LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO –
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type *Annual Report*	LIGO-T020078-00-Z	01 July 2002

LIGO DCC Archival Copy:

iVDGL Annual Report for 2001-2002

NSF Grant 0122557

iVDGL Collaboration

http://www.ivdgl.org

California Institute of Technology
LIGO Laboratory - MS 18-34
Pasadena CA 91125
Phone (626) 395-212
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Laboratory - MS 16NW-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

www: http://www.ligo.caltech.edu/

GriPhyN 2002-11
June 28, 2002
http://www.ivdgl.org/

iVDGL Annual Report for 2001–2002

The iVDGL Collaboration

NSF Grant 0122557

Paul Avery	University of Florida	Co-Director
Ian Foster	University of Chicago	Co-Director

Table of Contents

1 The iVDGL Project

1.1 Introduction and goals

The goal of this project is to establish and utilize an international Virtual-Data Grid Laboratory (iVDGL[1]) of unprecedented scale and scope, comprising heterogeneous computing and storage resources in the U.S., Europe—and ultimately other regions—linked by high-speed networks, and operated as a single system for the purposes of interdisciplinary experimentation in Grid-enabled data-intensive scientific computing. The project extends and applies the Data Grid research performed by the NSF-funded GriPhyN Project[2], and will deploy the Virtual Data Toolkit (VDT[3]) created by GriPhyN. The iVDGL project is also closely aligned with the DOE-funded Particle Physics Data Grid (PPDG[4]), as described later.

The iVDGL will drive the development, and transition to every day production use, of Petabyte-scale virtual data applications required by frontier computationally oriented science. In so doing, we seize the opportunity presented by a convergence of rapid advances in networking, information technology, Data Grid software tools, and application sciences, as well as substantial investments in data-intensive science now underway in the U.S., Europe, and Asia. We expect experiments conducted in this unique international laboratory to influence the future of scientific investigation by bringing into practice new modes of transparent access to information in a wide range of disciplines, including high-energy and nuclear physics, gravitational wave research, astronomy, astrophysics, earth observations, and bioinformatics. iVDGL experiments will also provide computer scientists developing data grid technology with invaluable experience and insight, therefore influencing the future of data grids themselves. A significant additional benefit of this facility is that it will empower a set of universities who normally have little access to top tier facilities and state of the art software systems, hence bringing the methods and results of international scientific enterprises to a diverse, world-wide audience.

This report is being written approximately seven months after the report of the initial award in September 2001. We have started substantial efforts towards the establishment of the laboratory, including the creation of five active Work Teams, deployment of early Grid testbeds, extensive interoperability work with our European colleagues, and substantial Education & Outreach activities. Much initial effort has also been spent on recruiting and hiring students, postdocs, scientists and staff; hiring a Project Coordinator, writing detailed planning documents; setting milestones; implementing a management structure; collaborating with other Grid projects and coordinating the experiments which compose iVDGL. During this time, we have had two collaboration wide meetings, as well as numerous smaller meetings for activities related to the Work Teams and Grid interoperability efforts. These topics are discussed below in more detail.

An important change in the Grid landscape occurred in Spring 2002 when the LHC Computing Grid (LCG) project was launched to organize Grid systems and other aspects of computing needed for the LHC. LCG is focused on a very large user community that has connections with most international data grid projects, including iVDGL. It thus serves as a special focus of our coordination and interoperability work, as discussed in Section 5.

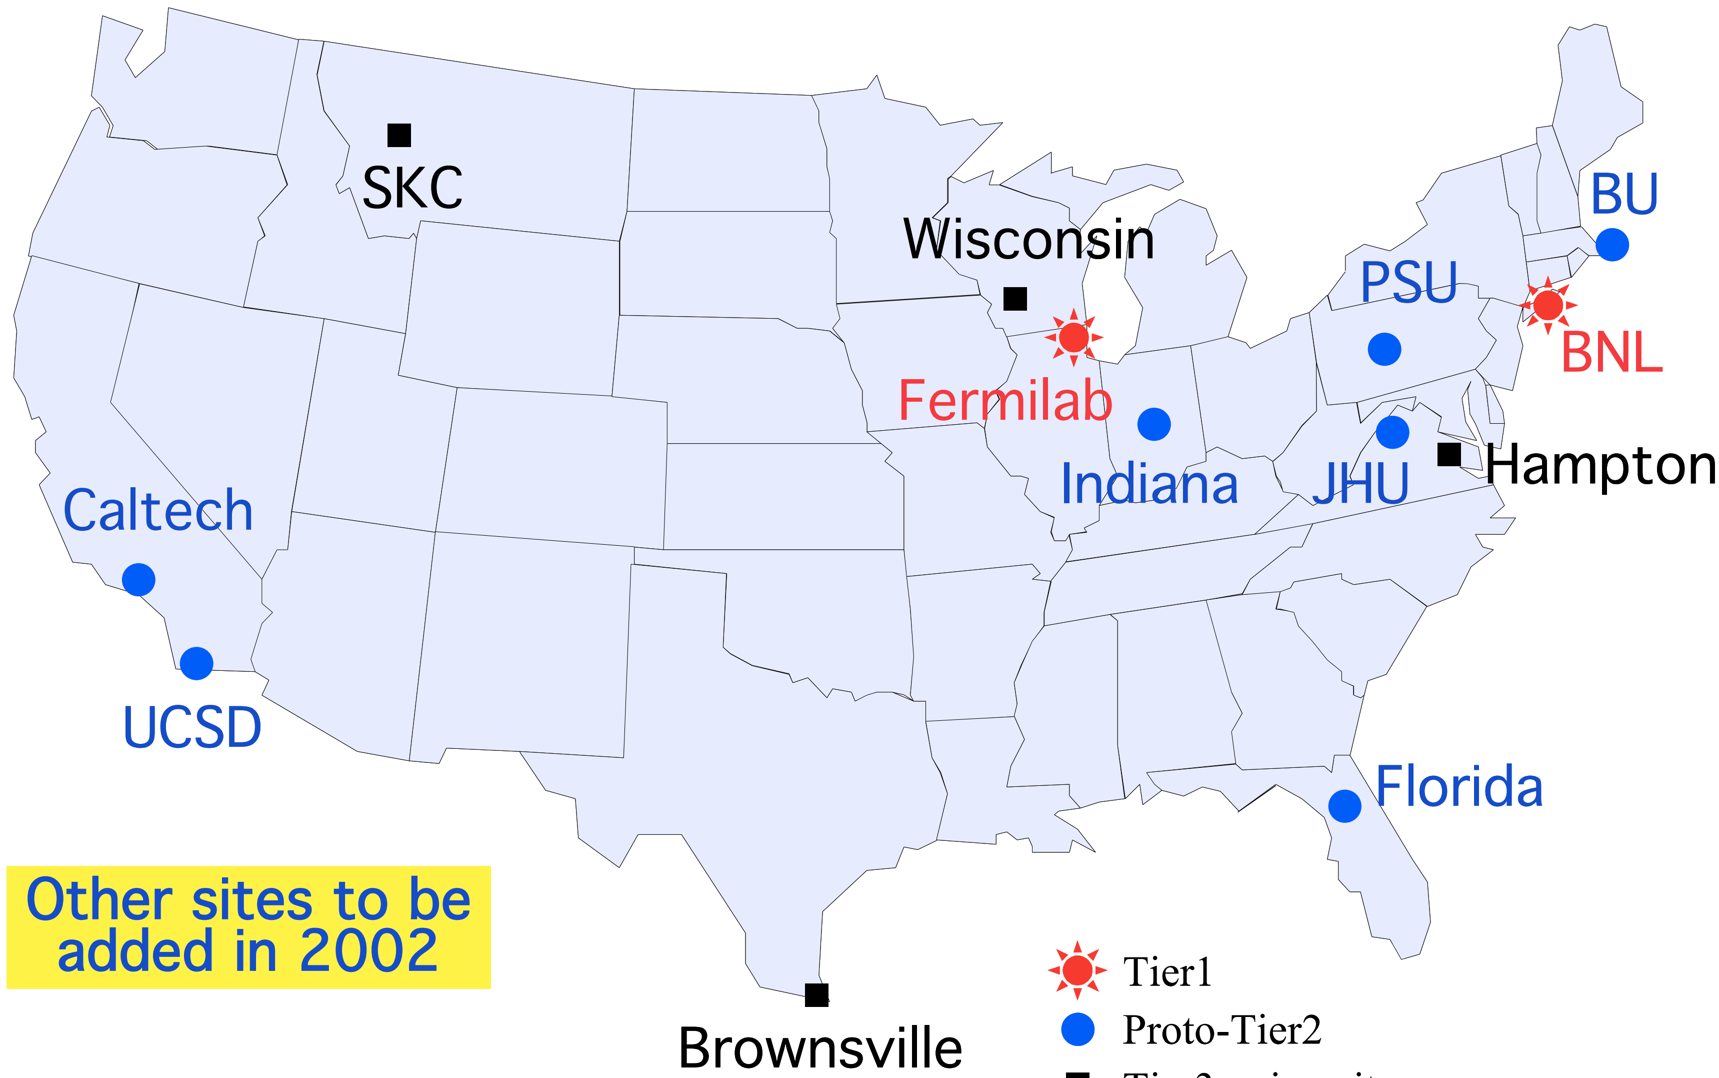

Figure 1: Map showing initial iVDGL sites in the U.S. More sites will be added by the end of Summer 2002.

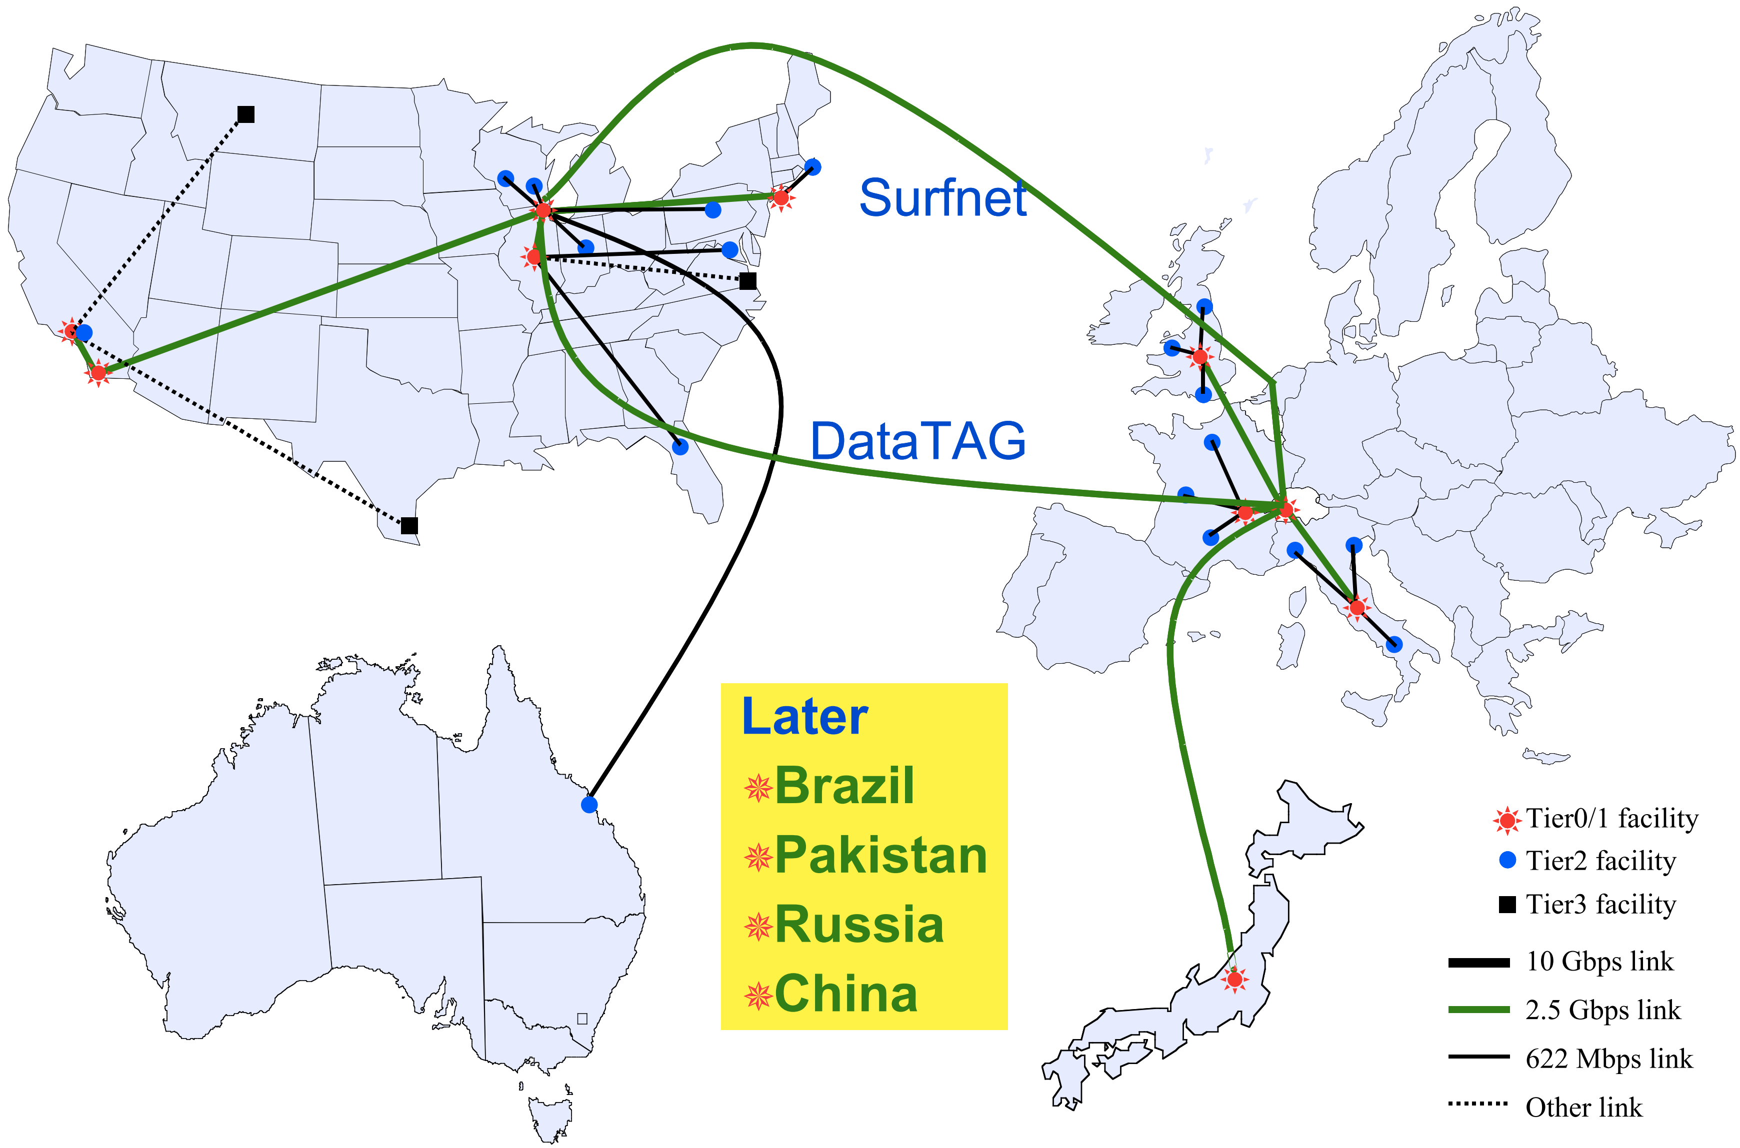

Figure 2: Map showing iVDGL resources in mid-2002. South American sites will be added by the end of 2002.

1.2 Participants and first year funding

A table showing a list of participants and their affiliations is shown on the web site[1]. We have expended much effort in finding and hiring personnel, but have experienced a somewhat slower rampup in personnel than we expected when the proposal was submitted. At the time of this report (June 2002) we have hired or made offers to approximately 75% of our full complement of hires, including people who were promised as matching contributions at Florida, Boston U and Indiana.

Funding for the U.S. part of this effort is provided by the National Science Foundation ($13.65M), with approximately $2M of in-kind contributions (primarily people) from member institutions. International participants pay their own costs. The participants in the NSF proposal are listed in the Table below.

Table 1: iVDGL institutions at the time of the proposal

Institution	Basic role in iVDGL
University of Florida	CMS, management
Caltech	CMS, LIGO
UC San Diego	CMS
Indiana University	ATLAS, iGOC
Boston University	ATLAS
U Wisconsin, Milwaukee	LIGO
Penn State	LIGO
Johns Hopkins	SDSS, NVO
University of Chicago	CS, management
U of Southern California/ISI	CS
U Wisconsin, Madison	CS
U Texas, Brownsville	Outreach, LIGO
Hampton University	Outreach, ATLAS
Salish Kootenai College	Outreach, LIGO
Fermilab	CMS Tier1
Brookhaven	ATLAS Tier1
Argonne	ATLAS, CS

We estimate that slightly more than 40% of our funds will not be expended at the end of the first year, though it is hard to be precise more than three months before that date (15 subcontracts need to be tracked and the dates of some hires are not precise). A major reason for lack of spending is our decision to delay purchasing hardware until better dual-CPU Pentium motherboards are available and RAID systems are better defined. We have set up procedures identical to those used in GriPhyN to track our total expenses every three months to aid our planning and reporting. Because of the vagaries of institutional reporting, the quarterly summaries will be in arrears by 1-2 months, which is considered acceptable.

1.3 Support for iVDGL Through the UK Fellowship Program

The UK e-Science Core Programme[5], run by the Engineering and Physical Science Research Council (EPSRC[6]), provides fellowships for EU citizens to work in the U.S. at iVDGL institutions. Up to six simultaneous fellowships will be available for young researchers (primarily those with a Computer Science background) for the lifetime of the iVDGL project. Each fellowship will consist of a 2-year computer science fellowship with a partner in the iVDGL, with the option of a following 1-year fellowship in the UK. More details, including qualifications and selection criteria, are available online[7]. A funding mechanism has been set in partnership with the NSF.

2 Management Structure

2.1 Overview

The U.S.-iVDGL project was conceived as a vehicle for testing and experimentation of capabilities developed by other Data Grid projects. The experimentation will be carried out on an international scale with partners from around the world who are interested in developing their own regional grid capabilities. Within the U.S. these goals led to primary relationships with the preexisting GriPhyN[2] and PPDG[4] Data Grid projects, leading to the so-called Trillium effort discussed in Section 5.5. A management structure has been defined that has representation from all three Grid projects in the U.S. Project Steering Group, which helps the Directors define and carry out policy that is consistent with the overall goals of the three projects. Although it draws on expertise and connections with other grid projects, the IVDGL project is funded by and reports directly to the NSF, and is completely responsible for carrying out its work program. The following description is updated from the iVDGL organizational document.

2.2 Meetings

"All-Hands" meetings are held twice a year, normally in conjunction with GriPhyN and even PPDG meetings. Other specialized meetings are held on an ad-hoc basis, including those sponsored by the Work Teams. Weekly meetings are held of the coordination group, and biweekly technical meetings of many of the Work Teams.

2.3 Management description

The diagram below shows the organization of US-iVDGL, including Work Team structure and relationship to the international community and other Grid Projects. We describe briefly the activity in the following sections.

