


Particle Physics Data Grid Collaboratory Pilot
Project Plan June 2003 – June 2004

PPDG Steering Committee,
 August 10th. 2003



1. PPDG, June 2003	1	3.1 Condor	10
1.1 Project and Team Plans.....	1	3.2 Globus.....	10
1.2 SciDAC Review	2	3.3 Storage Resource Broker	11
1.3 Follow up to Questionnaires.....	2	3.4 Storage Resource Management (SRM)	11
2. PPDG Cross-Project Plans	4	3.5 Virtual Data Toolkit (VDT).....	12
2.1 Data Management	4	3.6 LCG Support and Collaboration.....	13
2.2 Job Management	4	3.7 PPDG – IEPM-BW Collaboration	13
2.3 Production Grid Systems	5	3.8 PPDG- DOE Science Grid VO Management Project.....	13
2.4 Grid Analysis Tools (CS-11).....	5	4. Experiment Team Plans.....	15
2.5 Site AAA	5	4.1 Atlas.....	15
2.6 Web Services.....	6	4.2 BaBar	17
2.7 Executive Team.....	6	4.3 CMS.....	18
2.8 Collaboration with GriPhyN and iVDGL – Trillium.....	7	4.4 JLab	19
2.9 Joint JLAB-STAR Job Scheduling Project.....	8	4.5 D0	20
2.10 GGF Particle and Nuclear Physics Research Group	8	4.6 STAR.....	21
3. Computer Science Plans and Activities	10	4.7 PHENIX	22
		5. Funded Effort.....	22

1. PPDG, June 2003

This document gives the plans for the Particle Physics Data Grid SciDAC Collaboratory work for the period June 2003 to June 2004. During the past two years the collaboratory has made progress towards achieving its goals of introducing grid technologies into the participating experiments production distributed data processing systems and of increasing the share and quality of common grid technologies from the computer science and experiment groups. During the coming year the collaboratory will continue work towards these goals and also contribute to the transition to a service based grid infrastructure which is necessary to establish production experiment application grids based on a general grid infrastructure and technologies. The PPDG plan for 2003-2004 includes deliverables and milestones for each of the teams against which progress and success of the project will be tracked and coordinated.

The last two years have seen a significant change towards the acceptance of grids as a viable technology to address the experiments' distributed computing needs. In all the experiments participating in PPDG, developments and deployments of distributed computing and grid technologies involve many more people than those on the "PPDG team" per se. PPDG's dependence on external deliverables, the evolving requirements of the experiment collaborations, the increasingly broad responsibilities of the wide computer science groups and the need to communicate and collaborate both nationally and internationally with many outside projects and communities continue to challenge our ability to manage and coordinate success in the projects deliverables and goals. We will continue to try to rise to the challenge.

PPDG is focusing its efforts on

Data Management – file transfer and replication, storage, and storage management,

Job Management – job scheduling and execution, job descriptions and planning, and

Production Grid Systems – end-to-end application and system integration and deployments, together with those services needed to operation and run these systems – monitoring, errors handling and fault diagnosis.

Additional work will continue in

Authentication and Authorization - following on from the completed SiteAAA project through the efforts of the Site Security teams and collaborations with other grid projects especially the DOE Science Grid, Globus CAS, and the EDG/LCG security and developments groups,

Grid Analysis Environments – through a continuation of the "CS-11 Analysis Tools" working group and the development efforts within the experiment teams, and

Common Service Definitions and Standards – in meetings and working group discussions.

We are also collaborating with GriPhyN and iVDGL, the U.S. ATLAS and U.S. CMS software and computing projects, on a multiple-experiment functional demonstration grid for physics (Grid2003).

The goal of complete, robust and maintainable grid systems that deliver the performance and operating characteristics needed to meet petabyte scale particle physics experiments' data processing and analysis requirements is several years away from being fully realized. We hope to make additional progress towards this goal in the next year of the PPDG project and continue our understanding of what is needed to achieve it.

1.1 Project and Team Plans

In June 2003 each PPDG team submitted an individual plan for the third year of the SciDAC collaboratory, including identifying deliverables and milestones. The executive team held a phone meeting with each team to review and comment on these plans. The individual team plans area available from the PPDG web site:

Table 1. PPDG Team Plans 6/2003-6/2004

Team	URL
ATLAS	http://www.ppdg.net/docs/PPDGYear3/USATLASPPDGYear3-v2.doc
BaBar	http://www.ppdg.net/docs/PPDGYear3/BaBar-Year3-Plans.doc

Team	URL
CMS	http://www.ppdg.net/docs/PPDGYear3/CMS-Year3-Plans.doc
Condor	http://www.ppdg.net/docs/PPDGYear3/Condor-Year3-Plans.doc
D0	http://www.ppdg.net/docs/PPDGYear3/ppdg-d0-Year3-Plans-v3.pdf
Globus	http://www.ppdg.net/docs/PPDGYear3/Globus_year3_plan.pdf
JLAB	http://www.ppdg.net/docs/PPDGYear3/JLab-Team-Year3-Plans-2.doc
SRB	http://www.ppdg.net/docs/PPDGYear3/SDSC-Year3-Plans.doc
SRM	http://www.ppdg.net/docs/PPDGYear3/Team-Year3-Plans.SRM.doc
STAR	http://www.ppdg.net/docs/PPDGYear3/STAR-Year3-Plans.doc
Executive Team	http://www.ppdg.net/docs/PPDGYear3/ppdg%20year3%20mv%20plans.doc http://www.ppdg.net/docs/PPDGYear3/doesg-vo-ppdg-v2.doc
PHENIX	http://www.ppdg.net/docs/PPDGYear3/phx_ppdg.doc
STAR & JLAB Job Scheduling	http://www.ppdg.net/docs/PPDGYear3/Job Scheduling-Jlab-STAR.doc
US ATLAS & US CMS Grid2003	http://www.ppdg.net/docs/PPDGYear3/Grid3_v19.doc

This current planning document is drawn from these comprehensive Team Plans. We include only some sections of these plans here keeping as much as possible the original texts and wording.

During the year we will refer back to the individual plans when tracking and understanding the milestones and deliverables. In particular, most of the deliverables depend on other projects and people outside of PPDG. We have not included these dependencies in the overall PPDG plan. Most of the team plans also include specific issues and concerns that we propose to track and attempt to address during the year.

1.2 SciDAC Review

We have taken some specific steps to following the recommendations of the PPDG SciDAC reviewers in April 2003.

1. We are establishing closer collaboration with the DOE Science Grid project (SG), especially at LBNL and ANL. We see fruitful areas of increased collaboration in areas of support for authorization, monitoring, and operations. We are also inviting a Bill Johnston to be an additional liaison to the project.
2. We are making project and team plans that include milestones and manpower estimates.
3. We are focusing the project on a few specific and high priority areas: job management, data management, and production grid systems. Security and grid analysis environments also remain important areas of work.
4. The Executive Team is increasing its commitment in time and effort in order to follow up on project status and milestones.
5. We plan to define and document performance metrics and a process for software testing and release;
6. We are starting a longer term planning process for the project.

1.3 Follow up to Questionnaires

We are planning some specific steps to follow up on the results of the team questionnaires and phone meetings, which were held in preparation of the SciDAC review in the spring of 2003.

1. We will establish or restart cross-project coordination meetings on job management, data management and monitoring, and
2. We will encourage computer science presentations at and visits to experiment computing meetings, visits of Experiment team leads to computer science group meetings. We will continue the visits of the executive team to PPDG sites.

