

Particle Physics Data Grid Quarterly Status Report

October - December 2000

13 Feb. 2001

Contents

Introduction.....	2
Summary of activities related to "High performance site-to-site file replication service"	2
Summary of activities related to "Multi-site cached file access service"	3
Summary of activities related to "Development of Grid architecture and APIs"	3
PPDG Quarterly Status Report for Argonne National Laboratory, Mathematics and Computer Science Division, Distributed Systems Laboratory.....	4
Description of activities related to "High performance site-to-site file replication service"	4
Description of activities related to "Development of Grid architecture and APIs".....	4
PPDG Quarterly Status Report for Brookhaven National Laboratory,	5
Description of activities related to "Development of Grid architecture and APIs"	5
Description of activities related to "High performance site-to-site file replication service"	5
PPDG Quarterly Status Report for Caltech.....	6
Description of activities related to "High performance site-to-site file replication service"	6
Description of activities related to "Multi-site cached file access service"	6
Description of activities related to "Development of Grid architecture and APIs"	7
Description of other activities	7
Bibliography	7
PPDG Quarterly Status Report for Fermilab.....	9
Description of activities related to "High performance site-to-site file replication service"	9
Description of activities related to "Multi-site cached file access service"	9
Description of activities related to "Development of Grid architecture and APIs"	9
Description of other activities	9
PPDG Quarterly Status Report for Jefferson Lab.....	10
Description of activities related to "Development of Grid architecture and APIs"	10
PPDG Quarterly Status Report for LBNL, Scientific Data Management	11
Description of activities related to "Development of Grid architecture and APIs"	11
PPDG Quarterly Status Report for LBNL physics	12
Description of activities related to "High performance site-to-site file replication service"	12
Description of activities related to "PPDG Project Coordination".....	12
PPDG Quarterly Status Report for SLAC,.....	13
Description of activities related to "High performance site-to-site file replication service"	13

Introduction

This report follows the style of the previous report and summarizes the activities and progress of the Particle Physics Data Grid project for the period October 1 to December 31, 2000. Immediately following the introduction is a summary of activities across all participating sites. More detailed descriptions of these activities are available in the appendices of reports from individual sites (in alphabetical order) regarding activities at those sites.

Future reports on the PPDG project will be organized along lines that reflect the next stage of the project with emphasis on integrating and deploying initial data grid services to several HENP experiments as well as the core software work on modifying or developing the necessary middleware to support these services.

Summary of activities related to "High performance site-to-site file replication service"

Replication of data sets is necessary for performance reasons when the data, compute resources, and researchers are distributed over a wide geographic area. Last quarter the Globus team work on replica catalog and GridFTP code was just entering alpha testing. The alpha testing program has been quite successful based on feedback from the alpha users and our own experience at SC 2000. Of particular note is the addition of 64 bit file support for files greater than 2 GB in size and data channel cacheing which will reduce FTP overhead when multiple tranfers are made to the same site during a session.

During SC2000 the Globus team demonstrated our Data Grid tools and infrastructure. We moved 230.8 GB of data between the SC show floor in Dallas and LBNL in one hour, for an aggregate bandwidth of 512.9 Megabits per second. We also achieved a peak bandwidth of 1.03 Gigabits per second over 5 second intervals.

BNL and LBNL have begun planning a production data replication service for the STAR experiment which will allow data transfer between BNL and LBL (and possibly other sites) using grid services. Along with planning was some experience gained from installing Globus tools on some nodes at BNL and LBNL.

At the SuperComputing 2000 (November) event, the Caltech group demonstrated CMS object replication in the WAN. The demonstration utilized NTON links between Caltech's Center for Advanced Computing Research (CACR) and the SC2000 exhibit. The demonstration utilized the Globus toolkit for authentication and large window file transfers.

The Caltech group is working with CERN's External networking group on upgrades of the US-CERN line. These upgrades will allow us to test trans-Atlantic transfers with a goal of achieving 100 Mbits/second by Spring 2001.

Fermilab plans to work with the Cern and Caltech CMS GDMP project to integrate the first implementation of the Globus file replication services. As part of this work we will work with the Fermilab Strong Authentication project to understand and integrate support for interfacing the Globus Security Infrastructure to the Fermilab Kerberos realm.