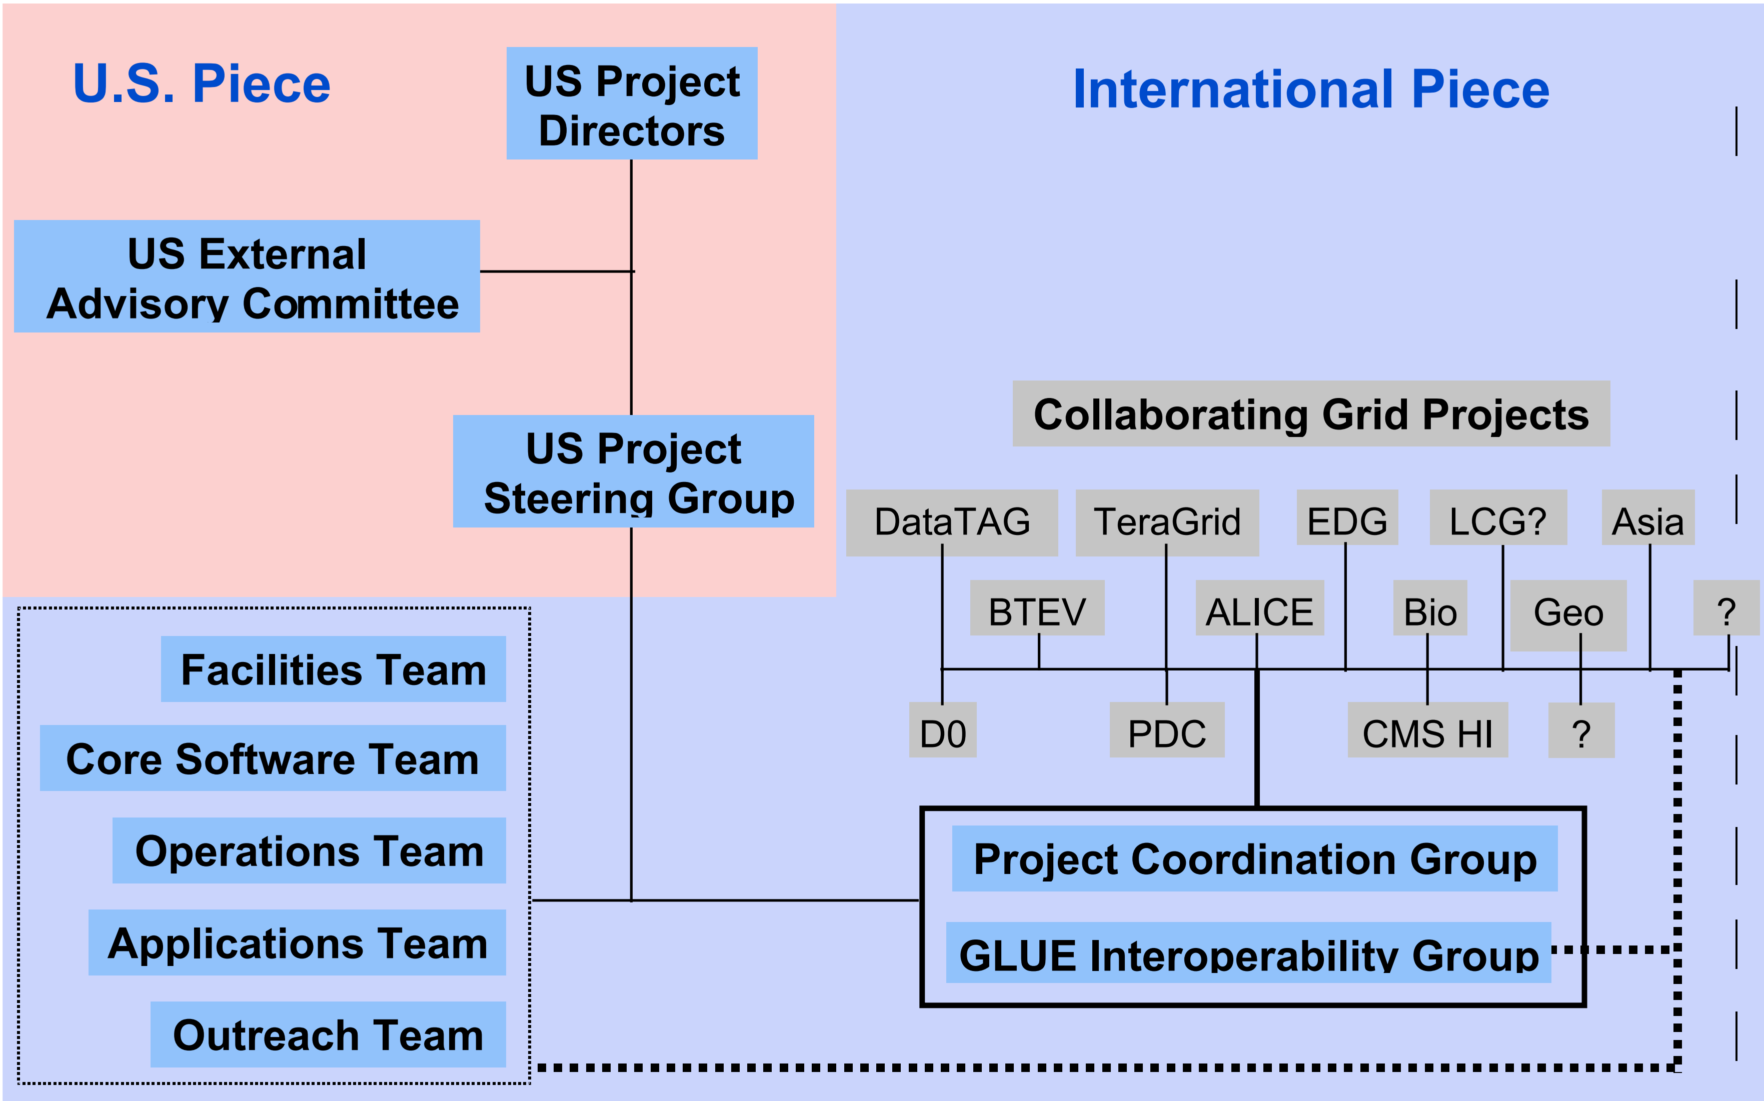

Figure 3: iVDGL management diagram, taken from Jan. 2002 management document.

2.3.1 Project Directors

As with GriPhyN, the Project PIs Paul Avery and Ian Foster act as Co-Directors and report to the National Science Foundation. The Co-Directors make resource allocation decisions jointly and are ultimately responsible for the success of the Project.

2.3.2 Coordinators

The iVDGL has several positions that support the Directors, as describe below.

Project Coordinator: The Project Coordinator assists the Directors and the Work Team leads in the day-to-day activities in support of the Project. The principal duties include coordinating the TestBed and Laboratory activities across all the Work Teams and collaborating experiments and projects – national and international – and identifying areas that require attention. Additionally the Coordinator calls meetings and provides minutes, is responsible for the project planning documents and reports to the funding agencies. Rob Gardner, presently at Indiana University, will be the iVDGL Project Coordinator starting July 1, 2002. He is leaving his current position and will commence employment by the Computation Institute of the University of Chicago on that date. Ruth Pordes of Fermilab has been acting as Interim Project Coordinator since the end of 2001.

Deputy Coordinator: The Deputy Coordinator assists the Project Coordinator by maintaining the web pages for the Project and the Work Teams, and providing technical and agenda support for meetings. These duties will be reviewed periodically and modified if warranted. Jorge Rodriguez is the current Deputy Coordinator.

International Coordinator: The International Coordinator is responsible for high-level coordination activities across peer Grid projects in Europe, America and Asia. The coordinator arranges meetings of the HICB three times a year collocated with GGF. The Coordinator tracks new projects as they are proposed and funded, and works with the iVDGL Directors on appropriate contacts and collaborative arrangements. Larry Price of Argonne is the current International Coordinator.

Interoperability Coordinator: The Interoperability Coordinator leads the GLUE Interoperability Group and ensures that IVDGL sites work seamlessly with sites in other parts of the world. Ruth Pordes of Fermilab is the current Interoperability Coordinator.

2.3.3 U.S. Project Steering Group

The U.S. Project Steering Group advises the Directors on important Project decisions and policies, particularly those that impact the operation of the U.S. component of the iVDGL and effect the wider goals of the project. These goals include: (1) enabling the construction of the Tier 2 data processing and analysis centers of the four experiments on the proposal; (2) creating and operating a U.S. Grid laboratory whose resources can be shared among the participating experiments and across a wider set of application domains (e.g. biology); (3) creating and operating an international testbed and provide seamless integration for applications running across the collective set of global resources of the US-iVDGL and its international collaborators; and (4) providing educational and outreach opportunities

Normally, members of the Steering Group attend Project Coordination Group meetings and it is expected to hold separate meetings only as the need arises. Steering Group members are listed in the Table below.

Table 2: Members of iVDGL Project Steering Group

Member	Role
Paul Avery	Project Director
Ian Foster	Project Director
Rob Gardner	Co-PI, Project Coordinator
Ruth Pordes	Interoperability Coordinator PPDG Project Coordinator
Harvey Newman	Co-PI
Alexander Szalay	Co-PI (SDSS/NVO representative)
Albert Lazzarini	LIGO experiment representative
Manuela Campanelli	Outreach Coordinator
Lothar Bauerdick	CMS experiment representative
John Huth	ATLAS experiment representative
Larry Price	International Coordinator
Carl Kesselman	Computer science representative
Miron Livny	Computer science representative
Mike Wilde	GriPhyN Project Coordinator
Richard Mount	PPDG Director

2.3.4 Work Teams

The main iVDGL work activity takes place within the five Work Teams. Each Team is headed by two co-leads (three in the case of Laboratory Operations) in order to share the workload. The current Teams and co-leads are

Table 3: Composition of iVDGL Work Teams

Work Team	Co-Leads
Facilities	Ian Fisk (UCSD) Bruce Allen (Wisconsin, Milwaukee)
Core Software	Scott Koranda (Wisconsin, Milwaukee) Miron Livny (Wisconsin, Madison) Alain Roy (Wisconsin, Madison)
Laboratory Operations	Jim Williams (Indiana) Carl Kesselman (USC/ISI) Ewa Deelman (USC/ISI)

Applications	Lothar Bauerdick (Fermilab) George Feteke (Johns Hopkins))
Education/Outreach	Manuela Campanelli (Texas, Brownsville) Keith Baker (Hampton)

2.3.5 Project Coordination Group

The iVDGL Project Coordination Group provides the forum for short term planning and tracking of the project activities and schedules. Its membership includes the Steering Group, leaders of the Work Teams, and representatives from our national and international partners. The iVDGL Coordination Group provides the forum for collaboration with TeraGrid[8], as well as the European and Asian Data Grid Projects – especially the European DataGrid[9] and DataTAG[10]. The Project Coordination group will have biweekly phone meetings and will meet face to face twice a year, normally during the All-Hands meetings. The Coordination Group includes the following people.

Table 4: iVDGL Project Coordination Group

Role	Member
Steering Group	All members
Work Teams	Co-leads of each Work Team
TeraGrid	1 representative
EU DataGrid	Bob Jones (Management) Federico Carminati (WP8 representative) Charles Loomis (WP6 representative)
LCG	Fabrizio Gagliardi (Grid Technologies) Les Robertson Matthias Kasemann
DataTAG	Olivier Martin Peter Clarke Antonia Ghiselli
Asia	1 representative
South America	1 representative
Deputy Coordinator	Jorge Rodriguez

2.3.6 GLUE Interoperability Group

A special Group has been established to provide the iVDGL work on the GLUE interoperability collaboration recently established under the aegis of the HICB and HIJTB (see Section 5.2). Each Work Team must consider the interoperability of its deliverables with resources and activities in the full Project. Each Work Team should appoint one of the co-leads to be responsible for this aspect of their work. The GLUE Interoperability Group effort has been established that includes peers from the national and international iVDGL partners as well as the representative from each Work Team. This Group works closely with the Project Coordination Group. Selected members of iVDGL partner projects will also to participate in this Group, which is expected to work closely with the Joint Technical Board of the High Energy Physics Intergrid Coordination Board (HICB-JTB[11]), which is also working towards interoperability of different Grid projects. GLUE members are listed below, and the effort is described further in Section 6.

Table 5: GLUE Interoperability Group Members

Person	Role in GLUE Effort
Rick Cavanaugh	US Liaison to DataTAG WP4.3, 4.4 – Interoperable Applications
Ewa Deelman	iVDGL Operations Team co-leader, US Glue Testbed coordinator, US Glue-GLOBUS coordinator
Rob Gardner	IVDGL coordinator
Alain Roy	VDT Team
Doug Olson	PPDG coordinator; US Liaison to DataTAG WP4.2 -AAA
Ruth Pordes	US Interoperability Group lead
Jenny Schopf	Co-leader of the US Physics Grid Monitoring Group
Brian Tierney	PPDG liaison to EDG ATF., US liaison to DataTAG WP4.1 – resource discovery
Mike Wilde	GriPhyN coordinator, US Glue-GLOBUS coordinator

2.3.7 External Advisory Committee

The External Advisory Committee (EAC) serves as a source of strategic advice for iVDGL. The EAC meets annually face-to-face and will participate in teleconference or videoconferencing reviews as needed. We have proposed that the EAC for the GriPhyN project serve in a similar role for the iVDGL. Current EAC members are listed in the table below.

Table 6: iVDGL External Advisory Committee Members

Member	Institute	Position
Paul Messina*	Caltech	Former Director, CACR at Caltech
Bill Johnston*	LBNL	Director, NERSC
Fabrizio Gagliardi	CERN	Project Leader, EU DataGrid
Jim Gray	Microsoft	Microsoft Research
David Williams	CERN	Former Head, CERN IT Division
Joel Butler	Fermilab	Former Director, Fermilab Computing Division

Dan Reed	NCSA	Director, NCSA Alliance
Fran Berman	SDSC	Director, SDSC
Roscoe Giles	Boston U.	Head of EOT-PACI

* indicates EAC co-Chair.

2.4 Relationship of US-iVDGL to international iVDGL

The international iVDGL is organized as a joint effort of US-iVDGL and a number of international partners. US-iVDGL, on its own, is organizing a testbed laboratory within the U.S. to test and integrate the data grid middleware of the GriPhyN and PPDG projects, and selected components of other Grid projects. This effort will be extended through International iVDGL to an international testbed with transoceanic links devoted to testing and integrating the middleware of many Data Grid projects, including at least EU DataGrid[9] and its associated national projects, DataTAG[10], GridPP[12], the LHC Computing Grid (LCG[13]) Project and Japanese Data Grid[14] efforts, in addition to the U.S.-based work of GriPhyN[2], PPDG[4], US-iVDGL, and TeraGrid[8]. A number of application experiments will be supported, including, but not limited to, high-energy physics experiments at the Large Hadron Collider, gravity wave searches, digital sky surveys and virtual observatories.

2.5 Procedures for joining iVDGL

Over time it is expected that additional projects, experiments, laboratories or groups, some of them representing additional scientific and engineering disciplines, will bid to join the iVDGL. Each case will be reviewed by the Project Coordination Group. If approved, collaboration with the new entity will be established through an MOU mechanism, and, if necessary, liaison will be appointed as an addition to the Project Coordination Group. New members are encouraged to join any of the Work Teams.

Because of the limited manpower and funds available within the iVDGL, new members are expected to bring sufficient personnel and resources to create and run their own facilities as part of the Laboratory. Help will be provided as much as possible to bring new sites and groups up to speed, but it is impossible to provide long-term support in the absence of further funding. Participation in Work Teams, meetings and testbed exercises will be the primary avenues for making contributions to iVDGL.

3 Laboratory Grid Activities

3.1 Creation and Operation of US-CMS Prototype Data Grid

3.1.1 Introduction

The creation of a test grid for CMS has been accomplished to validate the main premise of the VDT (which is to provide simple, effective tools for building grids) and to provide the necessary infrastructure for research and development of grid-enabled software. Much work has been invested to demonstrate that the test grid is able to provide the quality results required by the Production of CMS Monte Carlo simulated data. In addition, emphasis has been placed on stressing the Test Grid by submitting "at scale" production requests. This work successfully revealed several bottlenecks and scalability issues in the core-grid software that are currently being addressed.

3.1.2 Introduction

The CMS group within the GriPhyN-iVDGL collaboration includes Caltech, Florida, FNAL, UCSD, and Wisconsin, and has been actively pursuing several testbed activities that are described below. The sites are described in the Table below. Note that a South American site is expected to take part in the U.S.-CMS testbed by the end of Summer 2002. A map showing the U.S.-CMS testbed is shown in Figure 4.

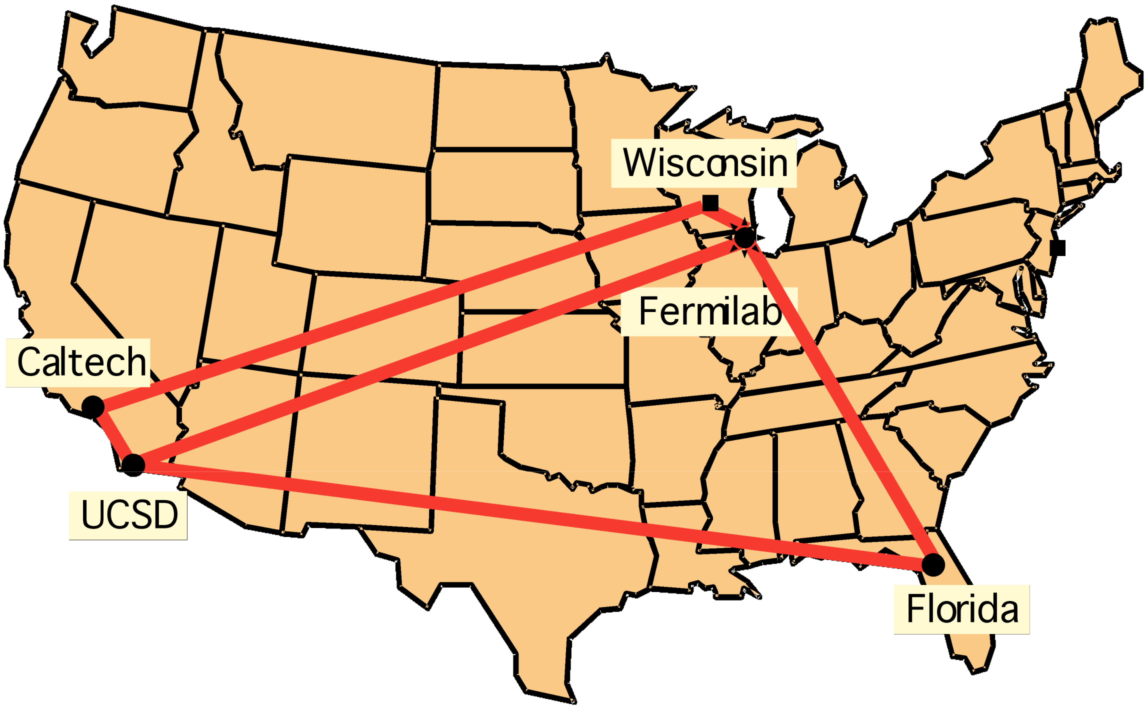

Figure 4: Map showing the U.S.-CMS Testbed, which forms part of the iVDGL computing resources.

Table 7: Institutions in U.S.-CMS Grid Testbed

Members	Status
University of Florida	Member
UC San Diego	Member
Caltech	Member
University of Wisconsin, Madison	Member
Fermilab	Member
Princeton	Pending
UERJ, Brazil (Albert Santoro)	Pending

A small scale US-CMS Test Grid has been built to provide the necessary infrastructure for research and development of grid-enabled software for CMS. The test grid offers approximately 30 dedicated processors (1 GHz each) distributed across the five different CMS institutions. In addition, the test grid is based upon the GriPhyN Virtual Data Toolkit (VDT) and was the first such grid to deploy the VDT. As such, it validated the main premise of the VDT, which is to provide simple, effective tools for building grids. Work to ensure availability of services and monitoring have resulted in a heart beat tool which uses the Grid Data Mirroring Package (GDMP[15]) to periodically move files around the test grid. In addition, ongoing deployment and evaluation of several monitoring tools is currently underway, including: the Globus Monitoring and Discovery Service (MDS[16]), the Condor Hawkeye[17] monitoring service, and the FLexible monitoring frAMework for scalable systEmS (FLAMES[18]) from Iosif Legrand of Caltech.

In the area of production of simulated CMS data, work has been performed studying and prototyping Virtual Data concepts and technologies into CMS Monte Carlo Production. This effort cumulated in a demonstration at Super-computing 2001 entitled "Virtual Data for Real Science" and showed that even an early implementation of the GriPhyN Virtual Data Catalog could already successfully track the dependencies of CMS data files and CMS application transformations between data files. In particular, the demonstration was able to regenerate any (missing or deleted) derived data file (across a test grid) on demand.

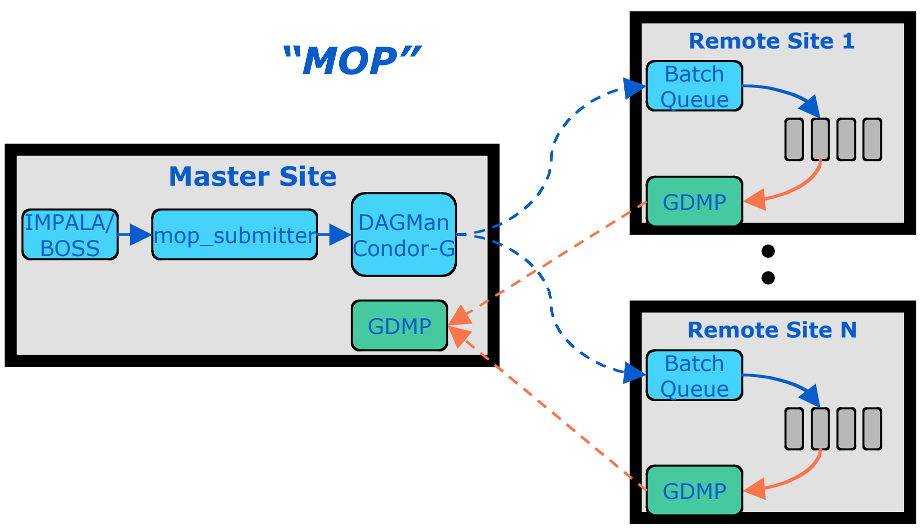

Figure 5: CMS Monte Carlo Production (MOP) prototype architecture

In addition, considerable effort has been devoted to deploying an operational tool on the test grid for the distributed production of CMS Monte Carlo simulated data, known as MOP[19] (**M**onte Carl**O** **P**roduction, see Figure 5). MOP is based upon Globus, Condor-G, and DAGMan[20] to "Grid-enable" the CMS Monte Carlo production system. IMPALA[21], a set of shell scripts, which are currently used by CMS for Monte Carlo production. GDMP[15] is used for file replication and file movement. The requirement for, and the pursuit of, a stable, robust grid platform have successfully provided valuable information to the core-grid middle-ware developers concerning scalability and reliability in an actual grid production environment. In June, the MOP system, running on the test grid, passed a milestone by successfully carrying out the first ever "at scale" production of simulated CMS data (over 100,000 simulated events) on a grid.

Studies in the area of physics data analysis led to a demonstration[22] at Supercomputing 2001 entitled "Bandwidth Greedy Grid-enabled Object Collection Analysis for Particle Physics." This work employed a client/server application that allowed particle physicists to interactively analyze over a hundred GB of physics event data stored at two "Tier 2 centers" which took part in the demonstration. The demonstration successfully showed several key elements of a future CMS Data Grid (see reference) including seamless, easy access to data using virtual data at the object level and grid technology.

Substantial activity has been devoted to the creation of a Grid-enabled data Analysis Environment (GAE) for CMS. Much of the work has focused on using web services to access CMS data. CMS data stored in an OODBMS are converted into reduced data formats, known as AOD or TAG data, and stored into an RDBMS. Physicists may then remotely query and analyze the AOD/TAG data in the RDBMS using various web services and optionally stored RDBMS procedures for improved efficiency.

Prototype systems have been developed both for Oracle9i and SQLServer, using large reconstructed simulated datasets produced by the ORCA[23] (Object Reconstruction for CMS Analysis) program. An analysis using traditional n-tuples to explore LHC physics involving "Jets and Missing Transverse Energy" signatures has been chosen, the data has been converted to TAGs and AODs with a schema generated from the Entity Relationship diagram describing the ntuples, and these objects have then been loaded into Oracle9i and SQLServer databases at Caltech and CERN. Prototypical web services have been written to access and query these object collections.

In addition, a remote data access framework, known as Clarens[24], has been developed to enable analysis of CMS data distributed over a wide-area network. Clarens is based on a client/server approach and provides Web-based services for physicists to run their own CMS analysis and AOD/TAG creation codes on CMS data. The Clarens server is linked against the standard CMS data analysis software libraries, while the (relatively light) Clarens clients are end-user specific; client implementations currently exist for several generic data analysis tools including, C++, Python, Java, PHP, and the Root Data Analysis Framework[25].

Particular attention has been given to documenting CMS requirements for grid computing. One resulting report [GriPhyN 2001-1] details the expected medium term grid infrastructure, which is to be developed by the Grid Projects in conjunction with CMS over the next two to three years. Another report [GriPhyN 2001-16] provides a reference model of how virtual data concepts will likely fit within a CMS data grid. Such research, which has proven vital to the planning of tasks, will continue as CMS refines and redefines its software environment.