2. PPDG Cross-Project Plans

As described in the original proposal, PPDG projects are experiment and/or computer science team based. PPDG encourages projects that include multiple teams. The goals of PPDG include experiment use of common grid components and sharing and reuse of grid technologies. PPDG cross-project activities are organized to help meet these goals. The joint US ATLAS and US CMS Grid2003 project and the JLAB-STAR job scheduling project each involve two experiment teams to date - and thus slightly break the mold of this section which is on Cross-Project activities. Their successes will hopefully be an example for broadening the collaborations on these and other projects in the future.

2.1 Data Management

All experiments use Globus-GridFTP for wide area, high throughput file movement. A new GridFTP will be delivered in several phases over the next twelve months. PPDG will coordinate the testing and deployment of this new version across the experiments. Both the new and older versions of GridFTP will be included in releases of the VDT. PPDG experiments will test the new versions as they become available in the experiment test grids, before the stable release is made for deployment in their production systems. This is expected to start in Fall 2003 and last for six to nine months.

Nearly all the experiments have adopted the SRM specification as the interface to their storage systems. Effort to implement the specification for the different storage systems comes from both inside and outside of PPDG.

PPDG includes about six implementations of Replica Management systems across the experiment and computer science teams. An effort in the last six months to document the interfaces and capabilities, with a view to defining and encouraging movement towards commonality, is incomplete. Once the document is delivered the Executive Team and Steering Committee will decide on what, if any, further steps to make. A related activity is providing a Replica Registration Service (RRS) to various catalogs. A coordinated effort was initiated to determine whether this functionality can be embedded into another higher-level Replica Management Service, or whether it should exist as separate service. A spin-off of this effort has been increased attention to interoperability across the two implementations of the Replica Location Service, Globus-EDG and the EDG-LCG. It is hoped that one result may be more commitment to avoid such divergences in the future. Identifying the needs of the experiments for data management services in the short and longer term and ensuring delivery and adoption of appropriate common and reusable technologies, continues to be an area of concern for the Executive Team.

An emerging requirement is for collaborators at separate sites to be able to build and manage their own collection of experimental results. Each site implements a separate logical name space onto which they register files. Peer-to-peer federation technology is used to cross-register files between the logical name spaces. The RLS technology implements a single logical name space as a hierarchical structure that is distributed across sites. The integration of the RLS replica catalog logical name space with the peer-to-peer federated SRB MCAT catalog logical name space is under development within the UK data grid. The result will be the ability to manage a data collection independently of the grid, while registering files from the collection into a replication service for distributed data processing within the grid.

2.2 Job Management

All experiments are using or will use Condor-G, GRAM, and ClassAds for grid job scheduling and execution. Additional planning and scheduling components are being adopted, developed, or will soon be in use by the experiments. Now is the time to increase cross-project attention to this area of work to better understand and coordinate commonality of approaches, capabilities, and interfaces.

We are establishing a working group to provide coordination across PPDG job management activities and projects. It will provide a forum for all experiments to discuss their system needs, plan and present their developments and deployments, and for the computer science groups to present and promote their common solutions. The group will discuss requirements and strategies for the support of dynamic scheduling of jobs across the application grids based on experiment policies and will understand the requirements and use cases for such policies. The working group will ensure PPDG has good communication and collaboration with GriPhyN and iVDGL work in this area.

Job planning, optimization and scheduling services are being developed in the context of GriPhyN and other the experiment specific data analysis systems. PPDG follows these activities, contributes to the requirements as it has

time and participates in working meetings. As these new technologies are deployed in the experiment data handling systems PPDG effort contributes to the integration and hardening activities.

PPDG has an ongoing collaboration with the EDG Workflow Management project (WP1). We are exploring further joint activities in this area with the LCG (and hopefully the EGEE) to evolve the capabilities and to support the necessary level of interoperability across the European and US sites for each experiments data processing and analysis system.

2.3 Production Grid Systems

Deployment of end-to-end applications over the grid and the integration and operation of grid-based systems continues to be a PPDG priority. Several more years of focused effort are needed to deliver and operate sufficiently capable and robust grid systems to satisfy the experiment wide data handling and processing requirements. All PPDG experiments plan to extend their deployments and use of grid systems for their data processing and analysis needs during the period from June 2003 to June 2004. PPDG management will continue to work closely with the experiment groups to prioritize the work needed to integrate and support the experiment grids and to address experiment milestones.

Monitoring is included as part of PPDG production grid systems activities. PPDG effort is insufficient to support the necessary development of monitoring services. Most experiments are using Globus-MDS as the core information management framework. The experiments use Ganglia, Hawkeye, MonaLisa, etc. for monitoring clusters, across a grid, and for the display of monitoring information. We hope that the Glue Schema will evolve and be adopted as the common information model.

PPDG will continue to develop end-to-end solutions for troubleshooting and fault diagnosis in collaboration with individual experiment groups. These solutions will be tested on the experiment grids and useful components added to VDT.

PPDG is part of the Grid2003 Project, together with iVDGL, GriPhyN, US ATLAS and US CMS (www.ivdgl.org/grid3). Grid2003 is a major effort to integrate and deploy a multi-VO functional demonstration grid of moderate scale. If it is successful it is hoped it can remain as a semi-permanent grid infrastructure used by the U.S. LHC and other communities. PPDG is contributing effort to the core project team as well as working with the experiment and computer science application demonstrator groups to prepare and run their applications across this multi-organization grid infrastructure.

2.4 Grid Analysis Tools (CS-11)

This working group will continue to provide a forum for communication and collaboration across the experiment groups as they continue the prototyping and development of grid analysis environments for physics. This is an area where continued understanding of the requirements, the architectures, the differences between “production” and “analysis” must all be explored. It is anticipated that this group will increase its collaboration with the LCG, experiment and other application development projects in order to reduce duplicate and parallel efforts, and ensure interoperability of services. CS-11 will work with the LCG GAG group on requirements and use cases for grid analysis.

The working group provides the interface to the Tech-X phase 2 SBIR on grid-enabled JAS, the LCG data analysis activities and the ROOT/PROOF activities. David Alexander of Tech-X states "Tech-X will develop a well-documented release version of a fully implemented portal using the latest Globus Grid Service middleware and the latest version of the Java Analysis Studio. Tech-X will work with the CS-11 Interactive Analysis Working Group of the Particle Physics Data Grid to develop a standard set of interactive analysis service interfaces and implement them in the Grid portal."

2.5 Site AAA

The PPDG SiteAAA project was completed in January 2003. Since that time the group involved in the project has maintained some continued its activities in several areas: Participation in LCG Security Group; leadership and participation in GGF Security Working Groups; testing of existing and emerging Authorization technologies: CAS, VOMS etc; and increased collaboration with the DOE SG.

2.6 Web Services

The migration of GRID based and distributed computing to a Web Services model is happening rapidly. Several releases of the Globus Toolkit based on OGSA/OGSI are becoming stable and ready to use by PPDG applications sometime over the next year. PPDG teams have several application and application middleware projects which will be based on Web Services. As a project PPDG will increase its evaluation of and attention to this technology and standards. We hope to identify some cross project effort to work at the technical level on consistency, interoperability and reusability issues.