On November 3, just before the opening of the SC2000 conference and exhibition, SLAC upgraded its connection to NTON from OC12 to OC48. This provided a strong incentive to meet or surpass the PPDG "100 MB/s" goal in sustained transfers between the SLAC/Fermilab booth at SC2000 and SLAC. During the SC2000 exhibition a transfer rate of 990 Mbits/s was measured sustained for several seconds.

The path from SLAC to CCIN2P3 Lyon continues to be used for transfers of over 300 GB per week. SLAC has been working with the Globus team and currently aims to bring elements of the new Globus file replication tools into production in mid 2001.

Summary of activities related to "Multi-site cached file access service"

The Caltech group has been working on the provision of services to automatically replicate databases between regional centers. The GDMP tools are Globus Middleware based database replication tools, which provide all the functionality of existing Perl based scripts in CMS, plus an improved interface, automatic web publishing of transfer progress, and the ability to resume a transfer from a checkpoint in case of network failure. The GDMP team have been interacting closely with colleagues at FermiLab to integrate the FermiLab HRM with GDMP. Once this work is complete, there will be a much more reliable and clean interface to the MSS. In January of 2001 a preliminary implementation of GDMP tools with the Globus Replica Catalogue should be available.

Summary of activities related to "Development of Grid architecture and API's"

The Globus project continues to help define the architecture, protocols, and APIs that enable grid computing including the Grid Security Infrastructure, Grid Information Service, GridFTP, and replica management work. We refined the Replica Management API and released it for final review. We have refined the security model for grids through the concept of a community authorization server (CAS). This will allow local sites to retain control, and reduce the local sites administration burden. It will also allow community permissions and policies to be enforced across different administrative domains.

FNAL completed the first phase of the test bed application integrating existing components of a grid architecture between Wisconsin and the D0 SAM system. This was demonstrated at SC2000. Collaboration with the LBL and Wisconsin groups resulted in many changes to the HRM interface defined at <http://gizmo.lbl.gov/ppdg/> (the current IDL is at http://cdcvs.fnal.gov/cgi-bin/public-cvs/cvsweb-public.cgi/ppdg_idl/). In collaboration with the CMS GDMP project we are including an HRM interface layer in GDMP and starting the work to extend the HRM interface to include writing of data.

As part of an ongoing activity of exploring the design of grid API's and a data grid infrastructure, Jlab has undertaken an investigation of several web technologies upon which to build and integrate data grid services. This includes the use of Java, Java servlets, and XML to create two of the components of a data grid, the Replica Catalog, and the Replica Host (file server front end). We are particularly interested in the use of Java technologies, as much of the laboratories (Jlab) disk cache and silo management software is now in Java, and the experience with it has been very favorable. The next step in this prototyping activity is to further integrate it with standard file transfer agents such as bbftp (already in use at Jefferson Lab) gsiftp, and gridftp. In addition, we have begun exploring the use of secure communications for servlet interactions, including the use of GSI (Globus Security Infrastructure).

The Scientific Data Management group at LBNL is involved in the development of "Storage Resource Managers" for data grid applications, which applies to both of the site-to-site file transfer service as well as multi-site cached file access service. The main activity during this quarter was the design of a Disk Resource Manager (DRM). The IDL for the DRM was developed to include 3 main functions: getRequest, getRelease, and getStatus. The DRM manages both "permanent" and "volatile" files in its cache. When space is needed, the DRM will only release "volatile" file space. The implementation of a minimal-DRM has already started. We expect to complete an alpha version in the next quarter.

PPDG Quarterly Status Report for Argonne National Laboratory, Mathematics and Computer Science Division, Distributed Systems Laboratory

Date: Dec 18, 2000

Participants: [Ian Foster](#), Steve Tuecke, Bill Allcock, Darcy Quesnel, Joe Bester, John Bresnehan, Sam Meder

Description of activities related to "High performance site-to-site file replication service"

Replication of data sets is necessary for performance reasons when the data, compute resources, and researchers are distributed over a wide geographic area. Last quarter the replica catalog and GridFTP code was just entering alpha testing. The alpha testing program has been quite successful. Several bugs have been fixed, the code base has become quite stable, and improvements implemented based on feedback from the alpha users and our own experience at SC 2000. Of particular note is the addition of 64 bit file support for files greater than 2 GB in size and data channel cacheing which will reduce FTP overhead when multiple tranfers are made to the same site during a session. Replica management work has been primarily in the design phase and is documented under the API section. Note that we differentiate between replica cataloging (bookeeping of what is where) and replica management (combining cataloging with data transport).

