Particle Physics Data Grid Collaboratory Pilot
Quarterly Status Report, April - June 2001

13 Sept. 2001

Contents

1. Project Overview

PPDG is in transition from the funding under the original NGI proposal, to that of the increased funding from the three year SciDAC funded project. Many sites received funding from the SciDAC project and are enthusiastically progressing with their work. We now have an opportunity for coordinated development, to look towards more significant development projects to bring Grid services into use in the experiments data processing and analysis systems and to contribute reused and common production components to the field.

Several Project Activities have been established and collaborative development work is proceeding. At the Collaboration meeting on 14-15th August we will cover the scope of the current work, discuss our goals for the next year and how to organize our efforts towards deliverables and milestones.

1.1. Project Management and Organization

Bi-weekly project phone conferences continued. The management plan as defined in the proposal (See Appendix) came into effect, with the Executive Team and Steering Committee holding periodic phone conferences. The Steering Committee started to address policy matters and interactions with other projects and activities.

The Steering Committee agreed that PPDG would approach and work with the DOE Science Grid SciDAC project to understand the potential for a common Certificate Authority, hosted by the Science Grid. Contact has been made with Bill Johnston to follow up.

The Steering Committee agreed we would support an SC2001 demo of the CMS distributed simulation production, working across the booths of several collaborators both on and off the SC2001 floor. Work is ongoing to define and understand the details.

We have started encouraging individual effort reporting. This is a work in progress to try and understand the most effective mechanisms to meet the needs of the individuals, project, and oversight groups.

1.2. Project Web Pages

The project Web pages have been established and are used to record documents and talks, meeting agendas and minutes, email lists and archives. A calendaring tool is used to good effect for planning and organizing purposes. By storing the pages in CVS and automatically updating the web pages on an hourly basis we are hoping to allow shared responsibility and multiple contributors.

1.3. Common Services Development and Integration

1.3.1. Data Replication and Catalog Services

PPDG work on the Globus Replica catalog and replication management services has progressed. GDMP as an extension of these services is being further developed and deployed by CMS for data movement. GDMP is a collaborative project between PPDG and the European Data Grid (EDG). Atlas and GDMP are using the Globus replica catalog and extensions to this service are in progress . Several experiments are testing GridFTP and are ready to use it when a production version is released by the Globus team. GDMP will test the Globus replication services after the EDG Project Month 9 integration milestone. SRB replication services are being interfaced to the BaBar databases as part of the Objectivity data replication needs. Jefferson Lab is working with the SRB team to define common XML schemas for mass storage and disk caching systems and for data replication and catalog services, so as to support future common data grid web services.

1.3.2. Job Scheduling and Management

Condor has been interfaced to the D0 SAM system by UK collaborators, CMS Tier-1-2 Simulation production will be deployed using Condor in the next few months (MOP PPDG project). Atlas is testing Condor for use in job scheduling across Tier-1 and 2 US sites.

1.3.3. Storage Resource Management

STAR is extending their existing file transfer activities between LBNL and BNL by interfacing through the DRM (Disk Resource Manager) to the HPSS MSS at LBNL. GDMP is testing the defined HRM interfaces through providing at interface to the Caltech HPSS and Fermilab Enstore systems. JLAB is also contributing to these definitions through the LBNL and Fermilab groups working as part of the Storage Resource Management SciDAC project.

1.3.4. Monitoring and Status Reporting

Discussions at the coming Collaboration Meeting will allow us to start to define concrete work and goals in this area.

1.4. Interactions with other Projects and Activities

PPDG pays particular attention to inter-working with other projects in the field providing distributed computing services – both hardware and software – to the PPDG experiments, the experiment data handling teams, and the general Grid development and standards bodies both research oriented and commercial. During the last quarter members of the project have participated in and contributed to:

- Grid Portals Workshop, SDSC, May 21-23
- HENP Networking Group, University of Indiana, June 1-2nd
- Euro-Globus workshop, Lecce, June 16-22nd
- HENP Data Grids Coordination meeting, Rome, June 23rd.
- European DataGrid Workshop, Oxford, July 2-5th,
- Global Grid Forum 2, Washington DC, July 15-18th
- GriPhyN applications meeting, Chicago, August 2-3rd
- HPDC 10, San Francisco, August 7-9th
- Globus Retreat, San Francisco, August 9-10th

1.5. Current Issues and Concerns

The announcement of PPDG funding was officially disseminated in June, but as of August not all the year 1 funding has been received. This is hampering the ability of some of our institutions to ramp up their effort on PPDG to the expected level, and is thus affecting progress.