2.7 Executive Team

The Executive Team works together on the coordination and management of the project, reporting to the PIs and Steering Committee. The executive team plan of work includes activities to:

1. Arrange phone meetings, all hands, monthly steering and collaboration semi-annual meetings, workshops and collaborative meetings with other projects.
2. Continue quarterly reports, web server and email lists, as in previous years. Maintain and update the PPDG web pages and calendars.
3. Monitor and review the progress and status of the project plans and adjust proposed deliverables and milestones as necessary.
4. Collaborate with DOE Science Grid VO management project and Site AAA project. The PPDG registration authority service with the DOE Grid certificate authority will be distributed to have agents at each of the DOE labs in PPDG. The agents are people authorized to issue new certificates and will be primarily responsible for handling requests from their own user communities. Bob Cowles (SLAC) and Dane Skow (Fermilab) help the executive team in coordinating this activity with Doug Olson.
5. Coordinate work on interactive analysis tools interface with grid (CS-11). Joseph Perl (SLAC) helps the executive team in co-coordinating this activity with Doug Olson.
6. Coordinate communication and discussion of work on job scheduling and management. Peter Couvares (Condor) will help the executive team in coordinating this activity with Miron Livny.
7. Coordinate work on data and storage management. Don Petravick (Fermilab), Reagan Moore (SDSC), Jennifer Schopf (Globus), and Arie Shoshani (SRM) help the executive team in coordinating this activity.
8. Co-chair and participate in Joint Technical Board, Trillium and U.S.-EDT/LDG/DataTAG Joint Projects. Mike Wilde, GriPhyN coordinator, Rob Gardner, iVDGL coordinator, and Alain Roy, VDT team lead, help the executive team in these activities.
9. Communicate and collaborate on joint Monitoring Projects (e.g. GLUE). Les Cottrell (SLAC), Jenny Schopf (Globus) and Brian Tierney (LBNL) help the executive team in these activities.
10. Provide interface and communications path with the phase 2 SBIR from SHAI, Inc. lead by Charles Earl on Agent Based High Availability computing. This project is motivated by the need to develop more robust and fault-tolerant distributed computing. Doug Olson (LBNL) and Arie Shoshani (LBNL) are the PPDG contacts with this project.
11. Develop a software testing and quality assurance plan.
12. Introduce and track performance metrics. The Team Leads will help the executive team in this activity.
13. Participate as members of the LCG project committees. The US ATLAS and US CMS software and computing managers, and the Globus team lead, as well as the executive team, participate as members of these committees.

Table 2. Executive Team Milestones

Date	Work Item(s)	Comments
8/2003	6	Start monthly job management coordination meetings.

9/2003	9 7	Restart monthly monitoring coordination meetings. Start monthly data management coordination meetings.
11/2003	3	Review progress of team plans and provide summary of status and proposed adjustments in milestones and deliverables.
12/2003	1, 12	Semi-annual PPDG collaboration meeting. Face to Face Steering Meeting. Agree on project and team performance metrics. Agree on process for integration of middleware into the experiments' software stack.
1/2004	11	Publish PPDG documents on software testing, quality assurance and middleware-application integration.
3/2004	3	Review progress of team plans and provide summary of status and proposed adjustments in milestones and deliverables.
3/2004	1	Decide on annual cross-project documentation or review process, as the next in the series of: questionnaires in 2003; project reviews in 2003; requirements documentation in 2001. Input to the GGF PNPA document on requirements of the HENP community to the GGF.
6/2004	1	Semi-annual PPDG collaboration meeting. Review of 2003-2004 performance of the project. GGF PNPA document summarizing GGF standards and recommended practices that could be adopted by the HENP community.
7/2004	3	Review progress of team plans and provide summary of status and proposed adjustments in milestones and deliverables.
8/2004		Informational document (or series of documents) that relates the HENP experiences to date to the goals and scope of the GGF,

2.8 Collaboration with GriPhyN and iVDGL – Trillium

The Trillium collaboration looks like it will continue to grow. Each of the Globus, Condor, US ATLAS and US CMS grid development, deployment and support teams frequently work as a single organization across PPDG, GriPhyN and iVDGL. The team leads and management of the three projects have significant overlap.

Joint projects are continuing for GLUE, Glue Schema, Monitoring, and VDT testing not only within Trillium, but also with the EU grid projects.

The PPDG team plans include deliverables for the Grid2003, (<http://www.ivdgl.org/grid3>) joint project in collaboration with the US ATLAS and US CMS software and computing projects, GriPhyN and iVDGL.

Some PPDG deliverables include deployment of GriPhyN technologies such as Chimera. As this occurs we will work more closely with GriPhyN on the evolution and support of these components.

Table 3. Grid3 Major Milestones

Date	Comments
8/2003	Grid2003 grid with four sites.
8/18/2003	Grid2003 Integration Week.

11/2003	SC2003 demonstrations.
---------	------------------------

2.9 Joint JLAB-STAR Job Scheduling Project

STAR and Jefferson Lab are beginning a prototyping effort of a Web service-based meta-scheduler, having already deployed SRMs and other XML technology. These two groups propose to develop a common batch web service interface definition, analogous to the SRM data management interface specification, as part of their respective Year 3 PPDG activities. The prototype implementations will immediately serve the respective communities, and will serve as valuable input to a larger community effort. Since both sites are currently engaged in XML and web services related batch developments, the JLab and STAR team therefore propose to work together to achieve this goal in a few steps.

1. Evaluation of the activities of the partnering team, with a view towards developing a common understanding of the problem to be solved
2. Work together on the definition and strengthening of a high level user Job Description Language (U-JDL) as needed to describe both multiple parallel-jobs (lattice calculations etc...) and large numbers of serialized jobs (statistically driven data mining). This work will be presented to PPDG as an input to a baseline for a high level job description language.
3. Use of Web services as the implementation / interoperability framework for a computational grid (loosely coupled, hence very scalable). A component design will be needed before the migration phase.
4. Use of SRM as an example of a successful approach via interfaces (gaining wide support) for dynamic data management approach

Table 4. JLAB-STAR Joint Project Milestones

Date	Work Item(s)	Comments
10/2003	2	Deliver the requirements and definition for the U-JDL for batch job description; gather feedback from the PPDG collaboration
11/2003		Design implementation completed and first version of WSDL for batch web service
12/2003	3	First implementations of site (not meta) interface, to map WSDL, XML onto PBS, LSF or Condor submissions.
3/2004		First implementation of meta scheduler.
5/2004	4	More robust implementation, including file location using Replica Catalog interface, file and/or job migration using SRM, output publication using SRM and Replica Catalog, (at least rudimentary) load & file location based dispatch (support for) and job accounting information going to a distributed database (Web Service).

2.10 GGF Particle and Nuclear Physics Research Group

The plan is for the Particle and Nuclear Physics Applications Research Group to provide a forum for discussion of issues related to particle and nuclear physics applications and production grids. The research group charter is currently under consideration by the GGF executive/steering committee.

There are three specific goals, all with the overarching aim of ensuring that HENP actively participates in setting requirements and defining standards to ensure that its needs are met and to bring its experiences in deploying and using large scale grids to the grid community as a whole.

- To bring the requirements of the HENP community to the GGF in order to explain and inform the wider grid community of the specific needs and issues of HENP

- To ensure participation of the HENP community in grid standardization efforts, particularly in those areas and services essential to the successful use of grids in HENP.
- To provide early feedback to GGF technical working groups on the success or failure of various grid software components as used in high performance production activities by HENP experiments.