During SC2000 the Globus team demonstrated our Data Grid tools and infrastructure. We demonstrated a climate modeling application, which employed our GridFTP and replica cataloging tools. During the Bandwidth Challenge we also demonstrated our ability to move large data sets between sites, for instance to create a new local replica. We moved 230.8 GB of data between the SC show floor in Dallas and LBNL in one hour, for an aggregate bandwidth of 512.9 Megabits per second. We also achieved a peak bandwidth of 1.55 Gigabits per second over 0.1 second intervals and 1.03 Gigabits per second over 5 second intervals.

For further information see: [Globus Replica Management API](#) (Under review and subject to update)

Description of activities related to "Development of Grid architecture and APIs"

The Globus project continues to help define the architecture, protocols, and APIs that enable grid computing. Of particular interest to PPDG are the Grid Security Infrastructure, Grid Information Service, GridFTP, and replica management work. During this quarter we refined the Replica Management API and released it for final review. Development of a reference implementation should begin next quarter. We have refined the security model for grids through the concept of a community authorization server. This will allow local sites to retain control, and reduce the local sites administration burden. It will also allow community permissions and policies to be enforced across different administrative domains. The CAS is still under design and has not yet been released for formal external review.

PPDG Quarterly Status Report for Brookhaven National Laboratory,

Date: December 14, 2000

Participants: Rich Baker, Tom Robertazzi, John Leita, Rich Ibbotson, Ognian Novakov, Razvan Popescu

Description of activities related to "Development of Grid architecture and API's"

We continued some work on an Objectivity database class that provides an object oriented API for on-demand database replication. This API is a wrapper around Globus file replication services with added functionality to handle Objectivity-specific database management. No significant developments were advanced during the October-December, although we did gain additional experience installing Globus tools on various nodes at BNL. During October-December, 2000, this activity consumed approximately 1.5 FTE-months of effort. No equipment was purchased.

Description of activities related to "High performance site-to-site file replication service"

We have begun planning a production data replication service for the STAR experiment which will allow data transfer between BNL and LBL (and possibly other sites) using grid services. The actual implementation of this service is expected to begin in early 2001, and we hope to have some rudimentary service in place by February. The effort to date has been less than one FTE-month.

PPDG Quarterly Status Report for Caltech

December 2000

Participants: *Julian Bunn, Mehnaz Hafeez, Takako Hickey, Koen Holtman, Iosif Legrand, Vladimir Litvin, Harvey Newman, Asad Samar*

Description of activities related to "High performance site-to-site file replication service"

At the SuperComputing 2000 (November) event, the Caltech group demonstrated CMS object replication in the WAN. The demonstration utilized NTON links between Caltech's Center for Advanced Computing Research (CACR) and the SC2000 exhibit. The software we demonstrated comprised a set of Java and C++ tools which allowed first a scan over a local Objectivity-based Tag database, and then the submittal of a request for the full event data corresponding to the chosen subset of the Tag events. This request was relayed to the CACR server over the WAN. Software on the CACR server then replicated the required full events out of a large Objectivity/ORCA database, packaged them into individual database files, and then shipped these back across the WAN to the client, where they were attached to a federated database. Once attached, a histogram was accumulated of an event parameter value (a parameter not available in the Tag). The demonstration utilized the Globus toolkit for authentication and large window file transfers.

The Caltech group is working with CERN's External networking group on upgrades of the US-CERN line. These upgrades will allow us to test trans-Atlantic transfers with a goal of achieving 100 Mbits/second by Spring 2001. The upgrades will require some modifications of the local network infrastructure at Caltech, in order to avoid bottlenecks there.

Description of activities related to "Multi-site cached file access service"

The Caltech group has been working on the provision of services to automatically replicate databases between regional centers. As production and analysis becomes more distributed between regional centers, efficient tools to move databases become more important.