We are still understanding the boundaries of responsibility and work between the various experiment and grid projects to identify those areas of explicit PPDG deliverables and focus. Various "cross-cut" activities are happening naturally, and others need steering committee direction. Given the activity and energy in the whole field at the moment it is important that we remain flexible in this area.

2. Project Activities

The major part of the PPDG work is organized as Project Activities (PPDG Sub-Projects). These are formed as working teams of Experiments and Computer Science Groups – with an identified Project Leader

and a Liaison/Coordinator from one experiment and one Computer Science Group (where the project leader may or may not be the same person as one of these). The goal is that each project defines deliverables, milestones and necessary effort to allow them to work semi-independently towards their goals. PAs report technical progress and details at the regular bi-weekly phone conferences.

2.1. GDMP (CMS-Globus)

GDMP has originally been created by one PPDG/CMS and one DataGrid/CMS member. In April, GDMP has become an official, mutual software project between DataGrid and PPDG. Thus, man power is provided by both projects: architecture, design and implementation is done in common and thus shall guarantee that the code base of GDMP is interoperable between both Data Grid projects. GDMP Project management is done by Asad Samar and Heinz Stockinger. During the quarter, also the contacts to and interactions with the Globus team have been increased.

The main effort in this quarter was to include work on a flat file support for transferring and replicating files of arbitrary file format. In particular, new code development has been done and some existing parts have been redesigned and updated for the flat file support. Flat files can also be transferred along with Objectivity files. As GDMP assumes that the files may not be available on disk, a new step of registering files with GDMP is added. Any file to be transfer via GDMP will first be registered with GDMP. Once registered, files can be moved to a Mass Storage System. For Objectivity files, attachment of dummy database files with an Objectivity federation and then replacement with the original DB files is also available. Finally, if Objectivity files are transferred, GDMP checks if the Objectivity lock server is running.

The GDMP source tree has been converted to build with the GNU Autotools. The new build system makes GDMP more portable and easier to configure at remote sites. The entire code of GDMP is accessible via a CVS repository at Fermilab: http://cdcvs.fnal.gov/cgi-bin/public-cvs/cvsweb-public.cgi/gdmp/

Another major task was to interface GDMP with the Hierarchical Resource Manager (HRM) server provided by LBNL and Fermilab. The work on GDMP-HRM integration is progressing well. Currently, HRM provides APIs for reading files from the Mass Storage System (MSS). The functionality for writing to an MSS is being incorporated through a mutual effort between Fermilab and LBNL. We have successfully integrated the read functionality in GDMP so now a GDMP server can talk to an HRM server and request some files to be staged from the MSS to disk. We are waiting for FNAL/LBL to standardize and implement HRM-write so that we can use that in GDMP as well. The HRM functionality will not be made part of the GDMP production release until we integrate both the reads and writes for consistency purposes.

A preliminary graphical user interface in Java has been provided and will be extended in the future.

DataGrid will have a major software milestone in September 2001 where a preliminary Data Grid testbed will be in place. A member of PPDG will take a major role to provide support for GDMP in the testbed activities which have already started now.

For progress on GDMP also refer to the Data Management Work Package (WP2) in the DataGrid project:

http://grid-data-management.web.cern.ch/grid-data-management/

2.2. D0 Job Management (D0-Condor)

The work of the project is to architect a system where Condor's job management is integrated with D0's data handling system. On June 20, 2001, there was an initial meeting to present and discuss status reports of the preliminary work done. P. Couvares visited FNAL for discussions. A study of data access patterns in D0 data analysis was performed http://projects.fnal.gov/act/sam_cache/sam_cache.html. Since then, no futher work has been done, on account of D0 not having filled the recently created job openings, as well as of UW to date not having received project funding. D0 anticipates to be ready to proceed with the collaborative work in late August or early September as the new hires are made.

2.3. CMS-MOP (CMS-Condor)

The CMS Monte carlo distributed Production (MOP) is working toward its first distributed test. The three primary components of MOP are file transport, provided by GDMP, job definition, provided by Greg Graham's CMS job tracking and creation (JTTC) system, and the job distribution system, mop_submitter. The architecture document was written during this quarter. It is available at <http://www.ppdg.net/pa/ppdg-pa/mop/mop.pdf>.