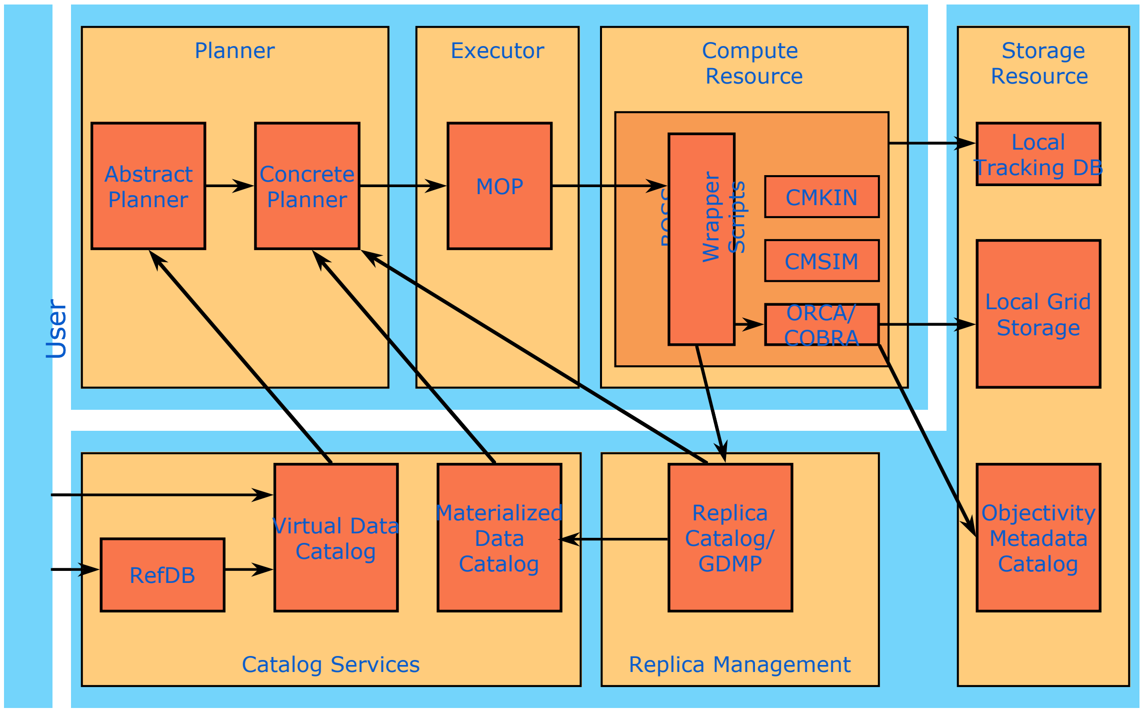

Figure 6: CMS prototype virtual Data Grid architecture

A short term CMS grid implementation plan for 2002 has been defined in [GriPhyN 2002-3]. As part of this short-term plan, a prototype Virtual Data Grid System (VDGS) is currently being developed based upon the documented CMS grid requirements [GriPhyN 2001-1] and upon the GriPhyN Data Grid Reference Architecture [GriPhyN 2001-12], or DGRA. The prototype system implements a virtual data catalog as well as simple planners and grid execution services from the VDT. The purpose of this VDGS prototype implementation is to provide a reference grid integration of current CMS software applications with existing grid technology and to test the viability of the DGRA for CMS needs. To ensure realistic CMS scenarios for testing and to discover unforeseen bottlenecks in a grid production environment, actual CMS Production requests for Monte Carlo simulated data will be processed through the prototype system. In addition, distributed data analysis tools, such as Clarens, will be integrated with the prototype VDGS allowing users to analyze the simulated data. This will provide valuable feedback to developers of the CMS Grid-enabled environment for data analysis as well as the infrastructures for production of simulated data. The CMS-GriPhyN team plans demonstrate the prototype Virtual Data Grid System at Super-computing 2002.

Finally, Virtual Organization issues related to account management and policy enforcement are being investigated. Once a centralized account management agreement (and infrastructure) is in place, Virtual Organization management scripts can be deployed to dynamically create gridmap files for multiple Virtual Organizations. A prototype VO management infrastructure is expected to be in place on the US-CMS Test Grid during Summer 2002. Parallel work involving policy enforcement and user authorization within a multiple VO environment is expected to be performed over the Summer with a simple implementation deployed on the US-CMS Test Grid by the end of the Summer. It is then hoped to combine, for testing purposes, the US-ATLAS and US-CMS Test Grids into a single multi-VO environment for a specified period of time.

3.2 Creation and Operation of US-ATLAS Prototype Data Grid

The US-ATLAS Grid Testbed is a collaboration of ATLAS U.S. institutions that have agreed to provide hardware, software, installation support and management of collection of Linux based servers interconnected by the various U.S. production networks. This project had its beginnings during Grid workshops beginning in June 2000, with a persistent testbed being functional since February 2001. The initial motivation was to provide a realistic model of a Grid distributed system suitable for evaluation, design, development and testing of both Grid software and ATLAS applications to run in a Grid distributed environment. The participants include designers and developers from the ATLAS core computing groups and collaborators on the PPDG and GriPhyN projects. The members are the U.S. ATLAS Tier 1 computing facility at Brookhaven Laboratory, Boston University and Indiana University (the two prototype Tier 2 centers), Argonne National Laboratory HEP division, LBNL (PDSF at NERSC), the University of Michigan, Oklahoma University and the University of Texas at Arlington (see Figure 7).

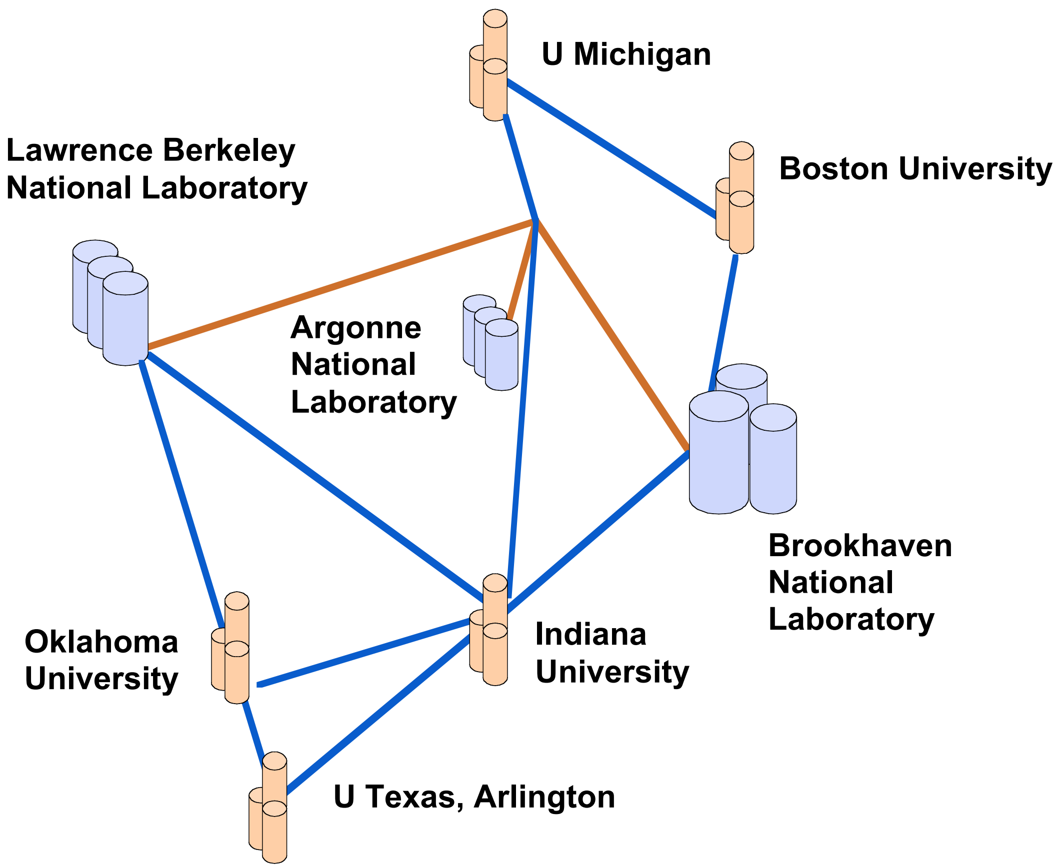

Figure 7: Map showing the U.S.-ATLAS Testbed, which forms part of the iVDGL computing resources.

The iVDGL project is contributing to this effort in a number of ways:

1. Developing two prototype Tier 2 data centers for US-ATLAS located at Indiana University and Boston University.

2. Integrating the two Tier 2 centers with the Tier 1 center at Brookhaven National Laboratory.
3. Supporting development of the Testbed effort.
4. Contributing to the planning and development of the US ATLAS distributed IT infrastructure.
5. Most importantly, supporting integration of VDT with ATLAS core application software.

During the past year the seven participating sites completed the upgrade of installed grid software with installations of:

- Globus 2.0 Beta
- Condor 6.3.1
- GDMP 2.0
- Magda[26] (**Ma**nager for **g**rid-based **Da**ta)
- Pacman (Software distribution tool. See Section 4.4)

The original deployment of the US ATLAS Testbed was completed before the GriPhyN VDT was available, and legacy deployment methods are still being used while an intensive testing and production exercise is completed this summer. The packages deployed on the testbed are the same as would be deployed using VDT, and the deployment is done with Pacman. Once the summer production exercise is completed, the testbed deployment will be redone using VDT. Jason Smith (BNL) is a member of the VDT support team and is working to ensure that the requirements of the US ATLAS Testbed are fully supported in future VDT releases.

To fully utilize the Grid computing model, an application must be capable of being run at any site where it is sent for execution. General high-energy physics applications depend in complicated ways on the configuration of the compute element where the job is run. Since a general Grid computing environment cannot be assumed to be homogeneous, a sophisticated job manager that is capable of checking and resolving dependencies is required. One effort developed on the US ATLAS Testbed is PIPPY which uses MDS to broadcast information about packages that have been installed at a site (See Pacman section below). There is also an effort underway to integrate the ATLAS Software Build process with Pacman so that the dependencies of a job can be determined. These two efforts should provide the necessary inputs to allow a job manager to resolve the dependencies to ensure that jobs execute as expected.

Currently, the job managers available within the Grid tool kits are not capable of resolving the dependencies between a job and the installed software environment of a compute element. An interim solution has been developed to create a "boxed" ATLAS application (GRATS) that has only minimal dependencies and can be run on virtually any computing element that is running a RedHat Linux operating system (RedHat 6.2, 7.1 and 7.2 have been tested). GRATS has been successfully run at all US ATLAS Testbed sites as well as many other sites, including FNAL (CMS Tier 1), University of Florida (CMS Tier 2) and Milan (EDG Testbed). GRATS runs ATLFAST[27] Monte Carlo simulation and reconstruction, creating n-tuples that can be analyzed by physicists. GRATS is now integrated into the Grappa portal tool described below.

Work also continues on the definition of site resources and publication to a grid wide information service, MDS. This will eventually include a hierarchical configuration of GRIS (Grid Resource Information Service) information providers reporting to a central server at Brookhaven

Lab. UT Arlington and Brookhaven lead this effort. In addition, conformance with other Grid projects, such as the EU Data Grid, iVDGL, PPDG, etc, is being assured through participation in the GLUE schema working group (described in more detail in Section 6).

Indiana University has continued to develop the GRAPPA[28] (Grid Access Portal for Physics Applications) job submission portal for the US ATLAS Testbed. A demonstration of this software was made at the ATLAS Software Week in March 2002. The demonstration showed the first instance of Athena grid-wide job submission on the US Testbed. Jobs producing Monte Carlo simulation events from the PYTHIA event generator followed by the ATLAS fast simulation program (ATLFAST) were submitted to several sites on the Testbed using a web-based job submission interface. Globus tools such as GSI for user credentials, GRAM for job submission, Condor and fork job managers were invoked on the prototype Tier 2 centers at Indiana University and Boston University, the Tier 1 grid test node at Brookhaven Lab, and an AFS-enabled grid node at Oklahoma University. Resources (sites, job queues) could be selected from a portal "notebook". Simple job monitoring was accomplished by queries to the site gatekeepers, which check and reported the GRAM status. Another important feature of Grappa is that output files and descriptive information (metadata) are registered with the Magda[26] grid-based data manager tool. This allows subsequent grid compute jobs to access data distributed across the grid. Physics results were written out in the form of n-tuples, and were fetched automatically back to the user's work area. Web reporting of histogram output was automatically generated by invoking PAW Kumac files from a Java servlet. Figure 8 shows the Grappa, web-based GUI used for the ATLAS testbed job submission.

Much of the near term ATLAS Testbed work will focus on completing the full integration of Grappa with other Grid tools that are currently evolving. The results of this work will be presented at the Super Computing 2002 conference in November in Baltimore. All US ATLAS Grid Testbed sites are contributing to this effort. The BNL group will present a demonstration focused on ATLAS physics production fully integrated into a Grid environment. The Indiana group will present a demonstration focused on the GRAPPA Grid portal work. The LBL group will present a demonstration focused on integrating performance monitoring with ATLAS data production. All of these demos will share a common code base and run the same underlying applications but will showcase different aspects. Central to these demos will be a large scale Monte Carlo simulation production running on the ATLAS testbed grid nodes. The simulation will be controlled via the Grappa portal and will make use of existing grid tools for CPU monitoring, registering data and metadata in replica catalogs, and network performance monitoring.

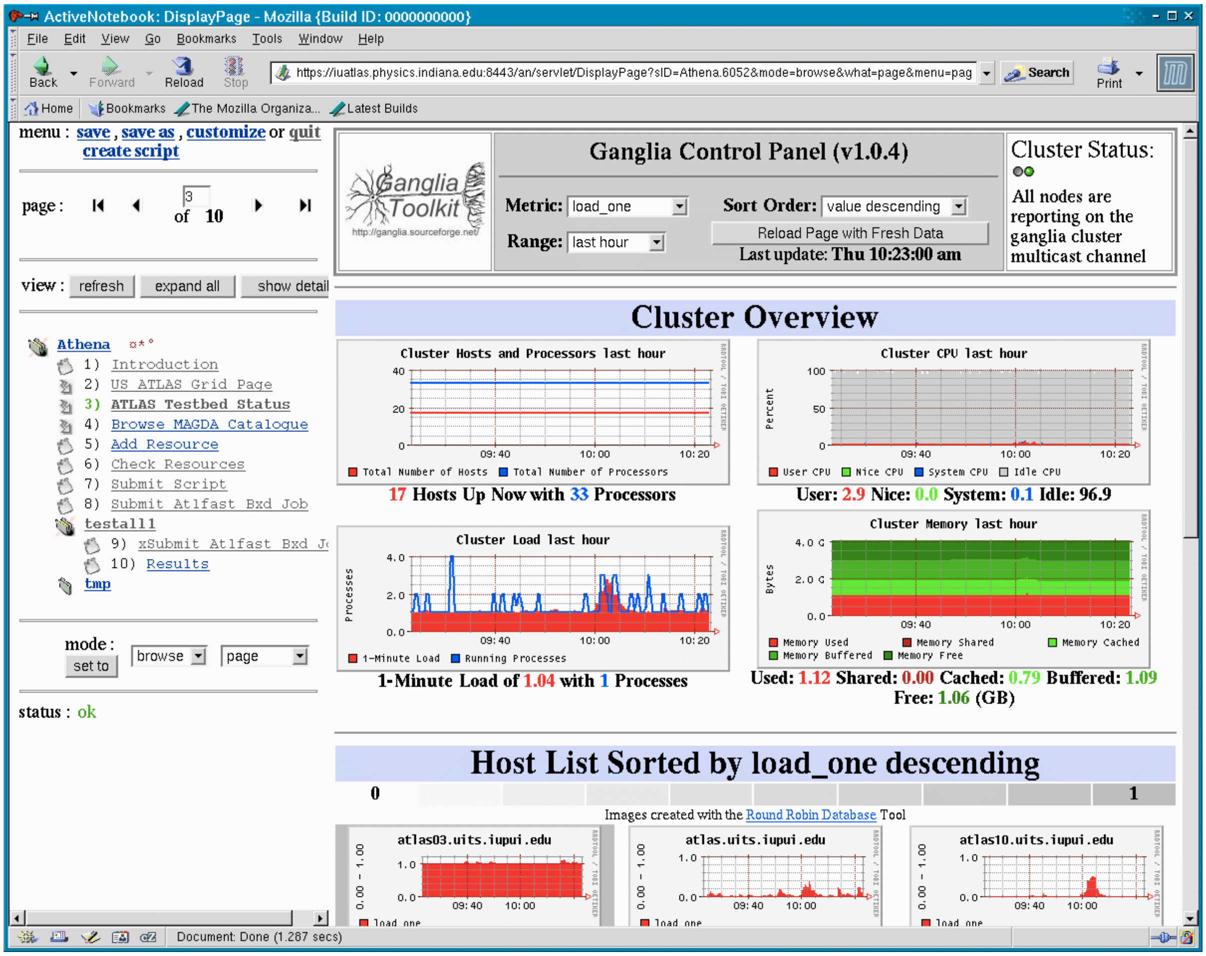

Figure 8: Screen shot of Grappa showing the interface to Ganglia[29] monitoring ATLAS jobs running on the testbed.

3.3 Creation and Operation of LIGO Prototype Data Grid

A number of international gravitational wave projects have begun a collaboration to exchange environmental site data in a lightweight frame format. The initial goal is to develop the infrastructure and protocols for exchanging and analyzing these data. Long baseline correlations are expected to be a limiting factor in the sensitivity that is achievable with a phased array of interferometric gravitational wave detectors. These initial studies, before the projects embark on the exchange of astrophysically meaningful data, will serve to establish the geophysical backgrounds. To date a simple network of servers has been set up at each of the participating sites. The system is called the Network Data Analysis System (NDAS). Participants include:

1. Prof. B.S. Sathyaprakash (GEO), Cardiff University
2. Dr. Karsten Kötter (GEO), Max-Planck-Institut
3. Mr. Benedict Cusack (ACIGA), Australian National University
4. Dr. Roy Williams (CACR), Caltech
5. Dr. Szabolcs Márka (LIGO), Caltech

6. Dr. Benoit Mours (VIRGO), LAPP, Chemin de Bellevue
7. Dr. Daisuke Tatsumi (TAMA), National Astronomical Observatory of Japan

We have developed and presently testing and improving a simple system based on standard UNIX tools and frame library code to transfer and merge data from multiple gravitational wave detectors distributed worldwide. The transfer and merger should take place with less than 20 minutes of delay and the output frame data are available to all participants.

The system is modular to allow future improvements and the use of new tools. As a natural course of the evolution for the project, we now wish to start the implementation of the second-generation system, which will require us to fully utilize recently developed Grid technologies (in particular, the VDT from GriPhyN) and CACR resources. The LIGO NDAS server at CACR (gridligo.cacr.caltech.edu) is dedicated to this effort and will be the target machine to be migrated to the grid based environment (see Figure).

The first phase of the LIGO Prototype Data Grid will to replicate the existing NDAS protocol for exchanging data with our European collaborating partners, GEO600 (UK/Germany) and VIRGO (France/Italy). We will start this effort with VIRGO and then replicate the proven prototype solution to the other participating sites.

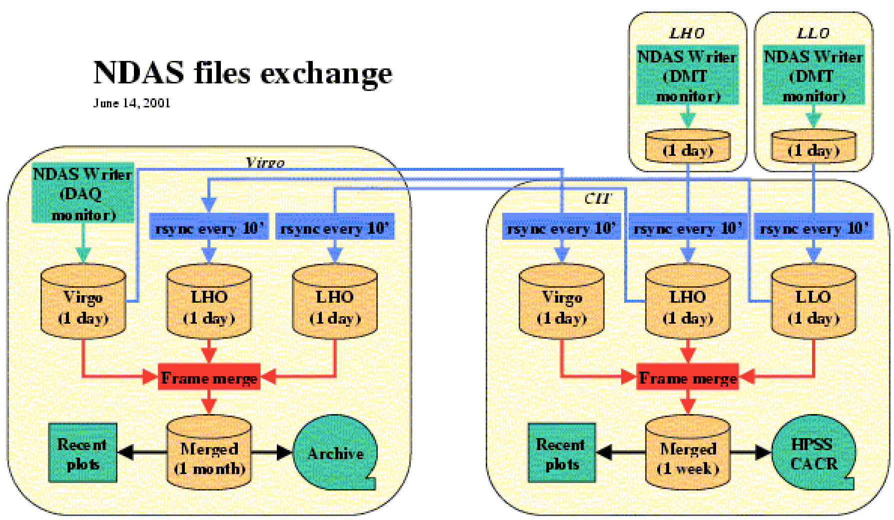

Figure 9: Diagram showing how files are exchanged between two sites using the NDAS system.

The goal of this migration to grid-based tools is to increase the robustness of the computing processes. Several problems have been identified.

- The system reliably produces double and triple coincident frames. However, frames of 4-fold and 5-fold coincidence are much harder to produce reliably.
- The remote frame generating processes worked reliably as long as they had access to the frames and their host computer was not saturated with unrelated tasks. The most prevalent failure mode occurs when remote processes (not under our control) prevent us from generating or transferring the NDAS frames to our central server. These problems are correctable, but require human intervention, leading to delays and data loss.

- At some remote sites the NDAS computer can lose frames when it is overloaded. Other issues include disconnected networks, rebuild/replacement of the computer, accidental use and saturation of our disk buffer by others or installation of new security features.
- We experienced problems associated with widely varying levels of security features at remote sites. Some sites also have limited bandwidth at certain times, leading to delayed transmission. The remote buffer size was increased to avoid data loss due to the erasure of older but not yet transferred frames. Large numbers of simultaneous RSYNC processes cause instabilities, so the maximum number of simultaneous RSYNC processes was set to 10.
- Varying security software, protocols and policies pose significant barriers, but we expect these to be overcome with Grid tools, which are designed precisely to work in different administrative domains.
- There is a large asymmetry between the two central (merger) sites (Caltech/LIGO and Cascina/VIRGO). Due to the very serious bandwidth limitations at the LIGO observatories, we are restricting access to LHO and LLO until the bandwidth increases. Since Cascina/VIRGO cannot directly access the LLO and LHO data, the Cascina site is delayed compared to Caltech and presently do not get the data within the 20 minutes target. In addition, since the Cascina site relies entirely on the Caltech site, therefore it cannot be treated as a backup site. This situation should be changed as soon as allowed by the circumstances.