The GDMP tools are Globus Middleware based database replication tools, which provide all the functionality of existing Perl based scripts in CMS, plus an improved interface, automatic web publishing of transfer progress, and the ability to resume a transfer from a checkpoint in case of network failure. In addition GDMP tools will be able to handle automatic replication and synchronization of registered sites. The functional prototype was completed in July of 2000. Version 1.0 was released at the end of summer 2000 and version 1.1 was released at the beginning of October.

The decision was made to use GDMP as the file replication system in CMS production for fall 2000. This required increased fault tolerance, more CMS specific features, and user support. The system was tuned to be able to run in a production type environment. This step included adding many new features to the existing prototype including checksum caches and dummy file attaches and support for parallel transfers to improve performance, resumption of file transfers from the latest checkpoint in case of bad network connection for added error recovery, and catalog filtering to provide users with more flexibility on which files to import from other sites or export to other sites. Added user support included setting up the GDMP Web page, adding a userguide to explain the main installation and execution steps and finally setting up a support list for users.

The GDMP team have been interacting closely with colleagues at FermiLab to integrate the FermiLab HRM with GDMP. Once this work is complete, there will be a much more reliable and clean interface to

the MSS (a need which again came up when a large ORCA federation was replicated from CERN to Caltech in preparation for the SC2000 demonstration ... see above).

The first implementation of GDMP database replication tools can only replicate and manage Objectivity database files, because they are tightly coupled to Objectivity applications. In January of 2001 a preliminary implementation of GDMP tools with the Globus Replica Catalogue should be available. This will allow file format independent replication and GDMP should replace the investigative prototype for all applications. Also in January an AMS plug-in which will implement a security protocol should be complete.

Description of activities related to "Development of Grid architecture and API's"

The Caltech group has been working on the establishment of a system to allow the submission, termination and monitoring of jobs as a group. The system currently has a scheduling mechanism that allows selection of processors based on processor type, load, and availability of data set and maintains replicated states, so that computations will not be lost if a server fails. The initial service was tested on 32 processors with the ORCA production software. Attempts to expand the system system to 64 processors ran into scalability issue, which have been solved.

A function prototype exists, when development is finished this should be converted to a product and integrated with the GDMP tools. The eventual goal is a complete system which can efficiently control and schedule jobs sending them to the data or the data to them.

Description of other activities

Installation of a large prototype Tier2 Prototype has commenced at Caltech's Center for Advanced Computing Research. The prototype consists of twenty dual-CPU Pentium III rack-mounting slave nodes of height 2U containing 30Gbytes of local disk, an HP ProCurve Ethernet switch with 24 ports 10/100 ports, and two Gbit uplink ports, a Dell PowerEdge 4400 rack-mounting server with dual Pentium 1GHz CPUs, 2 Gbytes memory, two Gbit Ethernet cards, and dual PCI busses, and two Winchester Systems RAID arrays of 0.5 Tbytes disk each. This system will be fully commissioned early in 2001, and will be used for various studies of LHC-generation simulation, reconstruction and analysis software, including Monte Carlo production, Grid-based file replication using GDMP (with new HPSS integration), task scheduling schemes, and PPDG related work. At the same time, a twin Prototype is being installed at San Diego Supercomputer Center, and the two systems will at some point be combined across NTON links to offer a combined partitionable system of double the size.

The Caltech group has been working on simulations of large scale production systems in order to help understand and optimize performance using the MONARC simulation toolset. Successful simulations of the Spring 2000 CMS ORCA production have been performed. Detailed monitoring of the production was performed for input to the simulation. Simulation will be performed of the 2000 Fall Production trying to build on what was learned. In 2001 there will be a simulation of tape access to evaluate modes and guidelines of tape access and what will be the quality of the service. Also in 2001 there will be a simulation of transactions to update meta-data in the data-base and the effect on production.

As an indication of the maturity of the simulation tool kit, it is now possible to evaluate job submission schemes on the a simulated production system. There are currently evaluations of self-organizing neural networks for job scheduling in distributed computer systems.