The file transfer system for this project is based on GDMP. However, at the beginning of the project, GDMP lacked a major features needed for MOP: replication of flat files, i.e., files not attached to an Objectivity database. Shahzad Muzaffar worked on this and other features of GDMP during this quarter; further information is available in the GDMP section.

The job definition system is the most CMS-specific portion of MOP. Greg Graham has been working on the JTTC system for job definition. A schema for job tracking has been developed and implemented using a MySQL database. The integration with the CMS production scripts is near completion.

During the past quarter James Amundson worked on the job distribution system. After consultation with Peter Couvares from the Condor team, mop_submitter has been implemented to share jobs between remote sites through the Condor flocking system. Condor G and Globus are used to interface to local batch systems.

Integration work for MOP is just getting underway. Multi-site testing is planned for August.

2.4. STAR-DDM (STAR-LBNL)

The STAR experiment at RHIC generated about 4TB worth of raw data in its first run in the summer of 2000. STAR's second run will start in the summer of 2001 and is expected to generate approximately 50 times more data than the first run, or some 200 TB. In PPDG year 1 the primary activity of STAR-DDM is to deploy and integrate Grid services for site-to-site file replication between STAR's tier-0 site at BNL and its tier-1 site at LBNL. In year 2 the focus will shift to job control and management of production computing, and in year 3 we plan to concentrate on coordinated storage and computation.

In the past quarter our efforts have been directed towards setting up the necessary infrastructure to begin using Data Grid services during STAR's 2001 run. Specifically, a STAR grid node has been created at BNL and efforts have been made to attain consistent Globus installations at both sites. Storage resource management components have been in development and will be adopted as they become available in the next two quarters. We hope to begin testing GridFTP sometime in the next quarter.

2.5. JLAB-Replication (JLAB-SRB)

Jefferson Lab has continued to prototype the use of XML based web services and client applications for managing data grids and batch systems, and has interfaced these prototypes to the existing lab developed JASMine file management software. Discussions were held with the SRB team to begin a common effort to standardize the XML schemas used to described data storage systems, with a view towards making web based client applications independent of the actual storage system used by a particular site.

A report detailing the development of advanced data schemas and user interfaces to data collections has been written by the SRB team with input from Jefferson Lab, and distributed to members of the JLab team. The report describes the general requirements for browsing and querying of large data repositories, presentation of query results including aggregating information about digital objects into range-based hierarchically ordered attributes, and support for management capabilities. The report also lists a set of deliverables for this PPDG sub-project, including software components to be collaboratively developed, documentation to be produced and presented to the Global Grid Forum, and a testbed at Jefferson Lab as an initial deployment. This plan of work will be further developed over the summer, with preliminary designs to be presented at GGF3 in Italy.

2.6. ATLAS distributed data manager (ATLAS-Globus)

The principal PPDG year 1 deliverable for ATLAS in the first year is the delivery of a production distributed data service deployed to users. The distributed data manager activity in ATLAS PPDG is developing the system which will be deployed to meet this deliverable. The production distributed data service is to exist between CERN, the ATLAS Tier 1 facility at BNL, and a number of US ATLAS grid testbed sites (some or all of ANL, LBNL, Boston University, Indiana University, U Michigan, U Oklahoma, and UT Arlington). The objective is a multi-point U.S. Grid (in addition to the CERN link) providing distributed data services as early as possible.

Development of the distributed data manager was initiated during the period, based on the DBYA prototype introduced in the last report as a starting point. DBYA provides ATLAS-specific infrastructure and metadata, and the infrastructure to vertically integrate the data manager from the grid toolkit components employed through to the ATLAS-specific interfaces by which the system is used. Grid toolkit components will be progressively integrated into the system. During the period a facility to load the file replica metadata of the system into the Globus replica catalog was developed. Over the next few months we expect to extend the functionality of the system from the present passive cataloging to active replication of files between grid sites using GridFTP.

We provided input to the Replication Requirements document in development, based in part on DBYA prototyping experience.

We are continuing to use ATLAS tile calorimeter test beam raw data in an Objectivity database to test the data transport and data replication tool sets.