The PNPA Research Group proposes to deliver: An informational document that defines the requirements of the HENP community at this time from GGF WG and RGs - Initial draft Mar 2004; an initial informational document directed to grid designers and users in the HENP community summarizing GGF standards and recommended practices that could be adopted to HENP computing problems – draft summer 2004; and an initial informational document (or series of documents) that relates the Application Area experiences to date to the goals and scope of the GGF – August 2004.

-

3. Computer Science Plans and Activities

All the experiments are already using Globus and Condor technologies. This means that these two computer science groups need to spend more effort in support and experiment specific integration. Given their limited effort available in PPDG these groups do not have defined deliverables and milestones – their activities are spread across all teams and will be prioritized by the executive team working with the Globus and Condor team leads.

Nearly all the experiments have adopted the SRM specification as the interface to their storage systems. Effort to implement the specification for the different storage systems comes from both inside and outside of PPDG.

3.1 Condor

The Condor PPDG team will continue to work with each experiment on the planning, design, deployment and support of their end-to-end applications. Development work will be done as needed in Condor-G, Condor, DAGMAN, Hawkeye, ClassAds, Computing on Demand (COD) and Data Placement (Stork) – both within the PPDG funded effort and the using other development efforts. The PPDG Condor Team will help experiments migrate to the use of Virtual Data Toolkit, already used by ATLAS, CMS and the LCG project, and start to contribute to the support and development of the VDT. The Team will extend and develop diagnostic, error reporting and fault handling tools in collaboration with first a single experiment, and then deploying these in other experiments as requested.

Members of the team will be assigned as needed to help specific experiments develop and deploy the production and analysis grid systems. The assignments will be prioritized to meet experiment wide milestones for grid based data processing.

3.2 Globus

Globus continues interactions in terms of coordination and support of the PPDG applications included weekly phone meetings and email lists for Atlas and CMS, following the grid emails lists of D0, and providing support for the Argonne-Chicago ATLAS team in their efforts to perform "data challenge on demand" event generation using VDT, RLS, and Chimera. Support through the discuss lists and Bugzilla are available to all experiments. Other work covers AAA and CAS, GridFTP, RLS, Due to manpower limitations no work is anticipated in support of the grid analysis, and work on a follow-on document to the Data Grid Reference Architecture that focuses on the approaches to and mechanisms for late planning will continue on this effort as time allows.

1. **Globus Toolkit 2.x:** updates and bug fixes : Bugzilla bugs will continue to be answered, but few additional code changes will be implemented on the 2.x code base.
2. **Globus Toolkit 3.0 :** Discussions need to be held with the various experiments to better understand the needs and requirements, and timelines, for use of the GT3.x code base.
3. **Site-AAA/VOMS;** Waiting for feedback from CMS. Once we have this, will evaluate the alpha version and make plans for incorporating this in the next release.
4. **GridFTP:** Current plans include a beta release of the new server as part of the GT3.2 Beta. The only guaranteed features at that point are data transport and security. Additional features will be released in the November time frame. This will include striping for sure, but no other features are guaranteed. A better evaluation from the experiments is needed, as well be users to evaluate the new code. We will also deliver a version of GridFTP for HPSS with their 5.2b release in late Q1 or early Q2 of 2004. The use of RFT by the PPDG experiments needs to be evaluated.
5. **Monitoring and MDS work:** Currently, discussions with iVDGL and DOE Science Grid for suggested MDS-2 deployments in those project are taking place, and this is likely to have a positive impact for the PPDG experiments in the near future by clarifying the data available and deployment issues. GLUE schema work is also continuing.
6. **Replica Location Service:** RLS will be modified to be an OGSA service. Feedback from the experiments will play a large roll in this work. This work will occur outside of PPDG and the support issues need some clarification since PPDG funding is not available.

7. **CAS** : Work has started on integrating CAS as a production service along with a CAS-enabled GridFTP server in the main Globus Toolkit release. It is unclear at this time if this work will appear in GT3.0 or the subsequent release.

3.3 Storage Resource Broker

The dominant computer science development efforts that have been requested by BaBar / PPDG are:

- Peer-to-peer federation of collection catalogs
- Specification of dynamic consistency constraints for access to collections
- Integration with Grid services
- Integration of data management systems with storage management systems (SRM)

The level of effort for the development of these capabilities substantially exceeds the support from PPDG. Additional funding sources have been found to support the first three activities.

To promote production use of the technology, an aggressive release schedule is being followed. Major releases are made on a yearly basis (the last major release was made in February, 2003, and included features that had been requested by BaBar), minor releases are made quarterly (the last minor release was made at the beginning of June, 2003), and bug fix releases are made within a week of identifying and fixing a problem. The goal of the SRB team is to have no outstanding bugs.

The success in having the SRB used in production within BaBar will depend upon the ability to resolve each of the BaBar operational concerns. These are discussed bi-weekly, with either revised code provided to BaBar, or a joint development of the requirements for a new feature. We will negotiate a set of system requirements for collection federation with the BaBar experiment and the UK data grid (which is also working with the CMS experiment). SDSC will build a peer-to-peer federation mechanism for linking independent SRB metadata catalogs. The effort to specify and manage consistency constraints is a multiple year effort, which is expected to take through FY05. The development of OGSA interfaces to the SRB collection is being supported through funding from other federal agencies/projects. The total level of effort devoted to the creation of OGSA interfaces is 2 FTEs within the SRB group.

Table 5. SRB Milestones

Date	Comments
10/2003	Requirements for BaBar production
12/2003	Initial implementation and deployment for BaBar production needs.
6/2004	Second implementation and deployment for BaBar production needs

3.4 Storage Resource Management (SRM)

The activities we plan to embark on in year 3 include activities relating to the specification of SRM v2.0, the development of DRM and HRM that are compatible with SRM v2.0, and the development of a Replica Registration Service (initially for the STAR project). The main action item we want from PPDG as a project is an opportunity to apply our technology - SRMs to experiments. This includes the use of SRMs as part of the request planning and request execution process. In addition, we'd like to interact with monitoring activities in order to provide current information on storage resources. The following tasks are planned:

1. **Complete the SRM v2.0 specification.** This task involved coordination of multiple institutions, and therefore requires a great deal of communication, discussion, and negotiations in order to reach agreement. - SRM v2.0 specification - standards document
2. **Developing a WSDL-based wrapper for DRM and HRM v1.0.**- WSDL-based wrapper - software

3. **Packaging DRM v1.0 for inclusion in VDT.** At a later time, a version that will work with NeST will be made available. - Packaging DRM v1.0 for VDT - documentation, user guide, software packaging
4. **Packaging of the DataMover for inclusion in VDT.** The DataMover is a command-line interface software that provides the capability to request File Replication from a given directory – the mirror directory is generated at the target. This version will work with any SRM v1.0 - Packaging DataMover for VDT - documentation, user guide, software packaging
5. **Develop Replica Registration Service.** Initially, it will be developed for the STAR project to write into their File Catalog. Later it will be developed to write into RLS.
6. **Development of DRM and HRM v2.0.** This will require a completer redesign and implementation of DRM and most of HRM, in order to support the new functionality. - Development of DRM and HRM v2.0 software.
7. **Integration of DRM v2.0 with NeST.** This version of DRM will take advantage of NeST’s lot to provide flexible dynamic space allocation and file assignment into spaces. We expect this development to require close coordination between the corresponding teams.