3.3.1 LIGO Data Replicator

Personnel at the University of Wisconsin, Milwaukee have designed, developed, and deployed a prototype of the LIGO Data Replicator (LDR). The LDR system is designed to replicate raw LIGO data from Tier-1/0 sites to Tier-2/3 sites within the LIGO Scientific Collaboration (LSS). The focus of LDR is on robust, high-performance, and fully automated replication.

LDR is built on top of the standard Globus Data Management APIs. In early work we used the Globus Replica Catalogue for cataloging collections of LIGO data. Because of performance problems associated with the LDAP database backend, however, we stopped using the Globus Replica Catalogue API and instead began using plain text files for cataloging. Recently, however, we have begun evaluating the new Replica Location Service (RLS) APIs developed by the ISI group with collaboration with Work Package 2 for the European DataGrid project, and we plan to integrate RLS into the LDR system.

Data transfers in the LDR system are accomplished using the Globus GridFTP API. In particular we have found it easy to do rapid prototyping using the pyGlobus version of the GridFTP API.

The LDR system is composed of a number of daemons that continuously run with each responsible for a particular, focused task. The daemons can be grouped into those that function replicating data between temporary storage "buckets" located at each site, and those that are specific to a site and function to move data between the local bucket and storage.

The daemons involved in site-to-site replication are:

- LDRrequest: examines catalogs to determine files needed locally and sends via GRAM a request to a remote bucket to have data staged and made ready for transport from the remote bucket.

- LDRfillBucket: examines remote buckets for needed files that have been staged or made available and uses GridFTP protocols to copy them to the local bucket.
- LDRbucketClean: maintains the local bucket or cache, purging old files.

The daemons involved with the local details of staging data to and from the local bucket are:

- LDRemptyBucket: moves data from local bucket to storage and updates local catalogs.
- LDRreconcile: examines catalogs and reconciles with the current state of the storage system.
- LDRstoragePurge: examines catalogs and purges storage of unregistered files.

In addition the LDRmaster daemon watches over the other daemons, restarting any if necessary.

3.3.2 Virtual Data in LIGO

We are also actively pursuing the realization of Virtual Data within LIGO. In November 2001, at SC2001, we demonstrated a simple end-to-end LIGO Virtual Data System, which allowed user to query for LIGO virtual data using simple LIGO-specific metadata attributes. The system included a request interpreter which translated the domain-specific request into a set of data movements and computations. The planner then chose where to execute the computation and where to access the data. The plan was then sent to Condor-G for execution. The resulting virtual product was then cataloged for future use. Although the system demonstrated an end-to-end virtual data materialization, it was only the first step in the building a Virtual Data Grid System. We are currently investing more sophisticated planner techniques and incorporating the Virtual Data Grid technologies within the VDT.

3.3.3 LIGO Tier 2 Center at the University of Wisconsin – Milwaukee (UWM)

The UWM LSC group plays a significant role in the LIGO data analysis community. It is home to several active developers of LDAS analysis code, plays a significant role in LSC software development, and hosts the Binary Inspiral Upper Limit Group, one of the LSC's four upper limit development groups. It also plays an active role in the GriPhyN project, and hosts Dr. Scott Koranda, a grid computing expert who works full-time on LIGO grid computing issues.

In addition, the UWM group hosts a large computing cluster (300-nodes, 24 TB of disk) which has been made available as a resource to the LSC. Through work by Scott Koranda and members of the LIGO Laboratory, this cluster has been grid-enabled and is now use by about a dozen non-UWM members of the LSC. The cluster mirrors a copy of the most recent and interesting LIGO data set, and provides a productive and flexible environment for data analysis developers. It is being used in the development and testing of the binary inspiral search code, the stochastic background detection code, and the pulsar detection code.

The UWM group received IVDGL funding in March 2002, which will support a part-time cluster System Administrator and a postdoctoral-level scientist. The positions are under active recruitment. The System Administrator position has a closing (postmark) date of June 18th, 2002, and by the 19th approximately 15 candidates had applied. The group should be able to hire a first-rate person as a System Administrator. The group has also been recruiting for the postdoctoral position. There is a shortage of qualified people in the grid-computing field, and the best hope seems to be to hire someone with experience in Numerical Relativity and a strong computing background. We made an offer to an attractive candidate, but on the verge of accepting the position, he withdrew from the search to take an opportunity in the Defense/Aerospace field.

The UWM group's principal IVDGL activities to date include:

- Bruce Allen (co-Chair) and Alan Wiseman serve on the IVDGL facilities group. This group is currently carrying out benchmarking and testing of large inexpensive RAID storage arrays and fast inexpensive CPUs. We believe that large, easily managed centralized disk farms for storage of data and intermediate analysis results will be important. The group has shown that one can build large central RAID file servers for under $4k/Terabyte, and is currently construction a 4TB system that should scale to 60 TB at below this cost.

  The IVDGL facilities group is also doing comparative testing of the P4 and Athlon CPUs and chipsets/motherboards. Towards this end, both the LIGO and CMS groups have contributed standalone benchmark programs in the form of static executables. The facilities group has been testing different CPU/chipset/motherboard combinations for several months and is building up a database of testing results. The successful UWM experience in bidding, commissioning and operating a large cluster is also serving as a model for other IVDGL members. Within the next few months, the Facilities group will provide a deployment plan for the first wave of IVDGL hardware acquisitions.
- Scott Koranda serves as co-Chair of the IVDGL/GriPhyN Virtual Data Toolkit (VDT) group. Among the other activities, UWM hosts the VDT web pages.
- As part of his iVDGL-GriPhyN activities, Koranda has written and deployed the grid-based data replication system LDR (LIGO Data Replicator) for mirroring LIGO data from the Caltech HPSS system to the UWM cluster. Recent advances at both the UWM and Caltech ends now allow data to be mirrored onto the system at average speeds exceeding 20MB/sec. This is more than fast enough to keep up with the data production rate at the Laboratory.
- Koranda represents LIGO on the IVDGL SC2002 working group.

3.3.4 LIGO Tier 2 Center at the Pennsylvania State University

The PSU LSC group plays a prominent role in the LIGO data analysis community. It is home to several active developers of LDAS analysis and system code, plays a significant role in LDAS software development. Finn co-chairs the Burst Analysis Group, one of the LSC's four upper limit development groups. PSU also hosts a LDAS development cluster: the first cluster outside the LIGO Laboratory to run LDAS. PSU is the home institution of the National Science Foundation Physics Frontier Center for Gravitational Wave Physics, the first of three Physics Frontier Centers funded by the National Science Foundation. Senior investigator Finn directs this Center, serves on the LIGO Scientific Collaboration Computer Coordinating Committee, which oversees LIGO Laboratory/LIGO Scientific Collaboration participation in iVDGL, and on the LIGO Scientific Collaboration Executive Committee.

PSU is the only LIGO "Greenfield" tier-2 site: the LIGO Laboratory its Observatories, and UWM have existing computing infrastructure that is being extended to support iVDGL activities. PSU is building from scratch a tier-2 center, which will support LDAS development and data analysis. (The existing, small cluster will remain as a stand-alone resource for development of LDAS code using the always current development snapshot.)

To support iVDGL activities at Penn State, senior investigator Finn has teamed-up with Vijay Agarwala, Director of Computing and Information Services for Penn State (formerly the Numerically Intensive Computing Group). The PSU group received iVDGL funding in early 2002, which supports hardware acquisition. Funding for 2002/2003 includes additional hardware acquisition, a part-time cluster system analyst and a postdoctoral-level scientist. The postdoctoral position has been filled with a start-date of 1 October. The analyst position is being recruited. As has been commented, there is a shortage of qualified people in the grid-computing field, and PSU is exploring a number of different possibilities. Until the full complement of personnel are on-board, PSU iVDGL activities are supported by Finn and Agarwala, and Agarwala's staff.

In the period covered by this report, activity at Penn State has focused on scoping the requirements for a greenfield Tier-2 center that will support LIGO data analysis and data analysis software development and recruiting personnel. (The status of the personnel search has already been described above.) In consultation with the LIGO Laboratory PSU has identified as a high priority maintaining configuration and architecture compatibility with the LIGO Laboratory Tier-2 centers, in order to minimize the number of new LIGO Data Analysis System configurations that must be tested and supported, and to maximize the ability of analysis software to migrate between Tier-2 centers. With this priority, PSU has identified a few basic hardware choices for key components and is proceeding to the purchase of a small prototype cluster, which will form the nucleus of the final system. This system should be in place by end July 2002.

3.4 Creation and Operation of SDSS/NVO Prototype Data Grid

The SDSS/NVO Data Grid will be formed between the two iVDGL sites at Fermilab/SDSS (Jim Annis and Steve Kent) and JHU (Alex Szalay). Construction of the sites awaits the completion of the hiring process, which will be complete by the end of July. Fermilab/SDSS will soon thereafter begin hardware purchases.

Hardware testing has been performed in conjunction with the SDSS production team. Dan Yocum of Fermilab/SDSS has put into the SDSS production system a data storage cluster consisting of 10 Linux box IDE-disk arrays, totaling some 20 Terabytes. He reports 200 Mbytes/sec read rates and 100 Mbytes/sec writes. He also reports success in a search for a Linux kernel/patch combination that supports the journaling file system XFS, NFS, the advanced RAID configuration, and the Objectivity database; the issue was stability, with uptimes of less than a week for most kernel combinations tested, and very slow NFS performance for the rest. The elements of an astronomy data grid have been explored by the replication of the SDSS Early Data Release data from Fermilab/SDSS to the Royal Observatory at Edinburgh, UK. The EDR data masses to a Terabyte. At 10 Mbytes/sec rates, it takes a day to download. ROE achieved nowhere near that rate, and the download took 2 months. Advanced software will not help this; data grids require a minimum hardware capability that is surprisingly ill-met in astronomy. Replication to the University of Chicago's Center for Cosmological Physics, on the other hand, took two days and was limited by bandwidth out of Fermilab. Interestingly, though the replication to ROE was slow it was successful, while an attempt using tapes met with failure. The future is clearly in the data grid, and the tiered model of high energy physics is the appropriate paradigm.

The Fermilab/SDSS group will build its iVDGL site around an existing research cluster, the Terabyte Analysis Machine. This is a 10 CPU Linux cluster designed to bring high compute power to bear on Terabyte scale datasets. We have long used Condor on this cluster. In the last quarter we have deployed the core iVDGL software of the VDT v1.0. Our plan is the build this cluster out into a larger iVDGL site over the next quarter, incorporating into the greater iVDGL network.

The SDSS application group (Jim Annis of Fermilab/SDSS and Yong Zhao of UC) constructed a prototype virtual data grid from a pair of machines at the Argonne National Laboratories, the Fermilab/SDSS Terabyte Analysis Machine, and the University of Chicago Condor pool. This prototype data grid was used develop a galaxy cluster finding virtual data application using Chimera (to be included in the next release of the VDT), with the resulting paper accepted as a technical paper at SC2002. We will next drive the application on a scientifically interesting data set using iVDGL sites, as project described more fully in the applications group section.

JHU is currently working on building its own half of the iVDGL SDSS site. Jan Vandenberg is actively involved in the Hardware group, and George Fekete in the Applications group. Our first demo application involves the computation of the three-dimensional power spectrum of the SDSS galaxy distribution from the redshift sample. This involves a very massive likelihood fitting of a high number of parameters (20), where each likelihood function evaluation includes the inversion of a matrix of at least 3000x3000, or larger. We anticipate the need to evaluate several hundred thousand likelihood points. We also need to perform this analysis over a large number of Monte-Carlo realizations of our data set. The data are stored in a relational database. We have already established a very efficient integration of the prototype with the database using Web Services. We have also been extensively experimenting with various forms of web services using the SOAP protocol, and WSDL. In collaboration with Jim Gray we have built several working prototypes. One of those, the SkyQuery is a distributed query engine, which successfully queries databases at FNAL, JHU and Caltech. This was built over 6 weeks, and won the second prize of the Microsoft Web Services Student programming contest.

3.5 Creation and Operation of an iVDGL Prototype Data Grid

A demonstration project called **US-iVDGL-1** has been proposed to drive an initial deployment and operation of the iVDGL laboratory on a short timescale. The project will involve running and support of the Laboratory for a period of a few days at the end of August 2002 using identified effort from each of the Work Teams to do the detailed integration, deployment and support for 2 to 3 weeks. The demonstration will include:

- Dynamic web pages showing US-iVDGL-1 applications being requested, scheduled and executed at all US-iVDGL-1 sites.
- Web pages showing automated installation and execution of different experiment and/or demonstration applications at a single US-iVDGL-1 site.
- Web pages showing problems and/or alarms encountered at each site where applications are deployed, a tracking system for operational support and problem resolution over the Laboratory.
- Application specific monitoring displays, graphical output, job reporting for grid-wide jobs.

As each of the different physics experiment collaborations are currently operating separate prototype data grids, another demonstration project, called **iVDGL-2**, has been proposed to integrate several of the existing prototype grids (including selected sites in the European Data Grid via DataTAG) into a single iVDGL prototype data grid. This project aims to test the ability to seamlessly run experiment-specific applications across an inter-grid environment. In addition, this effort aims to demonstrate, and ensure, that the VDT is interoperable with European Data Grid Middleware at the core-grid software level. This effort will draw on the experience of the Grid Laboratory Universal Environment (GLUE) Testbed (see below).

4 General Activities

4.1 Networking Activities

U.S.-iVDGL is working closely with network backbone providers, laboratories and expert groups within the HENP community to be sure the infrastructure is available to support the test-bed activities at the heart of the iVDGL project. These include, from the U.S. side ESnet, Internet2/Abilene[30], the U.S.-CERN Line Consortium[31] and the ICFA Standing Committee on Inter-Regional Connectivity which is chaired by iVDGL co-PI H. Newman. The leadership of and participation in the Internet2 HEPN Working Group[32] by many of the iVDGL investigators is an example of the increasing cooperation between iVDGL and network providers (Internet2/Abilene). IVDGL Laboratory Operations works closely with the Internet2 Measurement, QoS and E2E[33] Working Groups.

U.S.-iVDGL also has strong connections with the emerging TeraGrid[8] Facility that will be an important high-performance component of the domestic iVDGL testbed. ESnet is supplying additional infrastructure service as it has recently established a Certificate Authority that can be used by U.S. participants in regional and interregional grid testbeds. Tests of interoperability between U.S. and European certificates are currently underway.

The iVDGL has an important connection to the Asia-Pacific via the TransPAC Project[34] (NSF ANI-9730201). James Williams is the TransPAC Principal Investigator and a co-lead for the iVDGL Laboratory Operations Working Group. Mr. Williams has given a number of talks about the iVDGL to researchers in the Asia-Pacific, the most recent being at the Grid Forum Korea meeting in July 2002. We hope to use TransPAC to engage the Japanese ATLAS experimenters at KEK as a part of the iVDGL_2 testbed.

4.2 Certificate Authority Work

4.2.1 Authentication and Certificate Authorities

It is vitally important to have secure policies and mechanisms for authenticating users when they access distributed resources for a virtual organization. A Certificate Authority (CA) carries out this task, with the help of one or more Registration Authorities (RAs), by issuing digital certificates. The role of the CA in this process is to guarantee that the individual granted the unique certificate is, in fact, who he or she claims to be.

4.2.2 Recent CA and RA Work

A great deal of work[35] has been invested this year in ensuring that certificates from one CA are accepted in a wide variety of domains, including international ones. This work has been carried out by PPDG on behalf of the Trillium Data Grid projects: PPDG, GriPhyN and iVDGL.

In recent months, the DOE Science Grid Certificate Authority underwent some small changes and improvements that will favor long-term stability. The most noticeable change is use of the domain name http://www.doegrids.org/ as the public location of the CA. This change reflects the fact that this CA serves a broader community than just the DOE Science Grid Collaboratory Pilot, which uses http://www.doesciencegrid.org/ for its project URL[36]. The role and definition of the Policy Management Authority (PMA[37]) is becoming clarified and a document that serves as its charter is in draft and undergoing iteration.

There is current work to establish an LDAP directory to hold all of the certificates issued by the CA and to serve as the definitive source for public access to these certificates. Details of how this directory is organized are relevant to other work (under the Site-AAA area) and there are some discussions related to this. A test version of this publishing directory is available now and the production version should be in service in a few weeks.

Additional Registration Authorities (RAs) joined this CA recently, where the most significant (from a U.S. Grid projects perspective) is an RA for iVDGL. With help and encouragement from PPDG, and endorsement from DOE, the iVDGL was able to establish an RA with the DOEGrids.org CA and it is now beginning operation. Scott Koranda of Wisconsin, Milwaukee is the contact person. Koranda has already obtained an agent certificate and is now ready to begin processing certificate requests from iVDGL participants. This step is very significant not only because it solidifies the authentication infrastructure on the US side, among all of the experiments participating in iVDGL, it also sets the stage for a close collaboration with the European counterparts, which will recognize DOESG certificates on their resources.

There are some minor issues that develop and get resolved and some issues which are more significant. Probably the most significant issue is on how Certificate Revocation Lists (CRLs) are handled. The current Globus GSI infrastructure has only a token acknowledgement that one should be able to check the validity of a certificate but the infrastructure for managing CRL's is inadequate. It may well be that CRLs are not the best solution to the problem. At present, with some manual intervention on the part of the CA managers, up to date CRLs are published and available to the sites that need them.

4.3 Joint PPDG, Griphyn iVDGL Monitoring and Schema Definitions

iVDGL is part of the Joint Physics Grid Projects monitoring group which addresses Grid monitoring requirements. To date the group has defined a list of use cases, identified a set of information providers needed to produce the information needed. The group has then moved to address the more immediate needs from the Interoperability GLUE project. Interoperable or common resource discovery mechanisms are a key ingredient of interoperable grids. The GLUE Schema project has involved many people from iVDGL, DataTAG, EDG, PPDG and Globus/MDS to provide an upgrade to the MDS base schema that will support the HENP Grid projects on both sides of the Atlantic. This work is being incorporated into the MDS release 2.2 and will be delivered to the EDG Software release team, and VDT software release team by the middle of July. It appears this may also result contributing to TeraGrid as the GLUE schema are being discussed within this forum.

Initial steps in organizing a monitoring effort were taken during the 2001 – 2002 period. A web site was created[38], as were a mailing list and archives. This group is currently being lead by Jennifer Schopf and Brian Tierney. Jennifer is a researcher at Argonne National Lab, and is strongly involved in the Globus MDS and monitoring efforts in the Global Grid Forum. Brian Tierney is a researcher at LBNL, and was recently funded by Mary Anne Scott through PPDG. Brian has been working in the field of monitoring and Grid computing for many years now, including as a lead for the performance effort in PPDG and the lead architect for the NetLogger work.

4.3.1 Use cases

The first goal for this group was to define use cases for requirements gathering. We did this by first defining a template, and then by requesting use cases from the experimentalists involved in this effort. To date we have 19 of these covering a wide range of examples from testing a network for stability to evaluating the progress of an application. Jennifer Schopf presented this work at the Internet2 End-to-End Performance Initiative Measurement Workshop in January 2002, which is available on the web[39]. This work was also part of the talk she gave for the LCG kick-off meeting in March 2002[40].

4.3.2 Requirements document

Serious work started on developing a requirements document for grid-level monitoring issues. The current draft document is available online[41]. We are currently waiting feedback from several experiments, and hope to have a new draft of this by mid-July.

4.3.3 Meetings and Re-scoping

This group met informally as part of the GriPhyN meeting January 2002, and were tasked by the VDT people tasked us with defining a set of sensors to be deployed in the various application testbeds as part of VDT. However, at the PPDG meeting in February it was realized that that goal was really one of fabric management, and work in that area was being done by many members of the PPDG monitoring community already, and what was needed was a better plan for how to interface these various fabric-level monitoring systems to a grid-level monitoring system, such as the Globus Monitoring and Discovery Service (MDS). A new version of the charter has been advertised to the group, and is posted online[42].

This re-focusing will involve the use of the currently being defined unified schemas being developed by the Glue-Schema group (discussed below), the development of needed sensors or information providers to allow inter-operable deployment of this information, and a joint GIIS set-up for this group which has already been set up by Dantong Lu, BNL.

4.3.4 Glue Schema work

In early March the JTB started a group to define, publish, and enable the use of common schemas for interoperability between the EU physics grid projects (focusing on EDG and DataTag) and the US physics Grid projects (focusing in on PPDG, GriPhyN and iVDGL). Brian Tierney and Jenny Schopf are coordinating this effort[43]. This work is part of the Grid Laboratory Uniform Environment (GLUE) Phase I task[49].

This effort will provide a basis from which to understand short, medium and longer term needs and definitions, and will encourage coordinated progress, and increased communication between these groups. No one set of schemas are being adopted, rather a new, unified schema is being developed; with the goal of having schemas defined for use in LDAP, SQL and XML.

The first step is to define common schemas to describe Compute Elements (CE), Storage Elements (SE), and Network Elements (NE), to be used by the MDS and R-GMA Grid Information Services. The goal is to have common schemas defined, deployed, and tested in time for the EU DataGrid Testbed 2 release in September 2002. Common schemas for monitoring and notification events are being address by the Global Grid Forum DAMED working group, and will be addressed later by iVDGL and DataTag.

As of June 2002, the CE schemas are finished, and implementation of them in the Globus and EDG information providers has begun. A test of these is planned for late July. We have started work defining the Storage elements, and have recruited expert help for this from Arie Shoshani and John Gordon.

4.3.5 Work at Northwestern

The work at Northwestern University has focused on monitoring and analyzing the performance of the ATLAS experiment, in particular the ATLFAST[27] simulator. Using the auditors available within Athena, the ATLFAST simulator was analyzed in terms of different generators and identifying the time attributed to the different algorithms for the different generators. At this time, we are working closely with some key ATLAS researchers at LBL to provide performance feedback on the simulator. Another goal of this work is to use the Prophesy Infrastructure to develop performance models for the simulator that can aid in evaluating different execution environments by the replication manager. The performance data is available to researchers via a dedicated website and the Prophesy database.