A Globus replica catalog running at ANL with catalog query and update clients running at BNL was successfully demonstrated and exercised, both from a command line interface and from C programs. This is a first step in a program to demonstrate grid-enabled data access from Athena, the ATLAS experiment's control framework, in the coming quarter. An abstract for a paper describing this work was accepted for presentation and publication in the proceedings of CHEP'01, Beijing, September 2001.

See http://atlassw1.phy.bnl.gov/dbya/info for more information on the distributed data manager prototype.

2.7. BaBar Database Replication (BaBar-SRB)

The first year goal is to attempt to incorporate the SRB into BaBar inter-site data replication. The first step is to build a prototype to replicate databases from on federation to another. We also plan to carry out scaling and reliability tests to make sure that SRB based model meets BaBar's needs. The experience obtained from building and using the prototype will be used to put SRB into production provided it satifies BaBars needs. Currently, we plan to also investigate using the GLOBUS toolkit for intra-site BaBar data replication. We intend to setup a complementary prototype based on the GLOBUS toolkit to understand the issues involved in using GLOBUS for BaBar data replication as well as to understand how to interface GLOBUS with SRB.

The first version of the SRB prototype is ~2 FTE months away from completion. We expect that exposure of the prototype to members of the BaBar Data Distribution group will necessitate further modifications and improvements. We expect this iterative cycle to take ~1 FTE year to complete. People actively working on the prototype are: Andy Hanushevsky, Adil Hasan and a new-hire. Work on the GLOBUS prototype will start once the first version of the SRB prototype is finished. We expect prototyping using the GLOBUS toolkit will take ~1year FTE, people working on this effort will be Andy Hanushevsky, Adil Hasan and a new-hire.

The incorporation of SRB and GLOBUS into the BaBar data replication scheme will require a re-design of the BaBar Data Distribution tools. The re-design is currently being worked on by Yemi Adesanya and Adil Hasan and should be complete in 1 year FTE.

This effort should result in a suite of software capable of replicating BaBar data from one federation to another within a site.

Components: SRB, Objectivity, Tools based on Objectivity and SRB that allow us to replicate dbs from one federation to another. Tools to interface between the SRB and GLOBUS.

Group members: Yemi Adesanya (yemi@slac.stanford.edu), Makoto Asai (asai@slac.stanford.edu), Bob Cowles (rdc@slac.stanford.edu), Andy Hanushevsky (abh@slac.stanford.edu), Adil Hasan (hasan@slac.stanford.edu), Richard Mount (rmount@slac.stanford.edu).

Physics Project Leader: Adil Hasan, Architect: Andy Hanuschevsk, SRB Contact: Arcot Rajaseka, Bing Zhu (SDSC) 10% GLOBUS Contact: TBD.

Near-term goals: SRB prototype should be functioning by end Sept and should be robust enough for SC2001 (see the document at http://www.slac.stanford.edu/~hasan/ppdg/SRBPrototype.html). This page will contain documentation and results of the data replication effort.

3. Single Collaborator Efforts and End to End Applications

As well as the identified Sub-Projects, PPDG effort focuses on end-to-end applications and demonstrators. This is in keeping with our mission of a short to medium term focus, with working systems used to input requirements to current and future projects, feedback to ongoing designs and implementations, and a basis from which to discuss future needs and work. In some cases a fraction of the work to be done and reported is through off-project effort. PPDG benefits significantly from this synergistic work and uses it as input to the discussions and desicisions for on-project efforts and goals.

3.1. ATLAS

3.1.1. US ATLAS Grid Testbed

The US ATLAS Grid Testbed is now composed of 8 US ATLAS sites, ANL, BNL, LBNL, Boston U, Michigan U., Indiana U, Oklahoma U., and U of Texas at Arlington. Each provides a dedicated Globus 1.1.3/4 gatekeeper and additional computational resources for the testing and evaluation of GRID technologies as applied to ATLAS. During the period the Globus services have been in stable operation and the focus has been on the installation of ATLAS software and infrastructure at the testbed sites. During the summer and fall the focus will shift to applications testing.

See http://www.usatlas.bnl.gov/computing/grid/ for more testbed information.