Table 6. SRM Milestones

Date	Work item(s)	Comments
10/2003	1	
12/2003	2	
3/2004	3	
2/2004-6/2004	4	Replica Registration Service, for STAR – Feb 2004, for RLS – June 2004
7/2004-9/2004	5,6	DRM and HRM v2.0 – prototype – July 2004, debugged version – Sept 2004
9/2004	7	

3.5 Virtual Data Toolkit (VDT)

All PPDG experiments using Globus-GT2 will use VDT releases. The experiments provide additional testing of the VDT development and pre-production releases in their application and grid systems. The experiments will contribute to the testing of VDT versions which include GT2 or GT3. The transition to the support and use of VDT for GT3 will be decided as needed.

The Condor Team will work with STAR, JLAB, BaBar, and D0/CDF to migrate these experiments to include the use of VDT in their grid deployments. We will work with the VDT team to add necessary capabilities in support of this. This is expected to include: 1) Understanding of packaging and distribution needs of the experiments. Working with the experiments to ensure the VDT packaging and distribution supports their needs. 2) Support for multiple versions of VDT at any time. This set of PPDG experiments are in data taking mode and will have different schedules to be able to move to new versions of the toolkit and/or its components. 3) Addition of other configurations and/or packages needed by these experiments.

Table 7. Adoption of VDT

Date	Comments
11/2003	VDT based on GT3
11/2003	D0 using VDT
11/2003	BaBar using VDT
1/2004	STAR Using VDT

3.6 LCG Support and Collaboration

The European physics grid middleware is based on the same technologies as the U.S. physics grid projects. Globus, Condor, SRM and SRB technologies are all used by some or all of the EDG, LCG and experiment-wide distributed applications. In particular the EDG and LCG releases are based on the Virtual Data Toolkit (VDT). The experiments on PPDG have European collaborators who are working with and using the EDG and LCG grid deployments. It is a PPDG strategy to work across the collaborations as needed to facilitate and ensure interoperability between the U.S. and EU participants in the experiment grid projects.

PPDG will work with the VDT team on the support and deliverables for the LCG project in order to continue to allow interoperability between the experiment deployments in the U.S. and Europe. The PPDG computer science groups contribute to the evolution and support of software in the VDT and other software components used in LCG. The milestones for this work are tied to the LCG-1 experiment data challenges in 2004.

PPDG will also work with the new EGEE project as it emerges to understand areas of collaboration, interoperability needs etc.

3.7 PPDG – IEPM-BW Collaboration

The close collaboration between the IEPM-BW project and PPDG will continue. Less Cottrell and Warren Smith are contributing to the Glue Schema Network Element definition and implementation project; Network monitoring, analysis and display tools are being increasingly shared and deployed in the experiment grid systems.

3.8 PPDG- DOE Science Grid VO Management Project

This project is comprised of a set of activities and services developed and operated by DOE Science Grid, ESnet PKI project in conjunction with the computing sites of PDSF/NERSC, RCF/BNL, and JLab computer center. These services are being developed and run primarily on behalf of the nuclear physics experiments in PPDG. This is not to exclude the HEP experiments but they have largely established their VO management plans and projects and this project will communicate and interact with those as appropriate. Virtual Organization (VO) management generally consists of maintaining the membership lists of users belonging to a particular VO, their group memberships and roles in that VO, and a mapping of that information onto the local site authentication (authN) & authorization (authZ) systems for the computing resources that the VO is entitled to use.

The elements of this VO management project are the following work items:

1. **VO membership service** holding authoritative membership information and ability to generate proxy credentials holding VO group/subgroup/role membership information, based on CAS server
 - development and prototype service by DOESG (Von, Dhiva, Doug)
 - interoperability of proxy credentials with LCG/US-CMS/US-ATLAS VO service
2. **Gatekeeper with callout** to local authZ running at PDSF, RCF, JLab. The local authZ will have DN/VO membership – UID/GID mapping.
 - Steve Chan, Shane Canon LBNL, Jerome Lauret, Dantong Yu, Rich Baker – BNL, Chip Watson., Bob Lukens, Sandy Philpott – JLab
3. **Local site user account/registration database** includes X509 DN and experiment/VO membership
 - each site does this (NERSC, RCF/ACF/BNL, JLab)
4. **MyProxy with credential repository** (long term user’s private key storage)
 - Tony G.
5. **Document describing VO management practices** at each of the PPDG and DOESG labs, ANL, BNL, FNAL, LBNL, ORNL, PNNL, SLAC
 - Doug Olson, Keith Jackson (might contribute to GGF)

The detailed plans for items 1-4 will be developed as part of item 5.

Table 8. PPDG- DOE Science Grid VO Management Project

Date	Work item(s)	Comments

Date	Work item(s)	Comments
10/2003	5	
12/2003	4	
3/2004	1,2,3	item 3 is done or nearly done at each site already but included here since 2 depends upon it
3/2004-6/2004	1-4	operate VO service in pre-production mode (no SLA) and develop operations plan
7/2004	1-4	deploy production VO service

4. Experiment Team Plans

The text in this section is mostly selected from the individual team plans. Each section has a slightly different scope and focus. We refer you to the complete plans for more information and clarification.

All teams on PPDG are dependent on deliverables and components from the broader experiment application and infrastructure groups, and are also dependent on the support and modifications of grid components from the PPDG and/or other external computer science groups. It is understood that most deliverables will be effort limited. We list some few additional dependencies of these deliverables that need to be tracked during the coming year: **D0**: UK GridPP effort especially through RunJob and other contributions they make to the JIM deployment and testing; **STAR**: Deployment of the infrastructure and tools for delivering certificates for BNL teams and Condor-G web service developments; **US ATLAS and US CMS**: delivery of Grid2003.

PPDG's mission is to encourage and foster joint projects across the experiment teams. The PPDG experiment joint projects are: JLAB and STAR for VO management and Job Scheduling; US ATLAS, US CMS for Grid3 and VO management

4.1 Atlas

The current efforts of US ATLAS on PPDG focus on the following areas:

- **File and dataset cataloging, replication**, and metadata for ATLAS (Magda)
- **VO management** – a Grid User Management System (GUMS)
- **Grid monitoring** – Resource monitoring with Ganglia and MDS
- **Interoperability with EDG-based middleware testbeds**
- **Distributed analysis** of large datasets (DIAL)
- **End-to-End testing and evaluation of grid component environments** (GCE) which employ a number of grid software technologies, including, Chimera/Pegasus (Virtual Data System), RLS, VDT

PreDC2 will drive coordination of development of these projects to deploy in the Grid2003 environment, a joint project with PPDG, iVDGL, CERN (Tier 0), and the US CMS and US ATLAS Software and Computing projects. Experience and development from PreDC2 on Grid2003 will be extended to Data Challenge 2, which begins in Spring '04.

The PreDC2 data flow follows ATLAS mandates for computing model testing. The prototype Tier 2 centers will be used for event generation and simulation of Supersymmetry data. These data are to be cached at the Tier 1 where they undergo validation and quality assurance. Validated data are transferred to the CERN Tier 0 where they will be reconstructed. Reconstructed data will be transferred to the Tier 1 site for analysis exercises on the Grid. Details regarding data discovery, replication, and task formation and cataloging are still under discussion in both US ATLAS and international ATLAS. We intend to exercise prototype distributed analysis tools such as DIAL and Ganga.

The ATLAS persistency mechanism depends on POOL, which uses the CERN RLS for its cataloging mechanism. This introduces a number of issues regarding use of RLS for Magda and the GCE (Chimera/Pegasus) environment which uses the Globus-EDG RLS version, which uses different interfaces. Discussions on how to resolve the issue are underway. The timescale for having a solution in place is currently estimated to be the end of the year. After this time we anticipate a new development plan for incorporating RLS capabilities within Magda. Magda continues to supply functionality for ATLAS not yet provided by any of the available middleware packages, including VDT.

