 and UT Arlington (underway). Magda-based replication of ATLAS data between CERN and BNL continued, with ~300GB of data now replicated. Magda now catalogs files representing more than 6TB of data.

Near term plans include exercising Magda in a production setting in DC0 and feeding experiences back into the development cycle; integration with hybrid (RDBMS+object streaming) event stores; integration with application metadata catalogs; integrating GDMP in preparation for DC1; and further integration of Globus tools, particularly remote command execution for more flexible Magda usage at testbed sites.

During the period we developed (primarily off-project) a design and description of a 'hybrid' persistent data model consisting of data files plus a data management and metadata layer, the latter to be implemented using a combination of grid toolkit components and higher level metadata services. The hybrid data model is a proposal for managing the event data in an HEP experiment. It explicitly recognizes that the data is stored in files and separates the largely grid-based management and tracking of those files from the management of event data objects within the files. It addresses the problem of maintaining persistent references between event objects. The management of physics data collections (called datasets) is also discussed. Most of the work thus far is in design work directed at file-level management of and access to distributed event data, directly applicable to our PPDG program in distributed data management development. For details see <http://www.usatlas.bnl.gov/~dladams/hybrid>.

2.7. BaBar Database Replication (BaBar-SRB)

3. Cross-cut Activities

3.1. SuperComputing 2001 Demos

3.2. Certificate/Registration Authority

3.3. Monitoring

4. Single Collaborator Efforts and End to End Applications

4.1. ATLAS

4.1.1. US ATLAS Grid Testbed

Testbed sites continued to deploy and test additional grid infrastructure components, including Magda, the pacman package manager, GDMP, and Globus 2.0 (GridFTP, replica catalog, etc.). Work is underway to support the automated distribution via pacman (developed within ATLAS GriPhyN) of all components needed for deployment of a Grid-integrated testbed site capable of running ATLAS software. Regression tests distributed by pacman to validate various grid services (Globus, GDMP, Magda, etc.) on the testbed are also under development. We are becoming involved (Jerry Gieraltowski at ANL) in the ATLAS grid validation activities taking place within EDG. Network performance monitoring and tuning activities between BNL and US ATLAS grid testbed sites continued during the period.

4.1.2. Monitoring

ATLAS continued its involvement in the PPDG Grid Monitoring Project, developing use cases, defining the scope of the project, designing and developing linux farm monitoring using MDS. A prototype Linux farm monitoring tool using MDS was developed, based on the existing BNL farm monitoring system. The system can monitor up to 300 nodes, and the scalability will improve in future versions. Currently the system can answer limited questions which Grid users might ask via MDS, eg. "give me 40 least loaded nodes with Linux Kernel 2.4".

4.1.3. Distributed job management

Distributed job management activity focused on ATLAS DC0-directed development and deployment of a job management infrastructure integrating the use of distributed data management (Magda) and application metadata management (under development by Grenoble ATLAS) tools. Development of distributed job management proper will begin post-DC0, since DC0 does not involve distributed production. ATLAS DC1, commencing spring 2002, will involve distributed production, and we anticipate developing, deploying, testing, and iterating the development of distributed job management tools before and during the ~6 month duration of DC1. We will look closely at existing tools for adoption, in particular the MOP package of CMS/PPDG.

4.1.4. Data signature

During the quarter a preliminary design was developed for the 'event data history' classes that will constitute part of the data signature required for data equivalency tests or on-demand data regeneration. The goal of event data history is to record the history of data at the level of individual event data objects (EDO's), i.e. the collections of physics objects (clusters, tracks, electrons, ...) that comprise the event data of high energy physics. We require that the history be sufficient to reproduce the data at that level. We identify three levels of history objects:

1. Algorithm history
2. Job history
3. EDO history (includes pointers to its job and algorithm histories)

Classes describing each of the above were developed and can be found in the ATLAS repository under Control/AthenaHistory. For details see http://www.ustalas.bnl.gov/~dladams/data_history.

4.2. BaBar

4.2.1. Network throughput performance

4.2.2. Replica Catalogs in the Globus Framework

4.3. CMS

.

4.4. D0

4.4.1. Integration with GridFTP

.

4.4.2. Distributed Analysis

4.4.3. Test Harness

4.5. Jlab

.

4.6. ANL – Globus

.

4.7. NERSC – SDM

.

4.8. SDSC – SRB

4.8.1. SDSC - JLab Replica management interface

4.8.2. GridPortal project at SDSC.

.

4.8.3. BaBar Support

.

5. Appendix

5.1. SuperComputing 2001 demonstrations related to PPDG

CMS Simulation Production – IMPALA and GDMP	FNAL, Caltech	Demonstration of current CMS simulation production tools and GDMP replication tools
CMS Distributed Simulation Production (MOP)	Caltech, FNAL, Wisconsin, ANL, UCSD	Use of Condor-G/DAGMAN to automatically run CMS simulation production at multiple sites
Bandwidth Greedy Grid-enabled Object Collection Analysis for Particle Physics	Caltech, UCSD	Demonstration of the use of Grid tools and virtual data to support interactive physics analysis.
Reliable Transport	ANL	Extensions to the transport layer of GridFTP to support retry
Proxy Server Demo	ANL, SLAC	Demonstration of replica catalog proxy server
GriPhyN Virtual Data (CMS)	ANL, Florida	Generation of CMS simulation scripts from definition of physics parameters
Globus CAS prototype	ANL	Use of Community Authorization Service in Earth Sciences Grid
"Letting Scientists Concentrate on Science: Providing a Transparent View of Data on the Grid"	LBNL	http://gizmo.lbl.gov/~arie/sc2001.demo/slides/index.htm http://gizmo.lbl.gov/~arie/sc2001.demo/poster.pdf
"Bandwidth to the World"	SLAC/FNAL	The "Bandwidth to the World" project is designed to demonstrate the current data transfer capabilities to about 25 sites with high performance links, worldwide. (http://www-iepm.slac.stanford.edu/monitoring/bulk/sc2001/)
SDSC Grid Portals Architecture	SDSC	The SRB team has been working with Grid Portal Architecture group to use SRB in building Grid Portal services.