3.1.2. Monitoring

A simple tool kit based on iperf was created to facilitate performance monitoring and tuning among the different testbed sites. A low bandwidth problem between BNL and NERSC was diagnosed and a solution is in development. A web site gathering network performance information is in development at http://yu.rhic.bnl.gov/~cricket/cricket/grapher.cgi

3.1.3. Distributed job management

A secondary year 1 objective for ATLAS, of growing importance in the out years, is distributed job management. During the period an ATLAS GriPhyN effort was initiated (Saul Youssef, BU) to study and test the capabilities of Condor to manage a hierarchical job management infrastructure incorporating the various tiers of grid sites. The ATLAS PPDG team is in contact with this effort, and discussions are underway on a collaborative effort with PPDG.

See http://physics.bu.edu/~youssef/atlas/notes/ for more information on the Condor scheme being investigated.

3.2. BaBar

Investigations have been done and published at the HPDC 10 conference into new, novel and scalable approach for distributing a replica catalogue and resolving .le location information by using HTTP

redirection. HTTP redirection servers managing local catalogues allow for greater flexibility and local management autonomy whereas a global replica catalogue provides the necessary mapping of logical to individual sites. By distributing the catalogues a site can autonomously move for load balancing within a site without notifying a global replica catalogue. Our approach scales well to a large number of sites and entries and thus establishes powerful middleware service

3.3. CMS

Many people are contributing to these efforts. We provide a summary of the efforts here and more details are available through the individual efforts reports on the PPDG web site.

[1] We continued the development of the distributed service system, based on JINI & JAVA technology. The aim of this prototype system is to provide an information management framework as a flexible and reliable distributed service system, applied for specific data processing tasks in HEP

[2] Another activity t Caltech's Tier2 Prototype center is to make WAN performance meassurements, as a function of the TCP/IP stack chosen and the associated parameters. That includes throughput tests between systems located at Caltech, CERN's PoP at Chicago and University of Florida. There were some variations in network performance due to network congestion. Efforts are underway for direct connection of Caltech's node to Calren which will bypass all the shared campus traffic. Further measurements at higher speeds will be done over NTON (continuing earlier measurements by J.Bunn et al.), once that network is back in operation.

[3] Clarens is a client/server system for the remote retrieval of persistent histograms and tags stored in Objectivity database federations, as well as remote analysis of physics data using commodity access protocols. A web site describing this work is available at http://heppc22.hep.caltech.edu/. The Clarens server provides a distributed interface to existing HEP reconstruction and analysis software, including ORCA, CARF, and Anaphe. This leverages web-based transport, authentication and encryption technology included in the open source Apache web server. A functional prototype for data retrieval is running on the Caltech Tier-2 server, while the remote analysis functionality is in active development.

[4] Full CMS software installation was completed on the Tier2 prototype and made ready for production. 350k CMSIM events and 200k fully reconstructed events have already been produced during this period, for our local studies and for the CMS official production at CERN. The FNAL script-based submission system and CERN bookkeeping MySQL system are undergoing testing now. They will be ready for full scale tests during the Fall 2001 production, together with a "Testbed1" prototype from the EU datagrid. Large Scale Distributed Simulations: These continue on the CACR X-Class Exemplar, where we are generating two million background events for Higgs --> gamma gamma studies. The full CMS software has been used for local production on both the "RoadRunner" and "Los Lobos" Linux clusters in New Mexico The full CMS software has been installed and tested on the new Platinum Linux cluster in NCSA which is part of the recently approved (by NSF) Data Terascale Facility. After some tests full-scale production will start at the end of August.

[5] Robust Execution Service (RES) - The goal of the project is to create a job scheduling tool that aids high energy physicists in effectively using computational resources distributed worldwide. The novel features of the system are: a) Keeping track of large number of long running jobs. b) Supporting collaboration among multiple physicists. c) Conserving limited network bandwidth. d) Maintaining high availability. e) Tolerating partition failures that are common in wide area networks. In the first month of the quarter, the system was fully debugged to support running jobs under different user IDs given proper digital certificates

[6] The document 'CMS Data Grid System Overview and Requirements' (CMS Note 2001/037) was completed on July 16, 2001. This document specifies the long term requirements up to 2003 for the Grid projects that CMS is involved in, including PPDG. A demo of the use of grid technology by CMS (remote data extraction of CMS data using Grid tools) was done at HPCN Europe 2001 in Amsterdam (June 25 - 27, 200).