Bibliography

1. Tier2 Prototype Center in California: <http://pcbunn.cacr.caltech.edu/uscmssw/tier2.pdf>
2. GDMP workplan: <http://pcbunn.cacr.caltech.edu/uscmssw/samar.pdf>
3. Job scheduling system workplan: <http://pcbunn.cacr.caltech.edu/uscmssw/hickey.pdf>

4. Distributed Database Management and Replication at Object Granularity workplan:
<http://pcbunn.cacr.caltech.edu/uscmssw/holtman.pdf>
5. MONARC Simulation tool: http://www.cern.ch/MONARC/sim_tool/
6. MONARC ACAT paper: http://www.cern.ch/MONARC/sim_tool/Publish/SONN/
7. Mehnaz Hafeez, Asad Samar, Heinz Stockinger, "A DataGrid Prototype for Distributed Data Production in CMS", VII International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT2000), October 2000.
8. Asad Samar, Heinz Stockinger. Grid Data Management Pilot (GDMP): A Tool for Wide Area Replication, to appear in IASTED International Conference on Applied Informatics (AI2001), Innsbruck, Austria, February 2001.

PPDG Quarterly Status Report for Fermilab

Date: Dec 15, 2000

Participants: Jim Amundson, Phil Demar, Don Petravick, Ruth Pordes, Igor Terekhov Rich Wellner, Vicky White

Description of activities related to "High performance site-to-site file replication service"

Fermilab plans to work with the Cern and Caltech CMS GDMP project to integrate the first implementation of the Globus file replication services. As part of this work we will work with the Fermilab Strong Authentication project to understand and integrate support for interfacing the Globus Security Infrastructure to the Fermilab Kerberos realm.

Description of activities related to "Multi-site cached file access service"

The D0 SAM system is being interfaced to the IN2P3 bbFTP to support higher throughput transfer of Monte carlo files between Nikhef and Fermilab. Nikhef will configure their SAM system to store files in the local Storage Management system (SARA) and/or the Fermilab Enstore system. The SAM Metadata catalog supports the recording of multiple locations of files – on tape or disk. .5 man months has been spent on this work.

Fermilab participated in the Fall Production of CMS simulation data with files being distributed to CERN from the Fermilab disk cache and Enstore system. It is difficult to quantify the PPDG related work here - 1 man week is probably appropriate.

Description of activities related to "Development of Grid architecture and API's"

We completed the first phase of the test bed application integrating existing components of a grid architecture between Wisconsin and the D0 SAM system. This was demonstrated at SC2000. Collaboration with the LBL and Wisconsin groups resulted in many changes to the HRM interface defined at <http://gizmo.lbl.gov/ppdg/> (the current IDL is at http://cdcvs.fnal.gov/cgi-bin/public-cvs/cvsweb-public.cgi/ppdg_idl/). In collaboration with the CMS GDMP project we are including an HRM interface layer in GDMP and starting the work to extend the HRM interface to include writing of data. Two man months has been spent on this application.

Description of other activities

The CDF INFN group completed their first round of QOS tests between Fermilab and Italy, towards the goal of remote front end crate monitoring for the upcoming run - <http://www.cnaf.infn.it/~ferrari/quadis/>

PPDG Quarterly Status Report for Jefferson Lab

Date: Jan. 8, 2001

Participants: Chip Watson, Jie Chen, Ying Chen

Description of activities related to "Development of Grid architecture and API's"

As part of an ongoing activity of exploring the design of grid API's and a data grid infrastructure, we have undertaken an investigation of several web technologies upon which to build and integrate data grid services.

The world wide web has proliferated in part because of the simplicity of its infrastructure, namely a large number of autonomous web servers, with most content delivered using the (fairly simple) HTTP protocol. This protocol is easily supported in any language, and supports a rich variety of content (HTML). The needs of the marketplace have now resulted in the creation of a number of software technologies to leverage the web to deliver much richer content. Extremely sophisticated web services are now emerging in the area of e-commerce and business-to-business commerce (B2B), with complex content delivered in the form of XML (eXtended Markup Language) documents.

The data grid services envisaged by the PPDG project are no more complicated than those of B2B, and could potentially profit from incorporating those same technologies. Hence we have undertaken an investigation of the use of Java, Java servlets, and XML to create two of the components of a data grid, the Replica Catalog, and the Replica Host (file server front end). We are particularly interested in the use of Java technologies, as much of the laboratories disk cache and silo management software is now in Java, and the experience with it has been very favorable.