Future work is focused on detailed monitoring of the ATLAS experiments, identifying where time is spent within the different algorithms used in each experiment. This work entails utilizing the Prophesy Infrastructure to instrument some of the algorithms. The result of this work will be detailed performance models for the replication manager as well as performance feedback for the algorithm developers and ATLAS researchers. Work is also in progress for incorporating the monitoring work into a demo for SC 2002.

4.4 The Pacman Software Packaging Effort

The Grid model of computation relies on the ability to reproduce software environments at will on the machines where computations are to be performed. This is by no means a trivial problem for most of the large-scale scientific projects where Grid computing is attractive. For example, the environment required to do a Higgs reconstruction in the ATLAS experiment is a complicated mixture of ATLAS core software, CERN supplied software, Grid software such as Globus, third party software and specialized Higgs analysis software. Experience in many experiments has shown that construction and maintenance of such environments is relentlessly labor intensive and error prone. Traditional methods are already difficult even within a conventional computing model. In the case of Grid computing, this problem becomes critical.

Fortunately, the problem of defining and reproducing software environments can be solved with help from the Grid concept itself. Just as with computing and data storage, the expertise for installing, setting up and updating software is distributed across the internet. If we introduce a mechanism for an expert to define how software is to be installed, configured, set up and updated on all machines, we can use this distributed expertise to construct complete software environments at the push of a button. Using this mechanism, called *Pacman*, an expert defines a "software cache" containing instructions for fetching, installing and setting up software and defining any dependencies that the software has on other Pacman packages or on other externals. As a Pacman user, given explicit trust of a collection of caches, one can install, set up and maintain subsets of the corresponding software on any of the supported machines. As an example, software for the U.S. ATLAS testbed is distributed across the following caches as indicated in the Figure below

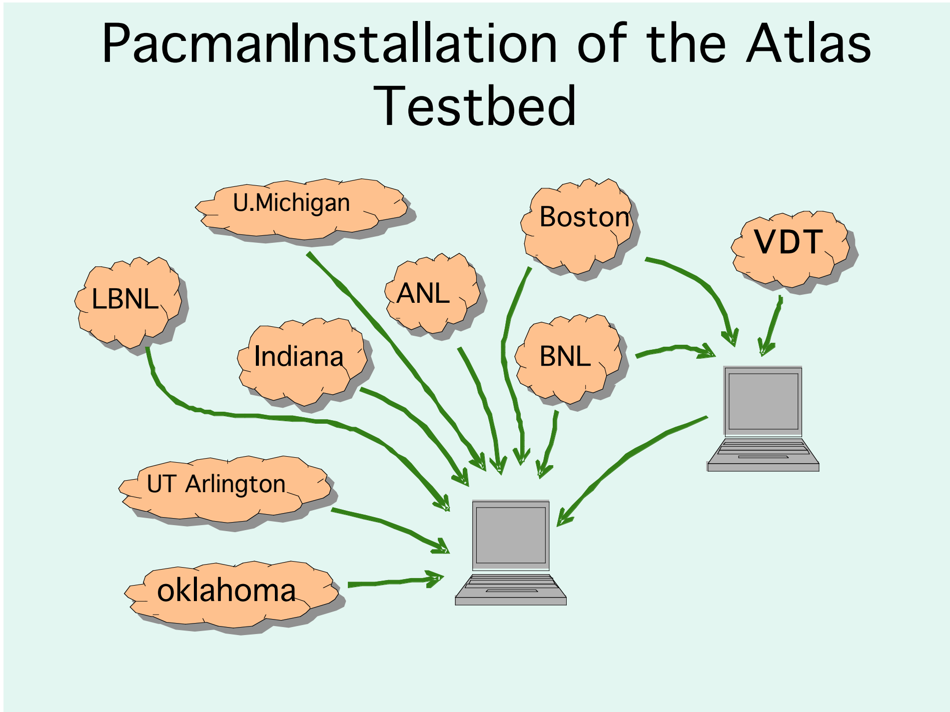

Figure 10: Pacman installation of ATLAS Grid software.

where each "cloud" represents one Pacman cache and where the arrows represent the action implicit in an installation command. In the case of ATLAS, the entire test bed is constructed by issuing the single command line

```
% pacman -get BU:Atlas
```

which installs ATLAS, Grid and third party software. Given this, we have reached a situation where our distributed experts can decide how software should be installed on all machines and their expertise can be broadcast throughout the test bed without each site working through problems individually. Having gotten the ability to reproduce our entire environment on a fresh machine with a single command, we have achieved an important prerequisite for Grid computing.

Pacman is designed with shared computing in mind. The fact that VDT is distinguished from the rest of the software above is meant to illustrate "relative installation" in Pacman. This symbolizes a feature of Pacman where installations can be dependent upon each other. This allows a single installation of common software like VDT to be used by multiple experiments without interfering with each other.

Pacman was developed at Boston University and is being further developed and maintained in collaboration with the VDT team at the University of Wisconson. Pacman has been successfully used at ANL, BNL, LBNL, Boston University, Indiana University, UT Arlington, U.Michigan, U.Wisconson, Oklahoma University. It has been successfully used by the VDT team to distribute VDT itself. The Pippy package from UT Arlington provides a mechanism to broadcast Pacman's database via MDS, so that information about what software is installed where can be used by other Grid applications. Given strong expressions of interest from ATLAS as a whole, from CMS, LHCb and European Grid projects, we are optimistic that Pacman will become a useful piece of the overall Grid computing picture. In particular, the CMS collaboration will embark on a large scale Pacman evaluation involving one US CMS farm, one EU CMS farm and one grid testbed.

4.5 Planning for SC2002

SC2002 demonstration plans include demonstration of experiment test grids, deployment of GriPhyN challenge problems and some specific demonstrations of interoperability between European and US sites. The SC2002 demos will be aligned with existing experiment milestones.

4.6 Meetings

The size, scope and interdisciplinary nature of iVDGL require building effective communication channels between the members of the collaboration. Early on we decided to use face-to-face meetings to establish and maintain such channels. Two All-Hands meetings, both in conjunction with GriPhyN, have been held. These took place Oct. 15-17, 2001 at USC/ISI and April 24-26, 2002 at Argonne. The next All-Hands meeting will be held September 9-13, 2002 at San Diego, in conjunction with GriPhyN and PPDG.

Most iVDGL meetings have been held via telecons and VRVS videoconferencing. A large number of such meetings have taken place in the past six months. The format, dynamics and frequency of these meetings varied from one activity to the other. However, they all contributed to a collection of extremely useful documents.

4.7 Web Site

A web site has been developed by a team at Florida and is being hosted at Argonne. The site has a description of recent news items, ads for open positions an archive for all IVDGL e-mail lists, links to IVDGL meetings (including papers presented at the meetings and registration information), pointers to External Advisory Committee meetings, and links to documents and other relevant projects.

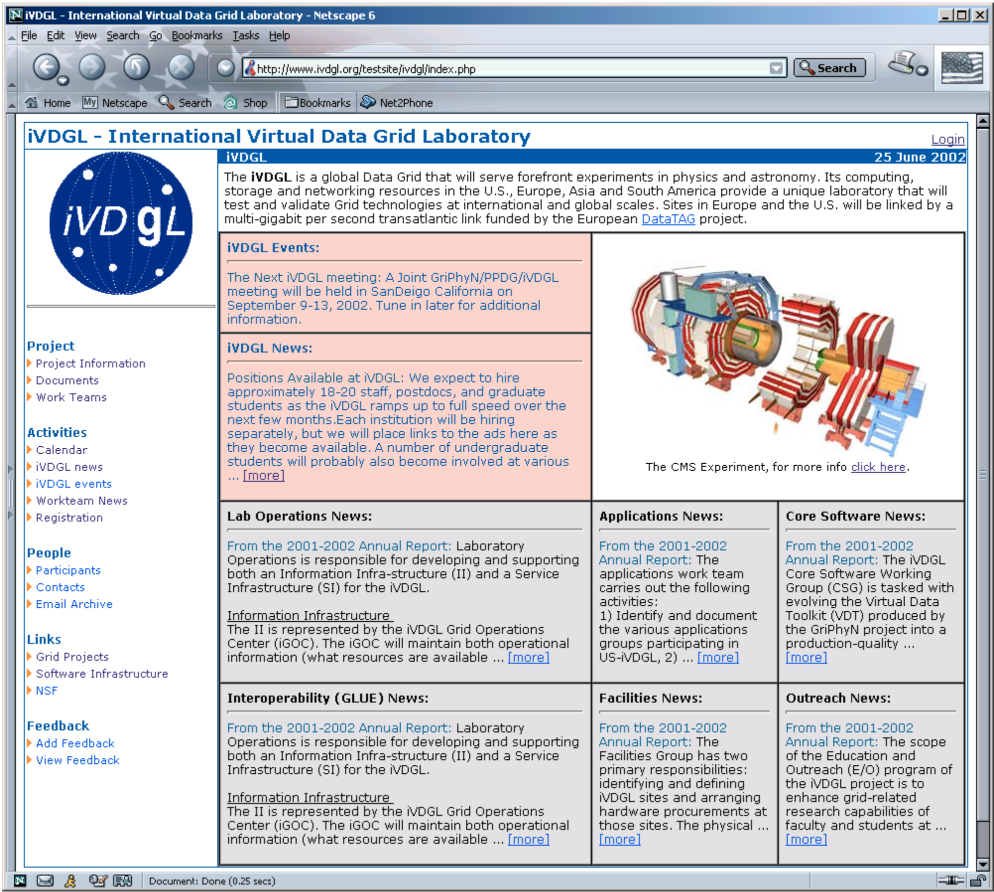

Figure 11: A screen shot of the new IVDGL web site. Cells are reserved for general news about the project (pink) and Work Team (gray). Each "[more]" link points to the full article. The image on the right is randomly selected from a collection of images with links to more information. On the top right hand corner is the link to the user and administrator login system.

The web site has proven to be a very good coordination site for the project, so much so that its deficiencies have become noticeable over the last year. It has also become apparent that maintaining the site with up to date information requires a significant amount of work. To remedy these deficiencies while reducing the maintenance the team at Florida redesigned the site with technology that incorporates the ability to control the information-content dynamically with built in logic. Dynamic information-content which includes: news articles, event postings and meeting information, are now displayed at different times in different locations depending on a predefined date and time attribute associated with the particular news or event article. For example, articles that are out of date are relocated from both the front page and a current location page to a past location page. In addition to the dynamic information-content, the new site features a user login system, a new document server and an agenda creation tool. The user login system was designed to control access to sensitive information and define the individuals who can administer the site. The agenda creation tool allows downloads, restricted uploads and is fully integrated into the new document server. Presenters without administrator intervention will now be able to upload their own presentation and add it to the GriPhyN/iVDGL document database at the same time. We have also designed a tool to facilitate and standardize the creation of registration forms for IVDGL meetings.

The front page was significantly redesigned to display the new dynamic information-content. It now consists of a grid of cells into which is displayed general information about the project, important IVDGL news and events and news from each of the iVDGL Work Teams and the Interoperability (GLUE) effort. All news and events articles contain information that is dynamic in the sense described above. There are also cells reserved for the display of images related to the project complete with links that point to more information about the image. Several images are stored on the site and selected for display randomly. The navigability of the site has been greatly improved by a menu bar that is located on the left side of every page displayed on the site. A screenshot is shown in the above Figure.

We plan to make some additional adjustments to the site before the end of Summer 2002. We will add more internal documents and links to other projects and bring up to date the new combined GriPhyn/iVDGL/PPPDG calendar of meetings and events.

4.8 Publications and internal documents

Because of our close ties with GriPhyN, we have decided to use a single document series that serves both GriPhyN and iVDGL. The following iVDGL related documents can be found on the GriPhyN document server[44].

- *U.S. iVDGL and iVDGL Management Plan*,
  January 2002
- *Start of Project Plans for iVDGL Year 1*
  Feb 2002
- *Progress of the Facilities Work Team*,
  Facilities Team, April 2002
- *Progress of the Core Software Work Team*,
  Core Software Team, April 2002
- *Progress of the Laboratory Operations Work Team*,
  Laboratory Operations Team, April 2002
- *Progress of the Outreach Work Team*,
  Outreach Team, April 2002
- *CMS Grid Implementation Plan – 2002*,
  Grandi, Bauerdick, et al., May 2002
- *Applying Chimera Virtual Data Concepts to Cluster Finding in the Sloan Sky Survey*,
  Annis, Zhao, Voeckler, Wilde, Kent, and Foster, May 2002. Accepted by SC2002.
- *GriPhyN and LIGO, Building a Virtual Data Grid for Gravitational Wave Scientists*,
  Ewa Deelman, Carl Kesselman, Gaurang Mehta, Leila Meshkat, Laura Pearlman, Kent Blackburn, Phil Ehrens, Albert Lazzarini, Roy Williams, Scott Koranda. To appear in HPDC 11.
- *Applications of Virtual Data in the LIGO Experiment*,
  Ewa Deelman, Carl Kesselman, Scott Koranda, Albert Lazzarini, and Roy Williams, Proceedings of the Fourth International Conference on Parallel Processing and Applied Mathematics (PPAM'2001). To appear in the Springer-Verlag in their Lecture Notes in Computer Science series.

- *Virtual Galaxy Clusters: An Application of the GriPhyN Virtual Data Toolkit to Sloan Digital Sky Survey Data*,
  Zhou, June 2002
- *Grid Laboratory Uniform Environment*,
  Ruth Pordes et al., June 2002
- *IVDGL Annual Report 2001–2002*,
  iVDGL Collaboration, June 2002

4.9 External Advisory Committee meeting

The External Advisory Committee (EAC) consists of nine members, listed in Section 2.3.7. The EAC met with GriPhyN and iVDGL on January 7, 2002, during the January 7-8 GriPhyN and iVDGL planning meetings. The committee charge, the talks presented at the meeting, and the EAC report, which was very positive, is available online[45].

5 Coordination with Other Data Grid Projects

5.1 LHC Computing Grid Project

The LHC Computing Grid[13] (LCG) project was recently launched by the CERN laboratory as a project to organize grid systems and other aspects of computing needed for its giant Large Hadron Collider (LHC) project. LCG is organized to draw on the broad range of data grid projects described above for grid middleware and practical experience in implementation and testing of grid systems. LCG is focused on one user community, but one that is very large and that has connections with virtually all of the data grid projects. As such it also serves as another focus of work on coordination and interoperability. LCG in turn is a strong participant in HICB and HIJTB.

Several areas are included in the LCG mandate, including distributed scientific applications; computational grid middleware, automated computer system management; high performance networking; object database management; security; and global data grid operations. iVDGL contributes to the Tier 2 centers of the US ATLAS and CMS software and computing projects. To them iVDGL contributes to the Grid Deployment aspects of the LHC Computing Grid project. The Experiment application Test Grids will be extended and merged with the LHC Computing Grid prototype deployments through collaboration with DataTAG and the US Experiment User Facilities.

The U.S. is represented in LCG management by several iVDGL members. Ian Foster is on the Software and Computing Steering Committee (SC2); Miron Livny and Ruth Pordes are on the Project Execution Board. Ian Foster and Carl Kesselman are additionally on the Architecture Task Force for the EU Data Grid project; Fabrizio Gagliardi from CERN is on the GriPhyN external advisory board. There continue to be many exchanges and visits of personnel.

5.2 DataTAG

The DataTAG project[10] funded in early 2002, is developing a wavelength-based transatlantic research network testbed between CERN and Starlight in Chicago. It is also working with SURFNet in Amsterdam to form a wavelength based triangular testbed for advanced network and Grid developments, using link speeds of 2.5 Gbps this year, with upgrades to 10 Gbps planned for 2004. Network developments will include level 2 services supporting Ethernet, and iSCSI (SCSI running over IP) that will allow distributed Grid resources to be integrated in novel ways.

The DataTAG network links will be the basis of a Grid testbed that will allow the U.S. Grid projects and European DataGrid to interoperate. DataTAG Work Package 4 and iVDGL work as peers in the deployment and interoperability aspects of their projects. Joint meetings are held between the projects on a regular basis. DataTAG and iVDGL are working together on delivery of GLUE (described in Section 6) and experiment-specific Grids that work across Europe and the US.

The major operational and technical goals of DataTAG are extend and complement iVDGL and the U.S. Grid projects and it is useful to summarize them here.

- To establish a global high performance intercontinental Grid testbed based on a high-speed transatlantic link connecting CERN and Starlight, with each side connected to appropriate national and international networks.
- To demonstrate advanced end-to-end network services across multiple domains, including: (1) Demonstration and optimization of reliable high performance data transfer over long distance high bandwidth-delay network connections, aiming for a minimum of 100 Mbit/s end-to-end throughput in the early stages and (5-10)x100 Mbit/s in later stages; (2) Provisioning and demonstration in a multi-domain network scenario of effective and consistent end-to-end network "value added" services including Quality of Service (QoS), managed bandwidth and advance network reservation; and (3) Estimation of the real usefulness of methodologies for traffic differentiation through tools for user-perceived performance measurement.
- To develop bulk data transfer validation and applications performance monitoring tools including: (1) Long-term infrastructure monitoring, (2) Monitoring of network performance delivered to the user (throughput and delay); (3) Application-level performance monitoring.
- To insure interoperability mechanisms between US and Europe Grid domains, and study the scalability issues which arise, including: (1) Interoperation of Grid security mechanisms; (2) Interoperation of information services to permit inter-Grid resource discovery.
- To provide the European research community with an intercontinental test infrastructure for advanced network and application test.

5.3 HICB and JTB

Starting in March 2001, an international data grid coordination organization has been created. Called the High Energy and Nuclear Physics Intergrid Coordination Board (HICB[46]), the group includes representatives of over a dozen grid projects from the U.S., Europe, and Asia, as well as people associated with ten or more current and future data intensive experiments. A total of four HICB meetings have been held, mostly in association with Global Grid Forum meetings. The next meeting is July 21, 2002, just before the GGF5 meeting in Edinburgh. In addition to presentations providing information about the many grid projects, meetings have focused on steps toward interoperability of middleware developed in the different projects and on common arrangements for licensing and sharing developed software.

An important working group created by HICB is the Joint Technical board (HIJTB), consisting of technically involved members of the different grid projects discussing steps toward interoperability and creating new working groups to implement those steps. The HIJTB meets monthly by remote conferencing and less frequently in face-to-face meetings when members are together. The major working group organized by HIJTB in its first few meetings is the GLUE (Grid Laboratory Uniform Environment), organized formally as a collaboration of iVDGL in the U.S. and the DataTAG project in Europe, which were organized with complementary foci to create an international testbed on which applications using data grid systems can be realized. The GLUE project has become an important component of the iVDGL program and is described in Section 6.

5.4 GLUE Effort

The Grid Laboratory Uniform Environment (GLUE) collaboration[47] was created in Feb. 2002 to provide a focused effort to achieve interoperability between the U.S. (iVDGL, GriPhyN and PPDG) and the European (EDG, DataTAG, CrossGrid, etc.) physics Grid projects with the goal to establishing Grid testbeds and production systems that span large world regions. GLUE's scope includes the definition of a set of software and test suites cover the basic interoperability requirements, and software will be deployed and tested on a testbed set up for this specific purpose. GLUE is described more fully in Section 6.

5.5 Trillium Collaboration: PPDG, GriPhyN and iVDGL

The increased coordination and synergy between the three US Physics Grid Projects – GriPhyN, iVDGL and PPDG – has led to the concept of the overall "Trillium" collaboration and presentation to our European and other partners. This strategy has strengthened the discussion of common and convergent goals; overall definition of scope and deliverables; and discussion of prioritization and milestones resulting in the movement of resources between the three projects as needed.

The fact that LHC is a significant focus of the US HEP program over the next few years, coupled with the fact that ATLAS and CMS are collaborators on all US Physics Grid projects, together with the shortfall in funding for the LHC S&C projects, together with the progress and successful collaborations in the US, has resulted in Trillium taking a major focus on the LCG project deliverables and milestones. This situation is leading to responsibilities and expectations for production deployment and middleware software support and maintenance that stress the expectations and plans of the Trillium projects to date. It is a measure of success of the projects that we are being approached and relied on for development, integration and support in these areas. It is however clear that the projects are stretching rather limited resources to carry out these duties. The iVDGL management must address funding and resource issues in the remainder of 2002.

5.6 Global Grid Forum

IVDGL participates actively in the Global Grid Forum (GGF[48]). GGF is a community-driven set of working groups that are developing standards and best practices for distributed computing efforts, including those specifically aimed at very large data sets and at high performance computing. Several members of iVDGL are actively involved in GGF activities.

5.7 Incorporation of New iVDGL Members

The iVDGL is expected to be a dynamic organization with new experiments, institutions, countries and disciplines being incorporated over time. Since Fall 2001 we have had many discussions with representatives of these additional activities and institutions. We expect to incorporate additional members starting in Summer 2002, but only after the completion of MOUs describing the duties of existing iVDGL experiments, activities and organizations. New members have been encouraged to seek additional funding as there are no iVDGL funds to support them.

The following Table summarizes the discussions we have had so far. Note that our European colleagues in DataGrid and DataTAG are considered to be members with whom we have established a working relationship. We expect to formalize the process of making them members of iVDGL after we complete the internal MOU process.

Table 8: Status of negotiations with new iVDGL partners

Member	Progress	Status
EU DataGrid	Existing working relationship	
DataTAG	Existing working relationship	
GridLab	Two discussions	Waiting for other MOUs to be completed
DOE Science Grid	One discussion	Need to elaborate some details, then develop MOU
BTEV expt.	Phone meeting, visit to Vanderbilt, other meetings	MOU being written; to be examined by iVDGL
ALICE expt.	Phone meeting, other meetings	Will probably join through EU DataGrid
CMS-HI expt.	Phone meeting, other meetings	Might join through CMS MOU
UT Arlington	Several discussions	Will probably join as ATLAS member, using ATLAS MOU
UT Pan American	Several discussions. Might join early as part of outreach effort	Will probably join as ATLAS member, using ATLAS MOU
Florida International University	Several discussions. Might join early as part of outreach effort	Waiting for other MOUs to be completed
PDC, Sweden	Several discussion, document discussing participation	Waiting for other MOUs to be completed
Brazil	Several discussions, meeting in Rio	Might join as adjunct CMS member initially
Japan	Letter of support, discussions at Grid meetings	Pending

Pakistan	Several discussions, two trips to Pakistan. Collaboration started between Caltech and NUST.	Pending
Romania	Trip to Bucharest May 2002, discussions with ministers. Collaboration started between Caltech and Polytechnica University.	Pending
Australia	Letter of support	No action yet
India	Discussion during CMS meeting in 2001	No action yet
China	Discussion during CHEP2001 conference, Sep. 2001	No action yet

6 GLUE Interoperability Effort

The following description of the GLUE effort summarizes a more complete document[49] still in preparation.