3.4. D0

D0 continued to deploy and support the SAM Data Grid system for the collaborations use. The system delivered several hundred Gbytes of simulation data produced offsite at Lancaster, Imperial College, Nikhef and IN2P3 to the Fermilab mass storage system using the bbftp software which supports multiple TCP streams for high speed file replication. Replication of data for analysis is starting to SAM stations at Fermilab and the University of Michigan. SAM was interfaced to the local mass storage systems at NIKHEF (SARA) and IN2P3 (HPSS). SAM now contains over 45,000 files and manages over 30 TeraBytes of data. Plans were started towards SAM developments in the PPDG context in a paper submitted to Chep 2001 (posted off http://d0db-dev.fnal.gov/sam_talks/) . Members of the D0 PPDG team attended the GGF2 meeting in Washington, and presented a paper on Resource Management in SAM at the 10th HPDC symposium (http://www-itg.lbl.gov/HPDC10/).

3.5. Jlab

Jefferson Lab has held discussions with Florida State University to deploy a prototype grid aimed at distributed data analysis of data coming from the CLAS detector, which produces over 50 TB of data / year. This system will be used to test file replication and eventually a full range of web services for data management and batch systems.

3.6. STAR

3.7. Wisconsin – Condor

3.8. ANL – Globus

Activities conducted by the Globus Project as part of its involvement in PPDG during the quarter April-June 2001 include:

A framework for replica catalog performance analysis was set up, permitting open source (OpenLDAP), Netscape/Sun/iPlanet, and Oracle servers to be tested. Initial results confirm that performance issues exist with the current schema, and schema revisions are under consideration. Finalization of the tests and publication and analysis of the results will take place in August.

Testbed systems were created within the ANL Globus Group to support PPDG experimentation, and to test the inclusion of ANL-MCS resources in the Atlas test grid. An important function of the testbed is to support activities to enhance the integration of Globus services in GDMP, and to further integrate GDMP as a Grid service for reliable file transfer.

Minor changes in GridFTP to support its integration into GDMP (to replace the earlier gsiftp service) were undertaken and are set for delivery to GDMP developers in August. Other changes to the replica catalog for later feature support in GDMP, including generalizing the flexibility of the mapping of logical to physical file names, was discussed at EuroGlobus and is under consideration by the Globus Project.

Support for Atlas activities including a CHEP paper on replication services within Athena; a class in replication for the CERN Computing School, and the introduction of replica services into the Atlas test grid have been underway. Also underway is an effort to introduce a fully distributed Grid Information Service based on the Globus MDS 2.0 service, with a two-level hierarchy of Grid Information Index Service and Grid Resource Information Service.

An analysis document was started to develop detailed use cases to analyze the requirements for replication within the PPDG experiments has begin and several revisions to this document have been circulated and commented on.

Design studies and discussions have begin towards the creation of a reliable file transfer service that would treat data transfers much like a scheduled job, with recovery in the face of client and server crashes, handling of space allocation and recovery from out-of-space situations, and the support of striped and parallel transfers and of replica catalog management. An early draft of a requirements and service feature

specification document is being reviewed within the Globus Project and has been shared with members of the Condor Project, with whom many discussions have taken place concerning the requirements for this service.

3.9. NERSC – SRM

People involved: Junmin Gu, Andreas Mueller, Alex Sim, Arie Shoshani

Version 2.0 of HRM was incorporated into the GDMP system. This installation was carried out by Asad Samar with help from Alex Sim. This version uses HSI to interface to HPSS, and eliminated the need for a pre-loaded file catalog. This vesrion was also adapted to run on Linux. Work has started on an enhanced version of HRM that includes a "write" capability. The design uses a single queue to support both "read" and "write" requests. The main concept is that a file written into HRM, first gets transferred into the disk cache, and then queued for migration to tape. A call_back is provided for both events: when the file has transferred to cache, and later when the file is migrated to tape. This version of HRM, version 3.0 is now in testing.

We have developed a new unified API for both HRMs and DRMs. The design considerations and the details of the API was discussed in a day long meeting with people from Fermi. It is the basis for the new versions of HRM and DRM we are developing.