The specific tasks which have been accomplished are as follows:

- (1) install Apache web server and Tomcat servlet engine (open source, easily replicated elsewhere)
- (2) design and implement XML-producing servlets for a ReplicaCatalog and ReplicaHost
- (3) integrate these servlets with the in-house databases to reflect the silo contents and disk cache state
- (4) design and implement XML style sheets to allow browsing by any web browser (no special client-side software installed)

A presentation of this work as presented at a collaboration meeting in December can be found at http://www.ppdg.net/mtgs/19dec00-fnal/XML_files/v3_document.htm.

The next step in this prototyping activity is to further integrate it with standard file transfer agents such as bbftp (already in use at Jefferson Lab) gsiftp, and gridftp. We will also be investigating the use of Java servlets to create third party file transfer agents, and to interact with a local disk cache manager using the interactions specified by PPDG. In addition, we have begun exploring the use of secure communications for servlet interactions, including the use of GSI (Globus Security Infrastructure).

This activity has consumed approximately 3 man-months in this quarter.

PPDG Quarterly Status Report for LBNL, Scientific Data Management

Date: Jan. 8, 2001

Participants: [Arie Shoshani](#), Alex Sim, Andreas Mueller, Ekow Otoo

Description of activities related to "Development of Grid architecture and API's"

The Scientific Data Management group at LBNL is involved in the development of "Storage Resource Managers" for data grid applications.

The main activity during this quarter was the design of a Disk Resource Manager (DRM). The IDL for the DRM was developed to include 3 main functions: `getRequest`, `getRelease`, and `getStatus`. We describe briefly below the design concepts of the DRM, in order to explain what will be implemented initially, which we call a "minimal DRM".

The design of the DRM includes the following functionality:

1) The DRM manages both "permanent" and "volatile" files in its cache. When space is needed, the DRM will only release "volatile" file space. There are 2 parameters that the DRM uses: the total size of the cache, and the maximum size that can be used for permanent files. For this purpose we added the functions: `makePermanent` and `makeVolatile`.

2) When a file is requested, if the file is found in local cache, and the file is volatile, it is pinned. Permanent files are not pinned. Pinning means that the file is marked with a timestamp. When the DRM needs space, some policy based on the timestamp is used to determine which file to remove. The typical policy is to use a time-out since last pinned, which is another parameter to the system. The details of various pinning policies and the reasons for them can be found in a powerpoint document:

<http://gizmo.lbl.gov/~arie/pinning/index.html>

and in the word document included in the ppdg message:

<http://www.ppdg.net/archives/ppdg/2000/msg00115.html>

3) When a file is requested, but not found in local cache, the DRM connects the remote DRM or HRM that has that file. It requests the file, and then issues an gridFTP get request to get the file into the local cache. If the remote site has no DRM or HRM it attempt the gridFTP anyway.

4) The DRM is multi-threaded and can handle multiple file requests concurrently. However, it has a limit for the number of parallel requests that it handles. If that limit is reached, then the DRM will queue the request. A request can be cancelled with a `getRelease` function.

5) The DRM can be called in a "blocking" or "non-blocking" mode. In the case of a blocking call, the caller waits till the DRM responds. In case of a "non-blocking" call, the DRM returns immediately and then calls back the caller in an asynchronous manner.

The minimal-DRM has the following limitations: a) While it handles parallel requests, it does not queue requests beyond its limit.

b) It performs pinning at level-1 only. See the pinning document for the meaning of level-0 through level-4 pinning.

c) It does not accept "non-blocking" calls and does not perform call-backs.

The implementation of the minimal-DRM has already started. We expect to complete an alpha version in the next quarter.

PPDG Quarterly Status Report for LBNL physics

Date: Jan. 8, 2001

Participants: Stu Loken, Doug Olson, A. (Sasha) Vaniachine

Description of activities related to "High performance site-to-site file replication service"

In preparation for deploying a file replication service between BNL and LBNL for STAR there were numerous tests of network performance between these sites over the production ESnet and LAN infrastructures at BNL and LBNL. Numerous changes to the network configurations at both BNL and LBNL along with an upgrade to the ESnet connection at BNL (to OC3) resulted in performance that was not well understood initially. The latest test result between a linux machine at BNL and a solaris machine at LBNL achieved a 5 MB/sec bandwidth for single ftp stream with 1 MB TCP window size. This is approaching what one would expect for this production WAN.