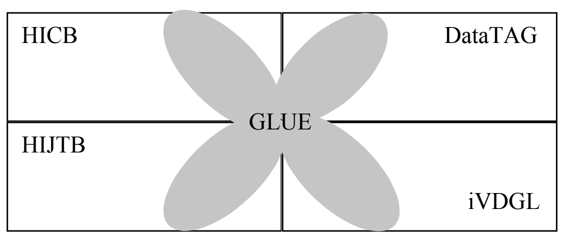

6.1 *Introduction and Goals*

The Grid Laboratory Uniform Environment (GLUE) collaboration was created in Feb. 2002 to provide a focused effort to achieve interoperability between the U.S. (iVDGL, GriPhyN and PPDG) and the European (EDG, DataTAG, CrossGrid, etc.) physics Grid projects with the goal to establishing Grid testbeds and production systems that span large world regions. GLUE's scope includes the definition of a set of software and test suites to cover the basic interoperability requirements, and software will be deployed and tested on a testbed set up for this specific purpose.

GLUE management and activities are provided by the iVDGL (U.S. side) and DataTAG (European side) projects specifically. The project reports to and gets guidance and oversight from the High Energy Physics Intergrid Coordination Board (HICB) and its Joint Technical Board (JTB), described previously.

Since the initial proposal for the GLUE project, the LHC Computing Grid (LCG) project was created at CERN to coordinate the computing and Grid software requirements for the four LHC experiments with a goal of developing common solutions. The LCG Project Execution Board (PEB) and Software Computing Committee (SC2) have endorsed the GLUE effort as bringing benefit to the project goals of deploying and supporting global production Grids for the LHC experiments. Accordingly, GLUE personnel in the U.S. and Europe are committed to deliver worth to the LCG and help understand the deployment and support issues resulting from its deliverables.

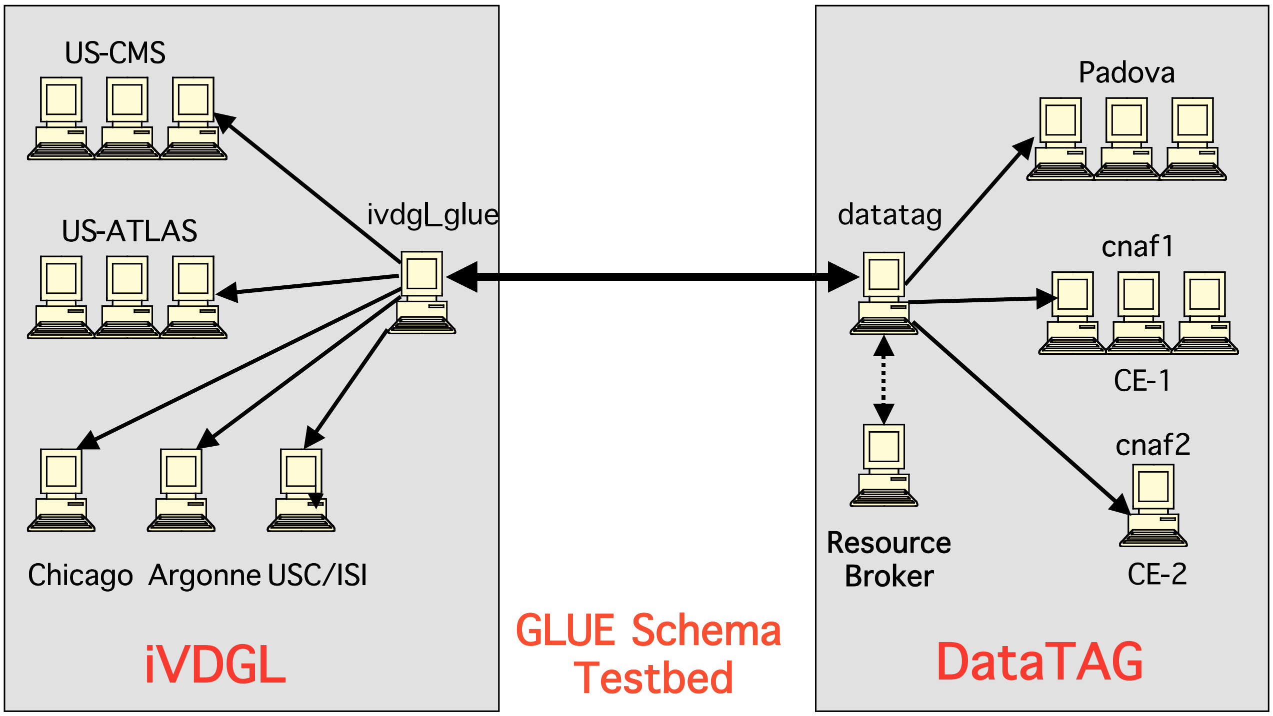

Figure 12: Diagram[50] outlining the chief components of the GLUE testbed.

6.2 Overview of GLUE Activities

The GLUE project includes a range of sub-projects to address various aspects of interoperability: tasks to define, construct, test and deliver interoperable middleware to and with the grid projects; tasks to help experiments with their intercontinental grid deployment and operational issues; establishment of policies and procedures related to interoperability; etc.

Once the GLUE collaboration establishes the necessary, minimum requirements for interoperability of middleware, any future software designed by the projects covered by the umbrella of the HICB and JTB must maintain the achieved interoperability.

The participating experiments will define use cases and deploy their applications to further test the interoperability of the underlying middleware components in more real-life and strenuous conditions. GLUE accepted solutions will thus be deployed on the Experiment, iVDGL and DataTAG testbeds to be validated by users. The collaboration includes tests on both dedicated hardware and as part of the experiment testbed systems. GLUE work includes:

- Definition, assembly and testing of core common software components of grid middleware drawn from EDG, GriPhyN, PPDG, and others, designed to be part of the base middleware of the grids that will be run by each project. GLUE will not necessarily assemble a complete system of middleware, but will choose the components that are necessary to satisfy minimum interoperability requirements. (Other projects may address some of these issues in parallel before the Glue effort does work on them)

- Development of configuration guidelines that would enable the existence of an interoperable Grid composed of EDG and GriPhyN/PPDG resources.

Experiments will be invited to join the GLUE effort to build and test their applications with the GLUE suite. The HICB and JTB will work with grid projects to encourage experiments to build their grids using the common grid software components.

Staging GLUE – both in terms of the resource sites available and the software infrastructure deployed – offers a natural mechanism to provide for tangible short term goals that GLUE will use for planning, scheduling and to see if we can actually meet its milestones. The current version of GLUE document only discusses Stage 1 – which is further subdivided into Phase I, II and III. Stage I is described further below.

6.3 GLUE Stage I Activities

The goal of GLUE Stage I is to achieve **i**nteroperability of the software components delivered and installed from the U.S. Grid Project distribution sites (currently VDT 1.1.1) and the EDG distribution site (TestBed 1 and incremental software releases as available). The deliverables of GLUE will be included in the EDG, VDT and other software distributions used by the Experiment application testbeds – and through these integrated into their testing and deployment efforts. What is sought from Stage I is basic, minimum functional interoperability and integration rather than high throughput and performance. We also expect to provide, acquire and/or otherwise stimulate the development of a permanent reference set of tests for new releases and new services as they are released following local integration and deployment tests, feedback from the experiment and grid project testbeds etc.

6.3.1 Phase I

- **Authentication infrastructure** – demonstrate the ability to perform cross-organizational authentication. While GSI provides the basic mechanisms that can be used for authentication of two parties, it is silent on the policy under which credentials are issued or acceptable. Thus authentication operations may fail because the credentials presented are not consistent with the policy of the entity to which they are being presented. For the proposed testbed to be viable, we must provide a minimal baseline policy that acceptable to all participating institutions and is trusted for a useful level of use at the participated sites (note that this does not preclude the use of other more stringent policies for operations that require a higher degree of assurance). Associated with this policy, each site must provide mechanisms for credentials compliant with this policy, along with administrative functions, such as certificate revocation.

- **Unified service discovery and information infrastructure** – demonstrate the ability to discover the existence and configuration of service offered by the testbed. This demonstration requires the definition of both a testbed wide discovery service, and places minimal requirements on information services supported by the participating sub-testbeds, or individual organizations. While the MDS 2.1 architecture and other information service prototypes in use define basic protocols for service discovery, they do not define the exact nature of the information schemas, nor do they specify how individual information services (e.g. GRIS and GIIS nodes) should be configured to form a distributed information service. The testbed-wide information service should define these structures so as to provide baseline (but not necessarily only) way of dynamically discovering services offered by the testbed.

- **Test of the Phase I infrastructure.** User jobs sent from an EDG location to run on an iVDGL location and vice versa. This requires the specific resource brokers to be able to discover and run a job "intercontinentaly". This also requires job scheduling and/or resource broker services to find the necessary resources both in EDG and iVDGL tesbeds. The tests will involve the following software: Globus Security infrastructure, MDS, GMA information services if available, GRAM, WP1 resource broker.

Many components, such as resource management, information services, authentication and authorization are needed to accomplish Phase I. It is anticipated that this phase might need to undergo further refinement of tasks and tests, for example defining a common schema is one of the necessary subtasks.

6.3.2 Phase II

- **Data movement infrastructure** – demonstrate the ability to move data from storage services operated by one site to another. This is a critical demonstration for the testbed in that the geographic distribution and potential for redundancy provided by the testbed become signification advantages with respect to technology demonstration. From the perspective of the testbed, we might define baseline configuration for GridFTP, Replica Location Service and GDMP enabled data services. This includes specification of schemas to define service characteristics, policy regarding scratch space, disk quotas, access control policy etc. [Not sure how this would be defined/implemented, this is more of a deployment issue].
- **Test of Phase II.** User jobs will be register a file on one EDG site and then retrieve it from an iVDGL site and vice versa. The "grid system" should be able to determine and use appropriate data files independent of where the storage resource is located. The software tested will include GridFTP, RLS and GDMP.

6.3.3 Phase III

- **Community authorization services** – demonstrate the ability to perform some level of cross organization, community based authorization. As part of this design process, we should determine if the community based access control model provided by CAS service or EDG Testbed 1 components are required for initial demonstration, and if so, include specification of the CAS and authorization servers as part of this task. Various high-level services such as replica management might be evaluated with the basic data storage and movement infrastructure proposed here at a future date.
- **Test of the complete computational service**. User jobs will be submitted from one of the resource brokers and will be sent to the computing element which is more suitable across the combined grid service. Logging and Bookkeeping, accounting and information retrieval should work across the domains. The software tested will include CAS.

In order to make a cross-continental, interoperable tested a reality, issues other than configuration guidelines and use of interoperable software need to be addressed. For the testbed environment, mechanisms by which participating organizations can make resources available to approved testbed users must be defined. The range of mechanisms could include generic accounts, community authorization or agreements to create accounts for every approved testbed user. Testbed operation also needs to include the definition of a user communities and participant. Membership and rights of these communities may be manually configured, however a more desirable solution would be to provide community authorization services as part of the baseline testbed services and to ensure that computational and storage resources are configured in such a way as to accept delegated credentials from the testbed communities. However, these issues are outside of the responsibilities of GLUE.

It is recognized that other software components might need to be evaluated and tested as part of the GLUE effort, however it is believed that these components are based on the software tested in Phases I-III and thus might be interoperable at the software level as the result of the GLUE effort.

- From DataGrid: WP3, WP1 (information services and resource broker), WP2, WP5, WP1 (full job scheduling and management) – in that order.
- From PPDG/GriPhyN – Globus V2.0, Condor v6.3.1, VDT 1.0 – in that order.

6.3.4 Progress to date:

Significant progress on the Authentication deliverable of GLUE has resulted in the ability to use US authorized certificates, specifically DOESG certificates, on the EDG testbed, and EDG certificates on the US Physics Grid Testbeds. Some details in areas of Certificate revocation remain. Support for interoperability of the US Tier 1 Site Facilities with these certificates is in progress through the PPDG SiteAAA project.

Significant progress on the Resource Discovery deliverable of GLUE has been made. A schema for Compute Elements encompassing individual compute hosts as well as more complex compute platforms, such as clusters has been agreed upon by the US and EU projects. This is a very significant step in the deployment of an international tested as the information presented to users and services across the Atlantic is now well defined and understood. In order for the resources to publish this information, new information providers still need to be developed. The GLUE Schema information providers and infrastructure will be delivered by the end of July. The GLUE testbed for testing this deliverable is in place at 2 sites in INFN in Italy, with ISI, ANL and some hosts from the Atlas and CMS testbeds in the US. Initial testing of the authentication, authorization, and accessibility of the information indexes has been done. Although this infrastructure is expected to be sufficient to test the interoperability of the US and European infrastructures in the computational domain, additional efforts are needed to characterize and put in place a testbed that would enable to design and test the interoperability in the storage management domain.

6.4 Coordination and Resources

6.4.1 Role of InterGrid Boards

The Joint Technical Board (JTB) is charged to understand interoperability, architecture, and middleware issues between the different data grid projects. Thus it is a natural forum for coordination, technical discussion, planning and review of the progress of this interoperability effort.

It is understood that the individual Grid projects will be working with the experiments on many different aspects of their data handling needs: e.g. CERN LCG project is creating a task force (RTAG) that extends the work done in WP8 of DataGRID to elaborate unified use cases and requirements between the four LHC experiments, to be implementable both with EDG and iVDGL middleware; PPDG and GriPhyN are working with the experiments on production (with attention to hardening, robustness and diagnosis) and development of support for Virtual Data (intelligent instantiation of data products) for deployment in the experiments data processing and analysis systems respectively; Several projects are working with BaBar and Run-II on support of their current data analysis and processing; etc. It is further recognized that the JTB will advocate and encourage the use of the software that results from the GLUE effort and will oversee that any new software developed within the projects will not break the interoperability of the middleware.

7 Facilities Work Team Progress

7.1 Scope and Goals

The Facilities Group has two primary responsibilities: identifying and defining iVDGL sites and arranging hardware procurements at those sites. The physical locations of the sites are relatively easy to identify. The proposal defined a group of funded iVDGL locations from the core experiments at Boston University, Caltech, Hampton University, the University of Indiana, Johns Hopkins University, Penn State, Salish Kootenia College, the University of California San Diego, the University of Florida, the University of Texas at Brownsville, the University of Wisconsin Milwaukee and Fermilab/SDSS. There were also three unfunded sites identified: Fermilab/CMS, Argonne, and Brookhaven national labs. One of the initial charges for the Facilities group was to identify additional sites from the four experiments defined in the proposal as well as sites from related experiments, like Alice, CMS Heavy Ion, D0 and BTeV. The University of Texas at Arlington, associated with the DO, experiment has been an active participant since the workgroup group was formed. We expect additional members and sites as iVDGL progresses.

Defining what is means to be an iVDGL site is more difficult. There are the allocation issues about how often and at what percentage the Facilities will be used for iVDGL common work and how much warning is needed. This is made more complicated by the existence of substantial amounts equipment at funded and unfunded sites unrelated to iVDGL equipment, which can be used for laboratory work. There are also technical issues about the way computing facilities are presented to the grid laboratory. We have largely left the details of the allocation issues to the authors of the MOU's and have tried to concentrate on the technical issues. We have been attempting to address whether it is possible and desirable to define a site not as a list of Facilities, but as a collection of services. For example instead of defining a site as having a number of processors, the site would publish that it has a batch queue or many batch queues with agreed upon configurations. This should make the facilities more manageable in the context of the grid laboratory by creating well-defined interfaces and boundaries.

The second primary responsibility of the Facilities group is arranging hardware procurements. We are tasked with ensuring that equipment purchased with iVDGL funds is compatible and consistent between iVDGL sites. We have reached consensus on the use of Intel and AMD systems with the Linux operating system. To ensure that equipment purchased for iVDGL sites is compatible with experiment applications we are compiling a set of executables that can be run by non-experts to evaluate new hardware. We are also creating a web page of generic testing tools and results. We will compile a list of hardware purchased, vendors used, experiences, and prices (if possible). We will attempt to compile a complete list of equipment needed to assemble a small facility, which could be used by non-experienced cluster builders. These lists are only valuable if they have been recently updated, so we will release an updated list regularly. We hope to strike a balance of supporting experienced facility managers without overly constraining them, while giving inexperienced cluster builders the support they need to succeed.

To meet all aspects of our responsibilities, the group has organized its activities around four co-ordinators. The first coordinator is the iVDGL Site Survey Database Coordinator: Jorge Rodriguez from UFL. The coordinator has now created and will maintain a database with up-to-date information about the different iVDGL facilities. The database has a web front end that displays information in a tabular form.

The second coordinator is the iVDGL Performance Test Coordinator. This individual coordinates testing of hardware components such as CPUs & chipsets, disks and disk controllers, and other hardware. This involves testing existing hardware, getting loaner hardware, arranging vendor tests, and finding sites to purchase equipment to be used for testing. The results are circulated and posted on the web. Since several of the iVDGL institutions have not yet completed their initial hiring, no specific individual has been appointed to this role. Until that is done, Allen and Fisk have been doing this work themselves.

The third coordinator is the iVDGL Benchmark Coordinator: James Letts of UCSD. This coordinator's job is to organize the collection and web posting of a set of simple statically-linked standalone executables that can be used for hardware testing. The results of the testing are posted on the web.

The fourth coordinator is the iVDGL Acquisitions Coordinator. This coordinator obtain acquisition plans from different iVDGL sites, maintains a global listing of who will be buying what, and when, and coordinate these purchases. This role has also not yet been filled, but we anticipate filling it by the end of the summer, when the acquisitions phase of iVDGL begins in earnest.

7.2 Milestones

7.2.1 April Milestones

The Facilities group expects to have published the initial set of experiment executables for benchmarking new hardware. Already executables from LIGO and CMS have been made public. The performance numbers for various hardware platforms will be published on the web page as they become available.

The initial set of scripts for performing a reasonably broad set of generic disk performance tests was published to the web page.

iVDGL Facilities members will give access to limited computing resources at their sites to designated contacts from other experiments. This will be used to determine the effort required to run the applications of one experiment at sites controlled by another experiment.

7.2.2 May Milestones

We will attempt to establish software environments for the experiment code at hardware facilities controlled by other experiments. Hopefully, non-privileged users can establish functional software environments. There is the possibility that some configuration will need to be performed by privileged users.

Site administrators will establish well-defined grid enabled batch queues for experiments using the Globus Gram interface. At the most basic, there will be an ATLAS-queue, CMS-queue, LIGO-queue, and SDSS-queue, but it may be necessary to have additional queues for specific experiment tasks.

7.2.3 June – September Milestones

After the MOU's have been signed, a preliminary schedule of hardware procurements should be released. It does not make sense to schedule hardware procurements more than six months in advance expect in completely generic terms, because the desired hardware will change very rapidly.

Members of the Facilities team will work with members of the Laboratory Operations team to design a schema for monitoring and discovering site services using information tools like MDS.

Facilities members should work with members of the Applications teams to configure grid enabled batch jobs to deploy the basic elements of the required software configurations from the batch job itself. This should reduce the amount of interaction needed between the site managers and the representatives from the experiments.

The first list of hardware for a complete iVDGL cluster should be released by the beginning of the second project year. The hardware procurements in the first year are primarily additions to existing facilities.

By SC2002 we expect to have demonstrated the use of grid enabled computing facilities outside the direct control of an experiment.

7.3 Current Progress

The iVDGL Facilities group has established a regular meeting schedule and had one face-to-face workshop[51] in combination with the LHC Common Projects Facilities meetings. We have had reasonable representation from almost all iVDGL sites at the regular teleconferences.

The Facilities group has carried out comparison testing of several different IDE RAID disk arrays. These are the most cost-effective large disk arrays. The systems tested include a standalone (Promise SX8000) IDE-to-SCSI RAID array and a internal 3Ware 7850 controller card. The conclusion was that the 3Ware card had approximately three times the performance as the SX8000, but it did not scale as well to very large file systems (more than 2 TB). Additional experiments are being carried out (at UWM) to see if performance can be improved by using LVM, hardware SCSI raid, or Linux software RAID to improve performance by striping across multiple SX8000, and to understand how file system type (ext3, ReiserFS, XFS, etc) influences this performance. For this purpose, working with Promise Technologies, four SX8000s are being acquired at cost. Each one will be populated with 1.2 TB of disk. This will yield a 4+ TB file system at below $16k in cost. In addition, the design will scale up to tens of TB on a single host.

We have also carried out testing of a number of dual CPU platforms, including one based on the latest ServerWorks chipset, that has four independent PCI buses and exceptional IO performance, and one based on the Athlon 760 chipset, that seems to offer exceptional floating-point performance at low cost. The facilities group has learned that the latest dual CPU systems do offer effectively twice the performance of the earlier generations. This appears to be because of improvements in the memory bus bandwidth, or the use of independent memory buses for each processor. This makes such dual CPU systems, for the first time, more cost effective for large general-purpose clusters.

Members of the group have posted two standalone CPU benchmark programs on the Facilities Group web site. One of these benchmarks (contributed by the LIGO Scientific Collaboration) is a test of FFT (Fast Fourier Transform) performance. The other (contributed by the CMS Collaboration) is an ORCA Benchmark. These have been used in the CPU testing described above.

Yujun Wu has posted a standalone Disk Performance benchmark program. This tests the Read and Write speed on character and block transfers in both sequential and random IO. It also tests file creation and access speeds. This benchmark has been used in the testing described above.

Jorge Rodriguez has begun populating the site database described above for the site coordinator, as described in Section 4.7.

The group has begun discussing purchasing strategies for large clusters. One method (that was successful for example at UWM) was to use closed sealed bids. In order for this method to work, however, an extremely detailed and unambiguous hardware specification is needed, which leaves no "wiggle room" for the vendor to make unacceptable substitutions. The Facilities Group is currently circulating several such specifications.

The Facilities group has not made much progress on establishing contacts, queues, and entry points for common iVDGL facilities and grid services. We expect this will progress more quickly as the iVDGL demonstrations are planning and executed.

7.4 Future Plans

The Facilities Work Team will participate in the execution of the iVDGL demonstrations for Fall 2002. We will work with the iVDGL coordinator and Operations group to configure and run a truly distributed computing environment.