We have completed a first prototype of a "basic DRM" that performs "reads" only. It was adapted to work with the new unified IDL. It has a "pinning" capability, and a time-out mechanism in case file get not "released" by the client. This DRM is now in the testing phase. A new general purpose DRM was designed, and is now being implemented. It is designed to support both "reads" and "writes" of requests having multiple files. We plan to use a restricted version of it as part of HRM for supporting the management of the HRM's disk cache. The first version of the DRM is expected to be completed by the end of the next quarter, and then it will be tested rigurously.

3.10. SDSC – SRB

The San Diego Supercomputer Center is collaborating on multiple projects that are developing data grid technology. The explicit PPDG activities include support for database replication for the BaBa experiment, and collaboration with JLab on the definition of replication attributes. Associated efforts that are funded through other projects that potentially can impact the PPDG data grid include development of a Grid Portal that is integrated with the Storage Resource Broker / Metadata CATalog collection management software, and collaboration with a United Kingdom data grid that will also access Babar data. Finally, extensions to the SRB are under development to support parallel I/O, improve error handling consistency, and improve user interfaces. Each project is described in detail in the effort reports on the PPDG web site.

4. Appendix

4.1. PPDG Management Plan

U.S Physics Experiment computing management and Computer Science project management have joined the PPDG proposal with a common vision to work together to develop and apply Grid software to the benefit of each experiment, and with a commitment to adapt, extend and reuse information technology and software used in one experiment for the benefit of other experiments. The PPDG management plan is designed to provide sufficient coordination of the specific project deliverables and milestones, as well as show our commitment to a process of working together towards the common goal.

The PPDG management strategy has top-down and bottoms-up components. The bottoms-up strategy involves exploiting the mission-oriented drive of the physics experiments, part of which is already directed towards strengthening collaborations with PPDG-CS and with U.S. and European Grid projects outside of

PPDG[1]. The involvement of the CS projects is driven by their desire to test and evaluate their technology in a real-life setting.

The top-down strategy must ensure the steering of the powerful bottoms-up forces to ensure that, while the individual goals of each of the experiments and the CS projects are met, most developments are of value to other PPDG collaborators, and with little or no additional generalization, to a wide user community. In addition, PPDG must join in the effort to steer collaborations of individual PPDG members with other Grid projects such that a common Grid architecture emerges and all funding is used effectively.

PPDG plans a two component management structure with:

1. An **executive team** composed of Ruth Pordes (PPDG Steering Committee Chair), Doug Olson (Steering Committee Physics Deputy Chair) and Miron Livny (Steering Committee Computer Science Deputy Chair). All three members of this team have extensive experience in managing software development and deployment projects and will each make PPDG a principal activity. They will be tasked to track the goals and work of the physics-CS activities teams and provide guidance to ensure overall coherence of the PPDG Collaboratory Pilot. Miron Livny will steer the project deliverables towards maximal commonality (and wide usability) in the components of each experiment's vertically integrated Grid software. The team will work together to advise the Steering Committee to best meet the short and long term goals of the Collaboration.

2. A **Steering Committee** (SC) comprising one representative from each physics experiment and one representative of each Computer Science group. The project PIs and the executive team will be ex officio members of the Steering Committee. In the interests of keeping the SC small and effective, the project PIs and members of executive team may also act as experiment or CS group representatives if they wish. The preliminary membership of the Steering Committee is:

Ruth Pordes, Chair/D0 Rep.

Miron Livny, Project PI/Computer Science Deputy Chair/U.Wisc CS Team Rep.

Doug Olson, Physics Deputy Chair

Richard Mount, Project PI/BaBar Rep.

Harvey Newman, Project PI

Lothar Bauerdick, CMS Rep.

Torre Wenaus, ATLAS Rep.

Chip Watson, JLab Rep.

Matthias Messer, STAR Rep.
Ian Foster, ANL CS Team Rep.

Arie Shoshani, LBNL CS Team Rep.
Reagan Moore, SDSC CS Team Rep.

The steering committee is structured to provide decisive management and to balance the Computer Science, software technology, physics-experiment needs and budgetary constraints.

Clearly, the success of such coordination and additionally the success of coordination with other HENP Grid projects will be based on an open discussion and decision making process. Both the Steering and Executive committees will publish agendas and minutes of their meetings and are charged with ensuring that successful communication and coordination exists throughout the Collaboratory Pilot through the three years of the project.

[1] Existing and planned collaborative activities involving PPDG, GriPhyN, and European Grid projects are described in the letters of support attached to this proposal.