ATLAS is participating as part of the joint Grid2003 project, thus gaining greater coherency in monitoring and VO management with other projects, and providing testing of the Distributed Analysis services in DIAL, with the Athena framework as the primary analysis tool. Testing of Grid3 will extend the tests of Chimera, RLS and VDT already in progress.

The ATLAS PPDG Year 3 deliverables will support the following:

1. Production system for simulation using Grid3
2. Staging/scheduling system for Tier1-Tier 0 connection

3. Coordinated system for reconstruction at Tier 0
4. Analysis tools that support a three-group analysis effort.

Table 9. ATLAS Team Milestones

Date	Work item(s)
11/2003	1
2/2004	4
3/2004	2,3

These activities have components and schedule as part of the US ATLAS software and computing plan as follows:

- 18-Jun-2003 Packaging, testing, validation of the GCE environment (Chimera, RLS, VDT, and various server and client libraries) to support ATLAS persistent grid challenges
- 27-Jun-2003 Grid2003 task force report
- 01-Jul-2003 Support for GCE 0.1 release
- 16-Jul-2003 Joint testing of the latest Chimera functionality for discovery and reporting of compute pool services prior to application execution (Pegasus MDS integration).
- 22-Jul-2003 DIAL – ATLAS combined ntuple with PAW as application as demonstration.
- 22-Jul-2003 Full definition of metrics, and goals for PreDC2
- 01-Aug-2003 Support for GCE 0.5 release: integrated MDS, initial workflow monitoring.
- 01-Sep-2003 Support for GCE 1.0 release which includes the below:
- 01-Sep-2003 Joint development with Chimera developers on interoperability extensions to support both US and EDG Grids.
- 01-Sep-2003 Reconstruction application running at Tier 0
- 01-Sep-2003 Simulation applications running at Tier 2's
- 01-Sep-2003 Staging, file transfer exercise for Tier 1- Tier 0
- 15-Sep-2003 Prototype version of DIAL with Athena as application, simple scheduler, POOL as data management scheme.
- 15-Sep-2003 Production interface with error recovery for simulation-reconstruction chain. (automated)
- 25-Sep-2003 Definition of Data Challenge 2 goals
- 01-Oct-2003 Full site installation and testing of tools for exercise of entire simulation-reconstruction-analysis chain
- 12-Oct-2003 Support for delivery of GCE 1.5 for Grid3/PreDC2
- 1-Nov-2003 Run full chain of simulation through analysis begins
- 15-Dec-2003 End of full chain of simulation test
- 03-Jan-2004 Definition of tools for Data Challenge 2
- Winter 04 Definition of analysis tools for results of Data Challenge 2, linkage with other distributed analysis efforts (e.g. DAWN, if funded)
- 01-Feb-04 Support for delivery of GCE 2.0: DC2 Alpha
- 01-Mar-04 Support for delivery of GCE 2.0: DC2

- Spring 04 Run data challenge 2

4.2 BaBar

The principal activity of the BaBar PPDG team at SLAC is collaboration with the SDSC PPDG team and with collaborators at BaBar's European "Tier-A" computing centers to deploy production data distribution based on the SDSC Storage Resource Broker, SRB. Continued work with the PPDG Site-AAA team to deploy workable Grid security solutions for BaBar is an important subsidiary activity. This work forms part of a larger BaBar Grid deployment program focused on the deployment of both data distribution and distributed job submission.

1. **Limited Deployment:** objectivity collections, limited management of data distribution, working only for ccin2p3 who will only use part of the system – essentially the metadata catalog. Automated loading of metadata. - Wilko, Adil, Cristina (SLAC), Liliana, Jean-Yves (ccin2p3)
2. **Phased integration with the existing JImport tools** replacing the pieces that perform crude SRB-like functions by SRB. With simulated p2p. - Wilko, Adil (poss) (SLAC), Anne-Marie, Jean-Yves, Liliana (France)
3. **SRB-based distribution of ROOT files** (depends very much on what the new BaBar bookkeeping system provides and how easy they make it to implement Grid tools. The new bookkeeping system is being designed at present). Finding files and loading into SRB and relating to the Bookkeeping. Have to make sure that bookkeeping have SRB in the front of their mind now. There is interest in collaboration with CMS to bring in more European effort. - People: Wilko, Jean-Yves, Liliana, Adil
4. **Testing of new SRB with peer-to-peer support.** Test new code, check functionality of reading from 1 Mcat writing to another. Test scaling. - People: Wilko, Adil, Liliana, Jean-Yves.
5. **SRB with VSC.** If we use the VSC then we need to get SRB to provide something on top of the MCAT to allow authentication. - People: SRB folks 2 people
6. **Production quality** (we should count this time from Deployment). All the scripts to monitor SRB servers, check consistency of catalog with what is transferred, scripts to bulk load, automatically register new users, manage disk cache space, staging requests. This latter part could possibly be ~6 months work. Also, scripts and apps to do deep-copying of selected events. Monitoring of transfer rates, load etc. Integration with LCG s/w. Interact with RLS or maybe use the federated MCAT. Allow jobs to query the SRB either directly or through BaBar bookkeeping tools. Mainly for data dist. - People: SRB + ~4 people
7. **Web services.** We need to understand how to integrate SRB into Web services. Very unclear the timescale or what's needed/required. Mostly exploratory initially. Low priority exploratory work will start now (July 2003) - People: ~2-3 people

Table 10. BaBar Team Milestones

Date	Work item(s)
7/2003	1
9/2003	2
3/2004	3
10/2003	4
3/2004	5
8/2003-2005	6
6/2004	7

4.3 CMS

The US CMS PPDG Team will continue to upgrade and support the US CMS Testbed infrastructure and software to deliver worth to CMS simulation production and analysis. Areas of work will include MOP, MCRUNJOB, and GRID2003. The major part of the Caltech effort will be concentrated on Grid Analysis Environment (GAE) work. The PPDG GAE deliverables are part of the GAE project at Caltech and depend on the other members of the group and the CMS collaboration. The Clarens deliverables are a collaboration with US CMS GriPhyN, MIT Proof and PPDG Analysis Tools efforts. All members of the CMS PPDG effort work as part of the US CMS software and computing project teams and help with the use and support of grids in the experiment.

US CMS Production Grid and Grid2003: To contribute to the experiment grid deliverables and deployments for pre challenge production, Grid2003 and DC04. Collaboration on necessary additional development to allow integration and deployment of experiments applications on the grid. Work with the Condor and SRM computer science groups on the use and testing of the middleware used for the CMS grid.

1. Working grid for CMS production as needed. This work is clearly only a small percentage of the total effort on these projects and is dependent on the delivery of the other components necessary for Grid2003 and DC04.

Grid Analysis Environment: To continue Grid Analysis Environment (GAE) efforts, including the GAE architecture, Clarens, the Distributed Heterogeneous Data Warehouses for Analysis (DHDW) and SOCATS and various GAE prototypes. The plan is to have some early users, including Rick Wilkinson, working on core software and/or PRS and/or Testbeam analysis. The following deliverables are those that are mainly the responsibility of the PPDG effort in the GAE project. However, they all rely on contributions from the other members of the project at Caltech and CERN. All durations assume 7/15/03 as the start date.