All of the funded participating experiments have aggressive acquisition plans over the next year to prepare for data analysis or data challenges. The Facilities Team will continue to work together to specify equipment that is acceptable to the participating experiments in terms of configuration, performance, and reliability. We will continue to perform benchmark tests and to use our combined purchasing power to work with vendors.

8 Core Software Team Progress

8.1 Scope and Goals

Originally the primary responsibility of the iVDGL Core Software Working Group (CSG) was to evolve the Virtual Data Toolkit (VDT) produced by the GriPhyN project into a production-quality software suite. Early on, however, the CSG realized that its scope and goals were tightly coupled to those of the GriPhyN VDT effort. We realized that turning the VDT into a production-quality toolkit required tighter integration between the GriPhyN and iVDGL efforts, and so we proposed that a subgroup be formed that spans the GriPhyN and iVDGL projects and whose purpose would be to produce, release, and support a production-quality VDT. Members of this new VDT Working Group would are composed of individuals from both GriPhyN and iVDGL. The key idea is to better integrate the design and construction of the VDT with the deployment, testing and support efforts.

The combined VDT group is led by (1) Miron Livny, liaison to GriPhyN/iVDGL management, (2) Alain Roy, focusing on development and packaging, and (3) Scott Koranda, focusing on deployment and production.

The scope of the VDT Working Group includes (a) development and packaging of the VDT, (b) development of packaging technology (Pacman), (c) developing testing suites for installation testing by users, interoperability testing between VDT components, and interoperability testing between the VDT and the EDG, (d) deployment and maintenance of an integrated user support structure including mailing lists, bug tracking, and web sites, and (e) development of documentation complementary to existing VDT component documentation.

8.2 Milestones

Key milestones include:

- The first release of the VDT (version 1.0) was in early March, 2002. Version 1.0 included the beta version of the Globus 2.0 Toolkit, Condor version 6.3.1, Condor-G 6.3.2, GDMP 2.x, and ClassAds 0.9.2. This first initial release was packaged using a specially modified version of PACMAN and included three flavors: VDT server, VDT client, and VDT developer.
- The VDT web site was deployed in conjunction with the release of version 1.0. The site is hosted at the University of Wisconsin-Milwaukee and continues to be updated. Included on the web site are links to all the software for downloading, a list of system requirements, installation instructions, an explanation of why we chose Pacman for packaging, and links for getting help and support with the VDT. The web site also functions for organizing the VDT working group and includes a list of group members along with the bi-weekly teleconference schedule and agenda.
- Also in conjunction of the release of version 1.0 two email lists for support and discussion of the VDT were set up. The two lists are vdt-support@ivdgl.org and vdt-discuss@ivdgl.org.
- Soon after the release of version 1.0 and the deployment of the vdt-support@ivdgl.org mailing list the list was configured to directly forward mail is into a bug-tracking system at the University of Wisconsin-Madison. Staff at UW-Madison respond and provide support for VDT users.
- In April of 2002 version 1.1 of the VDT was released. This version included an update from GDMP 2.x to GDMP 3.0, along with many enhancements to Pacman.
- Also in April of 2002 a decision was made between Saul Youssef, the creator of Pacman, and the VDT group at UW-Madison to jointly develop Pacman. A CVS archive for development was created at UW-Madison.
- In May of 2002 the VDT working group posted on the VDT web site a Release Policy, detailing how in the future the VDT will be released in two series: a development series and a more stable production series. Details of the Release Policy can be found at the VDT web site.
- Development of a testing mechanism to verify a VDT installation began in earnest in June 2002.
- In early July 2002 version 1.1.1-1, the first in the newly announced development series, will be released and will include the final release of the Globus 2.0 Toolkit.

8.3 Current Progress

A number of groups have installed the initial versions of the VDT and the working group is processing comments and feedback that have been made regarding the VDT, its packaging and installation. Detailed discussions have taken place between the VDT working group and representatives of both the CMS and ATLAS experiments on using Pacman to deploy all necessary software for a Tier-n site, rather than just core grid software.

We believe that a close collaboration between the teams will help us to on one hand to improve the overall functionality and robustness of Pacman and on the other hand broaden the spectrum of software distribution and installation "religions" that Pacman supports. Having CMS and ATLAS use Pacman is in our opinion a very important goal that will have significant implications on the success of "the Grid" concept in the HEP community.

Development of tests for verifying an installation has only begun. At this time only one person (James Letts from UCSD) is able to devote significant time to developing testing scripts. In order to make quicker progress we would like to find at least one additional person who could spend at least 50% of his or her time developing testing scripts.

The VDT working group holds bi-weekly teleconferences to discuss all aspects of development and deployment of the VDT. A schedule of meetings along with meeting agenda can be found at the VDT web site. We have held 8 teleconferences so far, and will continue meeting bi-weekly throughout the rest of the year.

8.4 Future Plans

Immediate future plans for the VDT working group include:

- Incorporation of the Chimera Virtual Data System developed as part of the GriPhyN project into a development release of the VDT, followed by a promotion to a stable VDT release after appropriate testing.
- Continued discussion with CMS Production regarding the use of Pacman for installation of a basic CMSIM-capable farm with Impala, Boss, and a CMSIM Darball, followed at the end of the summer CMS production with a more formal discussion about CMS adopting Pacman and feeding into its development.
- Continued development of testing scripts for verifying a VDT installation. Development of tests for component interoperability and interoperability with EDG software.
- Development and testing of automated configuration scripts to simplify configuration of the VDT.
- Drafting of initial documentation including VDT Getting Started, VDT FAQ, VDT User Guide, and the like.

9 Laboratory Operations Team Progress

9.1 Scope and Goals

Laboratory Operations is responsible for developing and supporting both an Information Infrastructure (II) and a Service Infrastructure (SI) for the iVDGL.

9.1.1 Long-term relationship between iVDGL Operations and other VOs

The Laboratory Operations Team will develop a model for infrastructure that may be deployed, either partially or as a whole, by individual VOs or Laboratories within the iVDGL framework (or outside the iVDGL for that matter, e.g. DTF). At this point in development, it seems reasonable that each VO will adopt/install a VO-specific Operations Center, using tools and techniques learned from the development of the more general iVDGL Laboratory Operations Center. In addition to supplying tools and overall operations guidance to the VOs, the iVDGL central Laboratory Operations Center may act as a top-level Operations Center, integrating the views of the individual VOs, participating in the iVDGL.

However, at this point in the project, the focus of work within the Laboratory Operations Team is to develop the necessary tools and procedures for the operations of a generalized gird structure. A description of our work to the present follows.

9.1.2 Information Infrastructure

The II is represented by the iVDGL Grid Operations Center (iGOC). The iGOC will maintain both operational information (what resources are available on the Grid and their status) and general information (who are the people and institutions involved in the iVDGL; how to participate in the iVDGL). The iGOC will act as both the private face of the iVDGL (face presented to the operational portions of the iVDGL) and the public face of the iVDGL, a role shared by the http://www.ivdgl.org web site. The II will have both an electronic component, represented by the iGOC web site and a physical component represented by the iGOC NOC staff (at present non-existent – see Resources)

9.1.3 Service Infrastructure:

The service infrastructure (SI) consists of services that are necessary to provide a view of the iVDGL testbeds.

- Security Infrastructure that enables users and resources to authenticate to iVDGL resources. As part of this work, it has been agreed at the project level that iVDGL resources will accept DOE Science Grid certificates. We have submitted a document describing the iVDGL policies to the DOESG Policy Management Board and it has been accepted with small changes. As a result the DOESG has agreed to provide certificates for the iVDGL community. ISI is developing and evaluating tools to manage the user community.
- Collaboration in an instantiation of a test grid for iVDGL Laboratory – this work has begun and is described as US-iVDGL-1 (a US only testbed) and iVDGL-2 an international testbed, encompassing the US and the EU with possible Asian participation.
- Grid discovery and status including device/service discovery, status and performance monitoring
- Grid Laboratory Error and Alarm presentation and handling
- Providing security solutions to the iVDGL, including the ability of testbed sites to acquire the information necessary for a gridmap file
- Other information enabling iVDGL sites to be appropriately configured

When developed and tested the SI will be maintained and operated by the iGOC.

9.2 Milestones

1. Development of iGOC web site (information infrastructure). This project is ongoing and results are available on the Laboratory Operations page, iGOC link, from the iVDGL main page.
2. Hiring Senior Grid Technologist. Permission to hire has been received from Indiana University, an ad developed and placed in a variety of publications. An acceptable candidate has been located and an offer will be extended on/before July 1, 2002.
3. Development of access control structure. ISI is developing an authentication control structure at ISI for later transfer to the iGOC. Preliminary documentation describing how users, administrators and EU participants can gain access to the iVDGL is available on the iGOC web pages.
4. Development of Lab Ops testbed. Provide an infrastructure in which hosts can register to the tested (setup a GIIS for the testbed), provide any additional information needed to configure the testbed host for participation in iVDGL. This is in progress with an August 1 completion date for US-iVDGL-1.
5. Development of a Laboratory Operations Working Group with a mailing list, regular meetings, minutes and a web page. Done.
6. John Hicks is working with Manuela Campanelli on development of iVDGL activities for REU students within the iGOC and Laboratory Operations in general. See the E/O report for details.
7. Participation in the planning and implementation of the Laboratory milestones – US-iVDGL-1 and iVDGL-2. Ongoing with a 1-August deadline for US-iVDGL-1 and a 1-November deadline for iVDGL-2.

9.3 Current Progress

9.3.1 Web site

This is an ongoing project. Development continues; good progress being made; information for many iVDGL sites is available online. Broad router monitoring information is available for all Abilene connected sites.

9.3.2 Hiring

Job posted. Candidate proposed. Hope to hire by 1-July.

9.3.3 Authentication

A document describing a support infrastructure for obtaining DOE certificates for use in iVDGL has been circulated. The document was written in compliance with the DOESG Certificate Policy and Certificate Practices Statement. Documentation is available on the IGOC web page.

9.3.4 Grid Development Laboratory

Initial testbed sites are being configured (ISI and ANL), a host (giis.ivdgl.org) has been configured to provide testbed information services. A dialog with the GLUE schema group has been established.

The Laboratory Operations web page has been developed and is linked to the iVDGL main page. Regular telecons are scheduled for the next 3 months (bi-weekly). Minutes are taken and available online.

9.4 Future Plans

1. Senior Grid Technologist position within Laboratory Operations must be filled. It will take the full attention of a professional employee to fully develop the II and to integrate the SI into the iGOC. Progress will be limited until this position is filled. July 1 deadline.
2. Closer linkage with performance monitoring group is necessary to develop mechanisms to display grid status and performance bottlenecks. John Hicks, from Indiana University will lead this effort, linking with Jenny Schopf's group. Ongoing.
3. Population of "resources" section of iGOC pages for testbed is necessary both to provide information as to the current composition of the iVDGL grid and to develop automated mechanisms for transferring GID information into the iGOC. This is directly related to the development of an iVDGL GIIS (see #4 above). Deadline of 1-August.
4. Development of an information structure within the iGOC to enable people and resources to participate in the testbed. Deadline of 1-August.
5. Provide a set of web-based interfaces for obtaining information about the testbed resources (ala Nordugrid). Deadline of 1-November.
6. IGOC Trouble Ticket system (TTS) has been activated. At present this is a web form which is e-mailed to williams@iu.edu. In the future the TTS will direct e-mail into a structured trouble management process. This will include (a) Knowledgeable 1st level support at the iGOC; (b) Tools required by the iGOC to deliver 1st level support (such as performance monitoring tools and site contacts); (c) Development of a 2nd level support process perhaps based on the site contacts; and (d) A set of site and management reports summarizing the TTs for various reporting periods. Note – this is a Trouble Ticket System, which is linked to operational problems (troubles). As a separate part of the iGOC a "user assistance system" will be developed to assist users with installation and preliminary operation issues (e.g. VDT setup, CA issues and such). Some preliminary documentation for CA issues is already available.
7. Consider issues of GIIS-GIIS communication among multiple virtual organizations.

In conjunction with the Facilities WG, develop a scheme for discovering site services and populating the iGOC Operational Information section (perhaps using GANGLION in combination with MDS).

10 Applications Team Progress

10.1 Scope and Goals

The applications work team carries out the following activities:

- Identify and document the various applications groups participating in US-iVDGL.
- Outline in document form a plan for Grid software integration for each

- Create a list of the application participants, including name, institution, phone and e-mail. For each application taking part, identify the contact person
- Develop and document a plan for inter-site tests for the first 6 months
- Create a work team list including name, institution, e-mail and phone number. Create a web page for all the above information

10.2 Milestones

The high-level milestones include the following. These will be broken down into lower-level milestones in a further refinement.

- Document the applications to be run on the Sites (June 1)
- Integrate the applications with the Core Software and test on experiments owned iVDGL sites (Sept 1)
- Deploy experiment applications on other experiment sites (Nov 1)

10.3 Current Progress

The work plan for the Application team is aiming at iVDGL Demonstration Milestone (US-iVDGL-1) in August 2002, with the goal of running and supporting the Laboratory for some days. The role of the Application Team is to plan and prepare Applications of at least two experiments to be run on iVDGL testbeds of at least two experiments.

A work plan concerning integration, deployment, support is being prepared:

- Establish that applications run on exp. own sites
- Look at packaging, configuration needs,
- Ensure that adequate configurations, diagnose tools and support is available
- Expand those installations to non-exp owned sites

For the test a sensible “lab-test” plan will be defined, including milestones and metrics to establish successful running of “lab test”. This may involve requests to individual experiments or other work teams for displays and graphs, documents, etc.

It appears now that all four major collaborating Experiments, ATLAS, CMS, LIGO, SDSS, will have something to contribute to the August US-iVDGL-1 milestone, so that there will be probably enough material to make an interesting lab-plan. Specifically, there is a strong interest from both the CMS and the ATLAS user facilities project to bring the ATLAS and CMS Application test-beds together for such a test. In addition to US-iVDGL-1 this could be a milestone on the way to the LHC Computing Grid production grid prototype in the US.

For CMS the applications are being documented and requirements have been listed in some detail. The US CMS test-bed has an application running now, which is the simulation of events in the CMS detector, including pile-up, and reconstruction of the simulated detector signals to physics objects. The components of this application include Impala, BOSS, GDMP, MOP, RefDB and MCRunjob.

ATLAS has successful deployed their application on the ATLAS testbed. A first inter-operability test has been started in running the ATLAS application on the CMS testbed.

Also for LIGO an application is a real possibility. They propose a code[52] that calculates "excess power statistic for detection of burst sources of gravitational radiation". For this calculation, jobs will be farmed out together with replicated data to do FFT analysis. The application is a "simple" executable, ideal for the Condor environment, and GridFTP would be used for moving data.

The SDSS applications group (Jim Annis, Fermilab/SDSS, and Yong Zhao, UC, along with Jens Voeckler, Mike Wilde, Ian Foster and Steve Kent) has developed an application that is ready to be deployed on the iVDGL. Progress was made on a GriPhyN/SDSS challenge problem, the galaxy cluster finding challenge problem described in a recent GriPhyN document. The science problem was to locate clusters of galaxies in the SDSS data, the computational problem being roughly compute- and I/O-balanced, and the computer science research was to describe this a virtual data problem. The Chimera virtual data system of Voeckler, Wilde, and Foster was used, as both the grid compute engine and a data providence system. A virtual data grid was constructed from a pair of master machines at Argonne, the Fermilab/SDSS iVDGL cluster TAM, and the University of Chicago Condor cluster. 200 sq-degrees of sky were analyzed, out of 40,000 on the sky, 7000 in the SDSS survey area, and 2500 available on disk. That means that enough data was being analyzed in this fasion to be interesting, but not enough to be ground breaking.

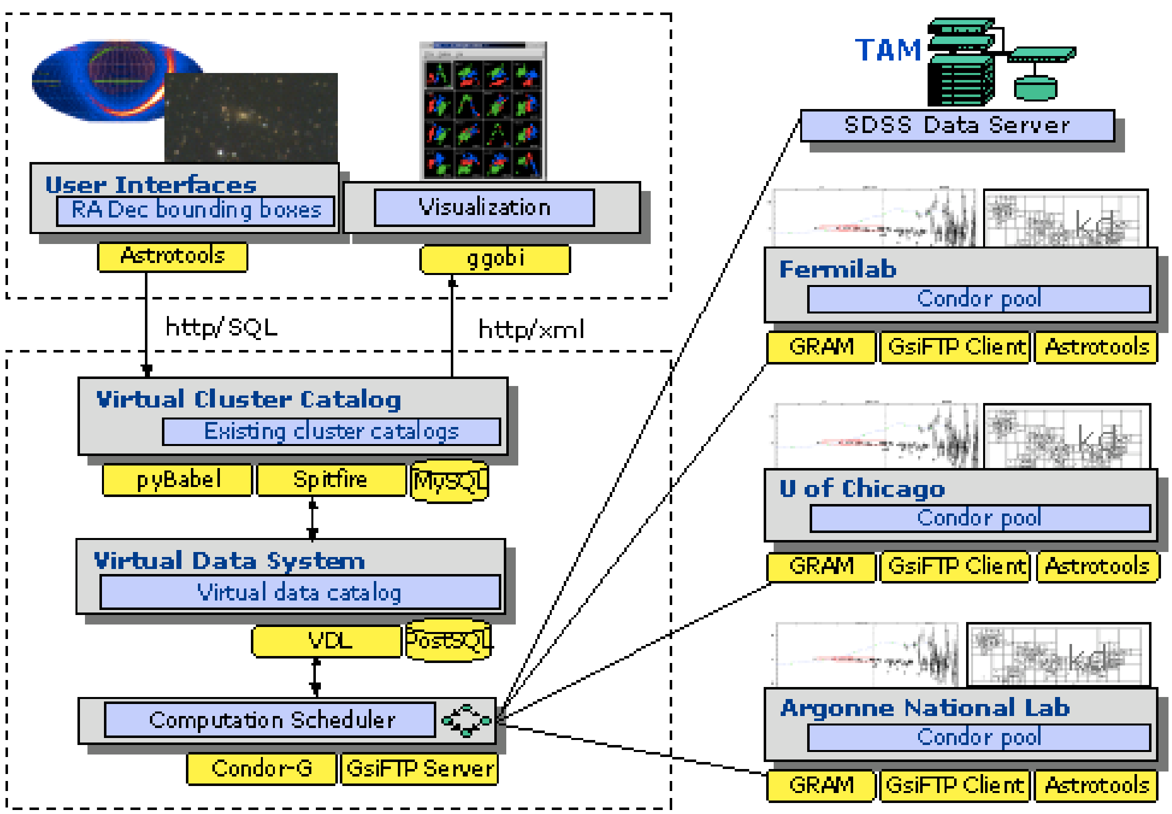

Figure 13: Architecture for integration of Chimera into SDSS environment for cluster-finding.

The translation of the science code into the virtual data language of the Chimera system provided some interesting results. The DAGs (a DAG, or Directed Acyclic Graph, describes schematically the steps that compose a job) computed were very complicated, as figure X shows for an absolutely trivial amount of data. The shape of the DAG had implications for the run time of the job. On TAM, a controlled environment, the time to completion of the nodes in a DAG were deterministic. On the UC Condor pool, the computers of which were in use by their owners, the time to completions of the nodes were stochastic, with the time to completion of a set of DAG nodes driven to 3 times that on TAM by individual long running nodes on the Condor pool. Grid weather. This has implications for scheduling: if the iVDGL grid sites are not dedicated entirely to a given application, then job completion times will have to be determined by running subsets of the job on the actual grid sites.

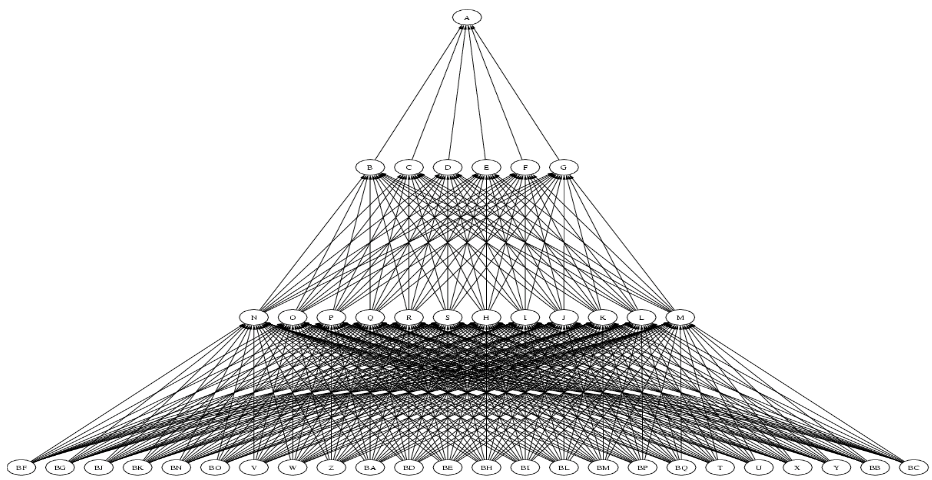

Figure 14: DAG for cluster identification workflow.

10.4 Future Plans

The SDSS application group will move to analyzing a scientifically interesting 2000 sq-degree area on the sky. This will be done by moving on to iVDGL sites proper: the iVDGL/CMS sites at the University of Wisconsin and University of Florida will be incorporated into a new virtual data grid to perform the computation. This will take some work: the science code has evolved, the Chimera code has evolved, we will for the first time be solving the problem of code distribution to sites we are not closely connected to. As our technical paper submission to SC2002 was accepted, we expect to a) perform the calculations so that the SDSS can use the results, and b) keep the virtual data grid alive long enough to use as a demo at SC2002. We would very much like to use this application to test code migration strategies as well: we believe that code should migrate along with data to grid sites for processing, as opposed to setting up a cluster with the science code and maintaining it across code version changes.

11 Education and Outreach Team Progress

11.1 Scope and Goals

The iVDGL Education and Outreach (E/O) program enhances grid-related research capabilities of faculty and students at other universities and institutions. In particular, it intends to promote learning and inclusion via the integration of faculty and (K-12, undergraduate, and beginning graduate) students at a diverse set of minority and under-represented institutions (MSI) into the scientific program of participating physics and computer science experiments.