There was a meeting with a number of BNL staff on Dec. 14, 2000 to discuss and agree on implementing the STAR file transfer as a "production" grid service for the RCF. As well as planning, there was some effort on installing and configuring Globus tools at LBNL (PDSF).

This effort was not significant during Oct. - Dec. 2000.

Description of activities related to "PPDG Project Coordination"

This activity involved holding meetings (virtual and real) as well as some travel.

This effort was about 1.5 person-months in Oct. – Dec. 2000.

PPDG Quarterly Status Report for SLAC,

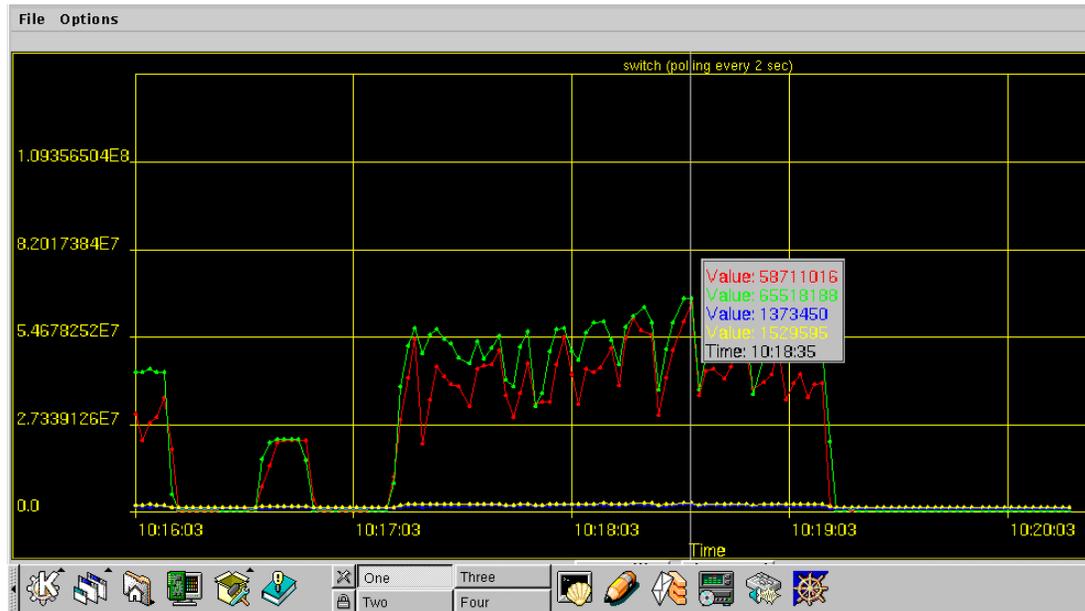
Date: Jan. 8, 2001

Participants: Bob Cowles, Andy Hanushevsky, Adil Hasan, Steffen Luitz, David Millsom, Richard Mount, Davide Salomoni.

Description of activities related to "High performance site-to-site file replication service"

Speed Testbed

On November 3, just before the opening of the SC2000 conference and exhibition, SLAC upgraded its connection to NTON from OC12 to OC48. This provided a strong incentive to meet or surpass the PPDG "100 MB/s" goal in sustained transfers between the SLAC/Fermilab booth at SC2000 and SLAC. During the SC2000 exhibition a transfer rate of 990 Mbits/s was measured sustained for several seconds.



The screen shot shows the peak performance that was achieved using two Intel PCs at SC2000. More details these test are available at <http://www-iepm.slac.stanford.edu/monitoring/bulk/sc2k.html>

BaBar Intercontinental Testbed

The path from SLAC to CCIN2P3 Lyon continues to be used for transfers of over 300 GB per week and plans are being made by IN2P3 to increase the bandwidth to OC12 as need rises in the next 18 months. Current data replication tools must be improved and the work is progressing with heavy PPDG involvement. Adil Hasan has made two visits to ANL to work with the Globus team and currently aims to bring elements of the new Globus file replication tools into production in mid 2001.

In addition a small project is underway to test the applicability of SRB/MCAT to immediate BaBar needs.

Personpower used at SLAC October - December: 2 person-months.

Equipment bought October - December: LuxN WavPortal system (placed at Sprint PoP for connection to NTON): \$19,635.