2. Analysis Object Converters: Continue development of ROOT2SQL (3 months); Complete development of FZ2SQL (3 months). User GUIs for ROOT schema management (6 months). Allow sub-selections of data – GUI (6 months). Complete SQL2SQL project (this fits with distributed Grid databases) (6 months).

3. Grid Analysis Environment Prototypes: Telecom 2003 (October), Supercomputing 2003 (November), World Forum (December), demonstrations. Handheld JAS – WIRED (Anjum, Bunn) (3 months). Demonstration SOCATS integration in Clarens (4 months). IGUANA client exchanging OI twigs (DC04) (6 months). Develop GAE Collaboration Desktop, including the above, Grid Views (MonALISA) and Interaction, and persistent collaboration.

4. Clarens <http://clarens.sourceforge.net/>: Continuation of existing work (support and incremental extensions). Increase P2P functionality. Target: evolve to fully distributed, managed, P2P system (6 months). Provide authorization and authentication for other services (e.g. DIAL, ORCA/COBRA/PROOF) (3 months). Develop stateful connection interface for SRB, DIAL, dCache (9 months). Flesh out CAIGEE system architecture e.g. interoperability with GriPhyN VDT (4 months – ready for Supercomputing). Develop service that Provides MonALISA monitoring information (1 year). Expand job submission services (e.g. cluster scheduler/ and PROOF) (1 year). Extend Javascript Web Browser interface (ongoing as new services are enabled). Make Java version of the Clarens server (NUST – 1 year).

5. Distributed Grid Databases: Deploy example physics analysis datasets in distributed RDBMS. Develop Web Service GUIs for data selection. Demonstrate distributed queries and aggregation of result sets from heterogeneous RDBMS (6 months).

6. STL Optimized Caching and Transport System (SOCATS): SOCATS is a general-purpose tool to deliver large SQL result sets in a binary optimized form. Create automated test scripts for SOCATS (3 months). Switch to Clarens RPC layer, from gSOAP (3 months). SQL Query storage in Chimera/VDT (6 months). Develop Java language binding (6 months). Investigate other bindings like CORBA, DCOM, .NET Integration with MonALISA and New services. Plan, gather, transport and monitor the processing of distributed object collections.

Table 11. CMS Team Milestones

Date	Work item(s)	Comments
------	--------------	----------

11/2003	1	working Grid2003 application demonstrator for SC2003
2/2004	1	CMS grid ready for DC04
2/2004	2,3	
4/2004-6/2004	4	Mar 2004 - p2p functionality, interface to SRB, Dial, dCache, etc. July 2004 – use cluster scheduler and proof, java implementation.
3/2004	5, 6	6 is ongoing beyond this timeframe.

4.4 JLab

The highest grid priority for the laboratory for both the experimental program and the theory program is a robust data grid. PPDG has allowed Jefferson Lab to participate in the Storage Resource Manager (SRM) project, and indeed Jefferson Lab has been a key contributor of ideas and prototyping for this data grid component. Expanding this effort into related data grid components (Replica Catalog, Meta-data Catalog) is a high priority for the coming year. This data grid, consisting of interoperable implementations of community-standardized interfaces, will support the movement of data analysis products and simulation products (detector and theory) alike. The PPDG effort at Jefferson Lab is a collaborative effort between the High Performance Computing group, whose main customer is currently Lattice QCD, and the lab Computer Center, which supports the experimental program. PPDG resources augment the programmatic efforts being expended towards their distributed computing goals, and provide for valuable interactions with others within NP and HEP doing similar work. The Jefferson Lab SRM implementation is being used to build a grid for the CLAS (Hall B) collaboration. The currently deployed system (Florida State University, University of Glasgow, Carnegie Mellon University, Old Dominion University, and University of Regina) is based upon the version 1 specification, and is being updated to reflect the enhancements within version 2.

Deliverables will be interface definitions and production level components for use in building a web services data analysis & simulation grid for the CLAS Experiment Grid (~dozen sites), as a prototype for Hall D (GlueX), LQCD (SciDAC) Production Meta-Facility and ILDG.

The components that we would anticipate being developed within the context of PPDG include:

1. **SRM improvements.**
2. **ReplicaCatalog web service interface;** migration of current JLab SQL implementation to this interface.
3. **Batch web service;** integration with Auger (batch interface, modified to be a client of web services for grid deployment). This is a joint project with the STAR job management and scheduling project.

Components which are more likely to be developed more within the LQCD context include:

4. **LQCD Meta-data catalog web service interface;** JLab implementation (or use one from UK or KEK).
5. **Performance metrics.** Philosophically, have a meta-facility which behaves like a local computer center. Similar percent utilization (above 90%), with data transfers limited only by WAN bandwidth, and possibly remote exploitation of resources similarly limited (but distributed simulation can probably prevent this from being a bottleneck).

Table 12. JLAB Team Milestones

Date	Work item(s)	Comments
8/2003	1	
10/2003	2	
1/2004	4	Prototype LQCD meta data catalog

4/2004	3	
6/2004	4	Prototype meta-data facility

4.5 D0

During the third year of the collaboration, we plan to stabilize JIM v1, expose it to high operational load. We will define and deliver Version 2 of JIM, as well as prototype an operational support model for the (remote) SAMGrid installations. Planned general deliverables for SAM/GRID:

- Devise a viable model for job input and output transfers of executables, associated flat files other than collider data (e.g., control files and calibrations), log and error outputs from the consumer, other consumer outputs which do not fit the model of data tiers
- Support Monte Carlo (MC) job distribution via the brokering service at a level of sophistication which replaces the current manual job distribution (i.e., uses the capacities of the production centers and implements the job priorities as input to the Monte Carlo request system
- Supply a monitoring and information service for servers and jobs, which will include better mechanisms for filtering and archiving than the current flat file logs, and allow retrieval of basic performance metrics
- Supply functional ‘virtual organization’ management system, either as further development of our current tools, or adoption of new technologies as they are available.
- Support reconstruction/analysis job distribution at the level of current functionality (manual choice of site), but making use of the job input/output model above
- Begin to understand brokering issues for reconstruction and analysis jobs.

The following developments will have to take place in JIM:

1. **Resource description for simulation (MC) environments** has to be frozen and published, using the JIM-adopted Condor ClassAd framework. Presently, a lot of ideas have been circulated and a prototypical implementation is available. The same applies to reconstruction.
2. **The interfaces and services between the Grid and the Fabric has to be clearly defined.** One of the novelties of JIM’s design is the use of the XML database on the border between the Grid and the Fabric. We will need to expose this topic to a broader Grid community so as to stimulate discussions and possibly standardize the interface between the Fabric and the Grid.
3. **The XML-based logger has to be fully developed.** Unlike the Logging and Book-keeping service from EDG WP1, our logging service is an important concept which will underlie historical data mining for both data transfers and jobs, and thus provide accounting at various levels.
4. **Evaluation of Web services** must also be complete by the end of year 3. While any particular grid-like system like SAM can do without web services, it is becoming increasingly clear that it is impossible to combine diversity in grid solutions with grid interoperability without a common language to describe services, and WSDL is the *de facto* such leader. Thus, unless the Grid community in general and Run II experiments in particular want to entrench into middleware consisting of proprietary (GTK 2) implementations of obscure, often criticized protocols (GRAM, GridFTP), we must arm ourselves with proper generalizations and describe our system in a language like WSDL
5. Perform stress-testing of Condor-G and feed back to the Condor team any issues. Work with the Condor team in the following areas of Condor-G development: 1) error reporting, 2) stability, 3) bug fixes, 4) structured jobs, 5) consistency in behavior, 5) better tools for analyzing problems, 6) confirmation that the Match Making benefits resource usage optimization.

Table 13. D0 Team Milestones

Date	Comments
10/2003	1) Finish installation of JIM at initial sites; 2) Begin load testing; 3) Exercise Monte Carlo operation; 4) Resolve issues with user input and output sandbox management; 5) Evaluate

Date	Comments
	VO management options; 6) improved reliability, 7) work with Condor team to transition to VDT releases.
1/2004	1) JIM V1.2 release; 2) Full logging service features ready; 3) Initial management of structured jobs; 4) Additional scheduling functionality; 5) Augmented resource advertisement; 6) Begin further evaluation of Web Services; 7) Round 2 deployment, 8) Updated support features; 9) improved reliability.
4/2004	1) JIM V1.3 release; 2) Move toward Globus and Condor Web Services implementations (pending their status); 3) New VO options (pending availability); 4) Full accounting statistics available; 5) Improved support and reliability.
7/2004	1) JIM V2; 2) Begin exploring distributed replica catalog solutions; 3) More work on web services.

4.6 STAR

A real challenge is ahead of the STAR collaboration for the experiment Year4 run. The long Gold on Gold running period will lead to an unprecedented amount of data which will require drastic change in the computing model, waiting for years for the data to be analyzed not being an option. In that regard, it will be a test of the strength of choices in the adoption of some of the Grid technology.

Currently helped by PPDG funding at the level of 1 FTE year, our current plan layout requires 47 months FTE equivalent. Only 12 months are covered by direct funding, the rest is planned to be covered by collaborative efforts with other teams (SDM, Condor, Jlab), internal manpower driven at great cost from within the experiment and external funding. The failure risk is high while the need for a transition to a Grid model for computing is within a year. This situation will ultimately lead to a dilemma.

Although foreseen, our general goals include

- 1. Consolidation and finalization of our File or data collection Replication mechanism:** To this end, we must: Complete and stress test, within the coming month, the replication registration service work between our two main sites. We are committed to provide feedback and participate to the improvement of new releases of the HRM products, including v2.0 when it becomes available. Work on the best use of network bandwidth and resources in general. Generalize the use monitoring of the transactions and transfers
- 2. Finalize our monitoring strategy and find/deploy/adopt a common set of tools** suitable for most (if not all) experiments at our respective shared facilities.
- 3. Finalize the convergence toward a site registration authority agent.** Work with facility personnel to define and deploy a STAR “VO” and evaluate existing tools for VO membership handling and account mapping.
- 4. Full migration of some workload to the grid,** two of which is within our reach: Introducing and test such an interactive user analysis framework as the Grid Collector. Several opened questions remains such as latencies for transferring files, resource allocation and resource management currently not being addressed. Achieving a user level job submission (batch) on the Grid during the coming year and eliminating “random” submission by providing to the users an interactive environment in which the evaluation of the outcome of a massive analysis can be made prior from using the full extent of the available resources. Monte Carlo simulations, usually self-contained, must transition to a Grid environment as early as possible. For reaching the above objectives, we must consolidate our existing framework: stress test submission using Condor-G (currently still an advanced tester and developer level activity), finalize our Monitoring and Discovery Service deployment and take advantage of its information for a site-level available resources. Most importantly, our current JDL approach must converge and rapidly be frozen (part of the user shielding policy to the Grid details).

While our focus must remain on the above objectives, GT3 recent release will ultimately lead to an increase usage in the community: its attempt to generalize services through the OGSA and Web Services approach using a potentially interoperable way to specify services (WSDL) makes it an attractive and awaited stable product we cannot ignore. While a full deployment of GT3 is seen as part of our Grid Year 4 activities and hardening, we will actively work on migration of our existing tools toward a Web Service oriented approach (WSDL/G-SOAP).

Table 14. STAR Team Milestones

Date	Work Item(s)	Comments
1/2004	1	Initial Replica Registration service (RRS) deployed
7/2004-9/2004	1	RRS consolidated, HRM V2.0 deployed to one (or more) new site.
8/2004	3, 4	Contingent on availability of sufficient effort from STAR and collaboration from Condor.

4.7 PHENIX

PHENIX presented their interest in collaborating on PPDG to the Steering Committee in June 2003. Following this we encouraged the experiment to collaborate on Year 3 Plans with STAR. The PHENIX computing coordinator has been added to the Steering Committee list. The PHENIX plan describes areas of joint interest and potential work with STAR and US ATLAS in the context of PPDG. The impact and increase in support load on the computer science groups due to an additional experiment participating in PPDG activities and meetings is under discussion.

Objectives of the PHENIX grid research efforts in the coming year include the adoption and/or development of tools: 1) supporting overall data management and the large-scale migration of data files between sites; 2) enabling job submission to more than one site; and 3) supporting management of those jobs by users.

5. Funded Effort

Clearly PPDG activities and deliverables benefit enormously from synergy with and leverage from the program of work of the wider computer science and experiment groups. In the PPDG quarterly reports we provide a broader list of those people directly participating in PPDG meetings and work. For the PPDG Year 3 plan we list those people funded in some part (>10% for the CS and NP teams, >25% for HEP teams) by the project, together with the areas of work.

TEAM	Name	F	Grid Systems	Job Mgmt	Data Mgmt	AAA	Analysis	Other
Globus/ANL	Ian Foster	Y	x	x	x			x
	General Support	Y	x		x	x		x
	Jennifer Schopf	Y	x		x			x
	William Allcock	Y	x		x			x
	Von Welch	Y				x		
	Stu Martin	Y	x	x				
ATLAS	David Adams	Y					x	
	Wensheng Deng	Y			x		x	
	Gerry Gieraltowski	Y	x				x	x

TEAM	Name	F	Grid Systems	Job Mgmt	Data Mgmt	AAA	Analysis	Other
	Dantong Yu	Y	x			x		
BaBar	Andrew Hanushevsky	Y			x			
	Adil Hassan	Y			x			
	Wilko Kroeger	Y			x			
CMS	Conrad Steenberg	Y					x	x
	James Letts	Y	x					
	Eric Aslakson	Y	x	x			x	
	Anzar Afaq	Y	x	x			x	
Coordination	Ruth Pordes	Y	x		x			x
	Doug Olson	Y	x	x	x	x	x	x
	Miron Livny	Y	x	x	x		x	x
	Joseph Perl	Y					x	
D0	Igor Terekhov	Y	x	x				
	Andrew Baranovski	Y	x					
	Gabriele Garzoglio	Y	x	x	x			
	Parag Mhashilkar (student)	Y	x	x				
	Vijay Murthi (student)	Y	x	x				
SRM/LBNL	Arie Shoshani	y			x			
	Alex Sim	Y			x			
	Junmin Gu	Y			x			
	Viji Natarajan	Y			x			
SRB/UCSD	Reagan Moore	Y			x			x
	Wayne Schroeder	Y			x			x
JLAB	William Watson	Y			x			x
	Bryan Hess	Y			x			x
	Ying Chen	Y	x		x			x
STAR	Gabrielle Carcassi	Y	x	x				
	Eric Hjort	Y	x		x			
Condor/U. Wisconsin	Peter Couvares	Y		x		x		
	Alan DeSmet	Y		x		x		
	Todd Tannenbaum	Y		x				
	Derek Wright	Y		x				