In Spring 2002, three MSIs (Tier3 centers) received funds from the iVDGL project for hardware and personnel to upgrade or construct small (32-node) clusters, thus bringing a large number of additional minority students directly into contact with large-scale grid research. These Tier3 centers are: The University of Texas at Brownsville (UTB), a Hispanic Serving Institution (HSI) involved in LIGO research (LSC member Institution), Hampton University (HU), a Historically Black College and University (HBCU) involved in ALTAS/CMS research, and Salish Kootenai College (SKC), a Tribal College that has recently joined the LSC (LIGO Scientific Collaboration).

Associated with each Tier3 center are also E/O work team members and several undergraduate and graduate students. To facilitate the coordination of the E/O activities among the Tier3 centers, Manuela Campanelli (UTB) is leading the E/O program for both the GriPhyN and iVDGL projects, integrating the two into a single E/O effort. Keith Baker (HU) is co-leading this effort. Other Members are: Tim Olson (SKC), and several undergraduate and graduate students (e.g. Jose Zamora and Sean Morris at UTB, Howard Brown at HU, etc.).

Team members collaborate to the overall effort by actively interacting with each other (through mail lists, regular telecons and meetings), leveraging existing E/O programs (such as SkyServer, EOT-PACI, QuarkNet, etc.), and collaborating with the E/O efforts in other national projects (e.g. PPDG, etc.) and international projects in Europe (e.g. DataGrid, etc.) and possibly in Asia and South America.

The scope of the Education and Outreach (E/O) program of the iVDGL project is to enhance grid-related research capabilities of faculty and students at other universities and institutions. In particular, it intends to promote learning and inclusion via the integration of faculty and students of a diverse set of minority and under-represented institutions into the scientific program of participating physics and computer science experiments.

Not only this, but also all senior personnel are meant to be engaged in this mission, with each person committing to lecturing and mentoring E/O activities at other institutions and universities.

11.2 Milestones

11.2.1 UTB contributions to outreach

UTB has many contributions to iVDGL outreach. These are are integrated with those of the GriPhyN project:

- UTB administration approved the hiring of a new faculty member, Manuela Campanelli, to serve as Education and Outreach coordinator for GriPhyN (Fall 2000-2001). Campanelli also participated in the iVDGL project, leading the E/O program for both GriPhyN and iVDGL projects (Fall 2001).
- Development of an E/O web site (see Figure). This project is ongoing and the results have been altready demonstrated at several All-Hands and EAC meetings (Fall 2001).
- Development of a E/O Working Group with a mailing list, regular meetings, minutes and a web page (Fall 2001).
- The first E/O meeting was organized at UTB (Spring 2002). UTB volunteered to host a facilities Working Group meeting (Fall 2002), and a GriPhyN/iVDGL All-Hands meeting (Fall 2003).

- Grid-enabling the UTB Linux cluster by two undergraduate students (Spring 2002). Intallation of newer releases and further testing of the VDT is an ongoing project. The final goal is to be able to participate to a grid testbed together with other major GriPhyN and iVDGL institutions (Fall 2003).
- Introducing two undergraduate students and a graduate student to grid computing by letting them participate to GriPhyN/iVDGL activities, meeting and conferences (Spring/Fall 2002).
- Exploring ways to collaborate with the EOT-PACI effort (Spring 2002).
- Preparation of a REU proposal to be submitted to NSF (Spring 2003).
- Creation and Editing of a joint `Newsletter' for the GriPhyN, iVDGL and PPDG projects (Spring 2003).

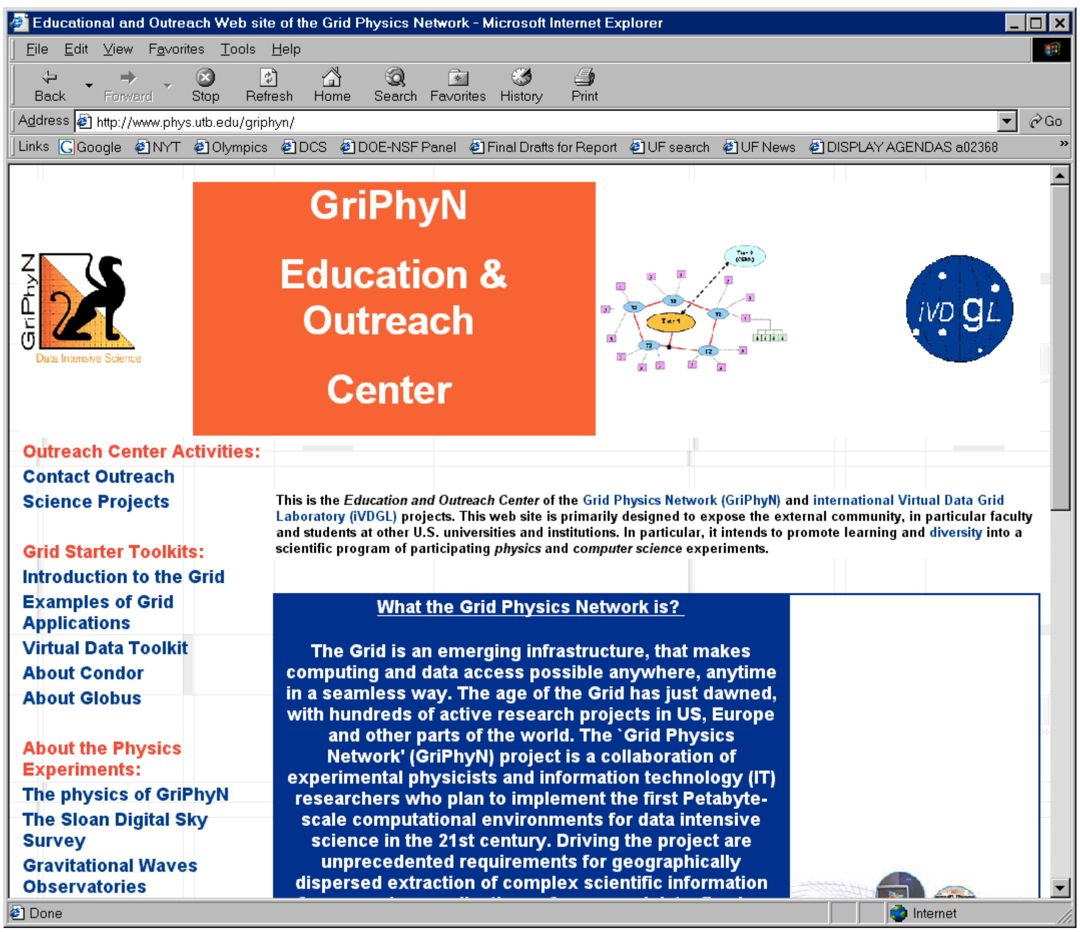

Figure 15: The iVDGL-GriPhyN Outreach page.

11.2.2 HU contributions to Outreach

HU contributes to grid development by being a testbed for iVDGL performance. Key short-term milestones are:

- Develop ways to collaborate with Quarknet and ATLAS outreach (Summer/Fall 2002).
- HU has expressed interest in hosting the next E/O meeting (possibly in Fall 2002) while high school teachers participating QuarkNet activities are still on campus.
- Complete GEANT4 simulation of simple detector compared with data from bench tests at HU (Fall 2002).
- Complete first physics simulations using ATLFAST++ (Fall 2002).
- Report written by K-12 student on iVDLG work at HU (Summer 2003).

11.2.3 SKC contributions to iVDGL outreach

- Tim Olson introduced SDSS SkyServer into teaching courses (Spring 2002). This is an ongoing project.
- Olson participated at Linux construction workshops (Spring 2002).
- SKC is planning the construction of a 32-nodes Linux cluster (Summer 2002).
- The actual contruction of the cluster is planned for year 2 (Spring 2002).

11.3 Leveraging Existing E/O Programs

11.3.1 SDSS Skyserver

SkyServer (see Figure below) is a project lead by Alex Szalay (Johns Hopkins University) and Jim Gray (Microsoft) to provide free Internet access to the public of SDSS data for both astronomers and for science education. This project is well under development, and it appears to be an excellent educational tool, containing some key reading material for teachers and several correlations to National Education standarts. Jordan Raddick (Web designer), who is the contact person for E/O in SDSS, has been working very closely with a high school teacher associated with the Fermilab/SDSS group, Robert Sparks, in order to develop ideas for the course material. Material to data has been at a relatively high level, more suitable for community colleges than high school, and Sparks is currently developing materials aimed more directly at high school students.

During the E/O meeting at UTB in March 2002, Carol Lutsinger and Andy Miller (local BISD elementary and middle school teachers in Brownsville) talked to Jordan about testing SkyServer projects in their classrooms. Tim is now also using Skyserver as an SDSS educational tool in his astronomy and astrophysics courses at SKC.

Figure 16: SkyServer Web page.

11.3.2 EOT-PACI

The EOT-PACI program (Education, Outreach, and Training Partnership for Advanced Computing Infrastructure) has expressed interest on ways to link the E/O activities of GriPhyN/iVDGL with their program. At this stage, Roscoe Giles (EOT Team Leader), Scott Lathrop, (EOT Program Manager), Mary Bea Walker (Associate Director for EOT at NCSA), and Valerie Taylor (PI of Coalition to Diversify Computing), who is also a senior investigator in the GriPhyN project, have been already in contact with Campanelli and Romano to discuss some initial ideas for a possible collaboration - (a) Joint participation at major educational conferences and All-Hands meetings. (A list of conferences is available here). (b) Exploring opportunities for workshops on GriPhyN tools/resources with the Advanced Networking with Minority Serving Institutions (ANMSI). ANMSI is helping to engage faculty from MSIs with HPC research centers and activities; workshops on cluster technologies are underway. The contact people for the ANMSI project are: Allison Clark (PI for the project for EOT-PACI) and Stephenie McLean (project manager). (c) Linking our web sites to on-line tutorials.

11.3.3 ThinkQuest

Last year, Campanelli and Romano began talks with ThinkQuest to develop special challenge projects based on the application sciences and grid technology. One idea is that GriPhyN/iVDGL would provide resources in the form of interesting data sets (e.g., SDSS images or LIGO data) and/or "sandbox" CPUs that students could use when creating innovative web-based educational tools for ThinkQuest competitions that will be directly linked to our E/O Web site.

Harvey Newman and Eric Alakson at Caltech have expressed interest in linking GriPhyN and iVDGL E/O with ThinkQuest. They have a VRVS reflector set up especially for ThinkQuest at www.advanced.org; in Armonk near IBM Headquarters. Further contacts with ThinkQuest are needed to explore a possible connection.

11.3.4 QuarkNet and ATLAS/CMS E/O

QuarkNet[53] (see Figure below) is an ambitious (and highly successful) NSF-funded program that introduces large numbers of high school students across the US in particle physics research techniques. High school teachers are invited to an eight-week summer program for intensive study at host institutions where they work closely with particle physics faculty and staff. These teachers then take their experience (and enthusiasm) from this program back to their classrooms after the end of the summer. HU was one of the founding institutions of the national QuarkNet program. In addition, HU is one of the leading institutions in the US for training large numbers of African-American students in the sciences. There is also substantial outreach to K-12 teachers and students, exposing them to forefront research and education in the sciences. Indiana University and The University of Florida are also active QuarkNet centers. Therefore GriPhyN and iVDGL researchers at these institutions are encouraged to provide a grid-related component to the already existing QuarkNet activities.

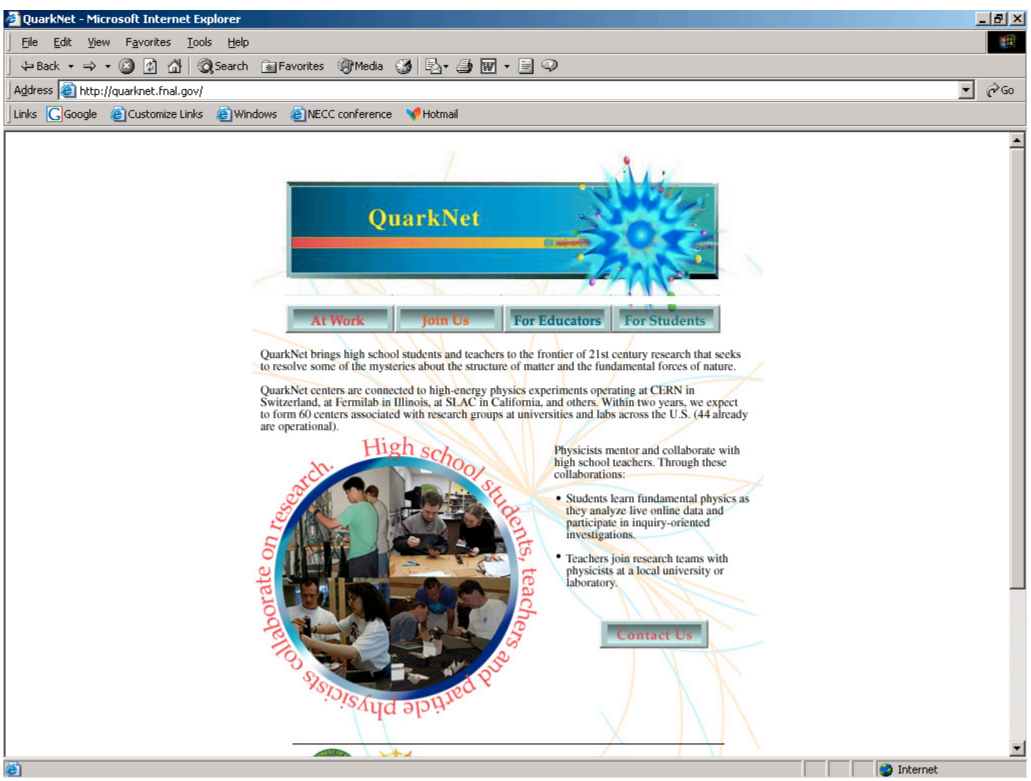

Figure 17: QuarkNet web page.

In particular, Keith Baker, who is the lead person for the HU, and a principal investigator for QuarkNet, will provide a direct link between the QuarkNet, ATLAS/CMS and GriPhyN/iVDGL E/O programs. Ken Cicere (a QuarkNet investigator) and Howard Brown (an undergraduate student) at HU will help Keith in this effort. The ATLAS group at HU consists of three faculty member, one postdoc, six technicians, and one engineering supervisor. The motivation for becoming a part of the iVDGL was to enhance the research capabilities within the ATLAS collaboration. The long-term goals are to:

- Get improved access to simulation tools for physics and detector simulations in the high energy physics community.
- Get improved access to ATLAS data once the experiment begins in 2007.
- Train Hampton University students and staff in the use of modern internet tools and methods.
- Contribute to the development of this revolution in internet use.
- Expose nonscientists, especially K-12 students and teachers to grid technology.

11.3.5 ATLAS Outreach

The Outreach Team has been encouraged to explore the following possible connection to ATLAS E/O activities

- Develop a summer program which explores ATLAS simulations on the grid. Longer term: develop simple web-based portal for simulation of physics events in ATLAS, showing an event display.
- Develop ATLAS – GriPhyN Outreach website which discusses contributions that GriPhyN and iVDGL are making for ATLAS and LHC grid computing in general. This could be connected to the ATLAS outreach website, and could provide a portal to all the major grid excitement, viewed from an ATLAS physicist's perspective.
- Build an iVDGL compute site - develop procedures for adding Tier 3 sites to the iVDGL. (Part of the mission of the new Physics Frontier Center at Hampton University is to develop a network of HBCU's that are meaningfully engaged in this kind of activity.) The contact person for ATLAS Outreach is Michael Barnett.

11.4 Current Progress

11.4.1 Hiring of Outreach Coordinator

In Fall 2001, The University of Texas at Brownsville (UTB) hired a full-time faculty member, Manuela Campanelli, to serve as a coordinator of the Education and Outreach (E/O) program of the GriPhyN project. To facilitate the coordination of the E/O activities among the various Tier3 centers, participating in iVDGL, Manuela Campanelli is leading the E/O program for both the GriPhyN and iVDGL projects, integrating the two into a single E/O effort.

Keith Baker (HU) is co-leading this effort. Other Members are: Tim Olson (SKC), and several undergraduate and graduate students who are receiving support from iVDGL (e.g. Jose Zamora and Sean Morris at UTB, Howard Brown at HU, etc).

11.4.2 Web site for E/O activities

As one of her first E/O projects, Manuela Campanelli created an active web site (GriPhyN Education and Outreach Center), linked from the main GriPhyN web page. At this stage, the E/O web site contains basic educational material about data grids and the participating physics experiments, and provides some technical support information (e.g., documentation, user manuals, how-to guides, etc.) for grid software (like Globus and Condor). During the following year, Manuela plans to expand this site to illustrate the concept of virtual data. The web site will also provide examples of scientific projects for students at various educational levels, which can be linked directly to the E/O activities of existing programs, e.g., Quarknet, ThinkQuest, EOT-PACI, which are described in later sections.

As one of her first projects, Manuela Campanelli developed a web site for E/O activities[54], linked from the main GriPhyN and iVDGL web pages. This web site currently contains basic educational material about data grids and the participating physics experiments, and provides some technical support information (e.g., documentation, user manuals, how-to guides, etc.) for grid software (like Globus, Condor, and the VDT). In the last 6 months, this site has been continually improved to provide more detailed information about the E/O activities in GriPhyN and iVDGL and other similar educational programs. It also contains a list of talks given by various GriPhyN and iVDGL members and course development activities that are on-going at various GriPhyN/iVDGL institutions. The E/O web site will soon be expanded to provide examples of scientific projects for students at various educational levels, which can be linked directly to the E/O activities of existing programs, such as SkyServer, EOT-PACI, QuarkNet, which are described in later sections.

11.4.3 Construction and Grid-enabling Tier3 Linux Clusters

During Fall 2001, UTB completed the construction of a 96-node Linux cluster. Although constructed primarily to analyze LIGO data, it is also being used as a testbed for GriPhyN/iVDGL software, thus introducing Hispanic minority students at UTB to distributed computing and grid-related technology. During Spring 2002, UTB undergraduate students Jose Zamora and Sean Morris (with the help of Scott Koranda at the University of Wisconsin, Milwaukee) learned how to install an initial version of the VDT on the cluster. UTB also hired another graduate student, Santiago Pena, who will be working jointly with Jose Zamora and Sean Morris, to test newer versions of the VDT. The aim is to be able to participate in joint LIGO-grid testbed together with other interested Tier 2 centers. UTB has also been in contact with Ahktar Mahmood (UT Pan American) and Kaushik De (UT Arlington), who are interested in joining iVDGL, in order to help them grid-enable their clusters.

During spring and summer 2002, all Team members will participate at the iVDGL facilities work team telecons to discuss and plan issues related to the construction of the Tier3 centers in preparation for next year. In addition, Tim Olson (SKC) has been attending several cluster construction workshops and plans to attend more meetings this upcoming fall. HU faculty and students are using a 50-node Intel-Linux computing facility at the US-ATLAS host laboratory for their simulation work. The cluster is connected to Brookhaven National Laboratory (BNL) via Internet2.

11.4.4 Meetings and Workshops

On March 1, 2002, Manuela Campanelli organized the first GriPhyN/iVDGL E/O Meeting at UTB. The meeting successfully spread the news about GriPhyN and iVDGL to UTB students and faculty, and to local high school and middle school teachers and students. There were a number of talks given by GriPhyN/iVDGL researchers (e.g., Paul Avery, Jordan Raddick, Scott Koranda), and an inauguration of the UTB cluster “Lobizon” attended by UTB administration and local newspaper reporters. The meeting also served as an opportunity to begin to coordinate the E/O activities among the various Tier3 Centers participating in GriPhyN/iVDGL.

11.4.5 Dissemination

In December 2001, Manuela Campanelli released an interview about GriPhyN, iVDGL, and grid computing in general that has been published in several magazines in Germany, among them an article in the Financial Times. In March 2002, she also released an interview[55] with the Brownsville Herald and wrote an article with Patrick Brady (University of Wisconsin, Milwaukee) in *Matters of Gravity*[56] to inform the relativity community about Grid computing.

11.5 Future Plans

11.5.1 Research Experience for Undergraduates (REU) Supplement

In order to give undergraduate students the opportunity to participate in grid-related research at several GriPhyN/iVDGL institutions, Campanelli plans to submit a proposal for an NSF Research Experience for Undergraduate (REU) supplement, at the end of Summer 2002. The idea is to request support for 10 to 20 undergraduate students doing grid-related research during the summer months. Students themselves will apply and choose the mentoring institutions on the basis of the research projects proposed by each institution. At the end of each year, the students would present posters or give talks at a conference specifically designed to showcase their work. The preparation of an REU proposal will require the input from GriPhyN/iVDGL institutions interested in mentoring students during the summer. Institutions interested in participating in this program should indicate (a) who is willing to mentor students and (b) what projects they think the students could work on.

So far, the following institutions have expressed interest:

Indiana University: James Williams and John Hicks at Indiana University propose to develop a simplified "view" of the grid via the iGOC. One idea for a summer project would be to have students develop a scaled-down, simple to install VDT that high school teachers and students could use to grid-enable a small cluster or single computer. The primary goal of a summer REU at the iGOC would be to collect, organize and publish (via the web) information to make the iVDGL (and the iGOC) more accessible by the Minority Serving Institutions (MSI). The rationale for this request is that the MSIs have fewer technical resources than the larger institutions. As a beneficial side effect should be simplification of access to the iVDGL for everyone.

University of Wisconsin, Milwaukee: Bruce Allen at University of Wisconsin, Milwaukee (UWM) has offered to mentor several undergraduates. His group already had one REU student working on their Beowulf cluster last year. He proposed some possible projects titles: (1) "Automated hardware monitoring on a large Linux cluster"; (2) "Automated software propagation on a large Linux cluster"; (3) "Automated account management on a large Linux cluster"; (4) "Networking performance testing on a large Linux cluster"; and (5) "Automated file replication using grid tools". However, detailed descriptions of these projects need to be provided.

Fermilab: Ruth Pordes at Fermilab has offered to mentor a REU student to work on the "Graphical Displays of Use and Installation Validation of VDT". A possible subproject is the "Display Status of GDMP Heartbeat", consisting of the following steps: (1) Define the set of hardware where VDT is to be installed. (2) Install VDT on those machines. (3) Run the GDMP heartbeat across the VDT installation. (4) Develop a web page to graphically display locations of GDMP in the Grid and show status of the heartbeat, e.g. Green - heartbeat working, Red - heartbeat not working. (5) Show graphs of history of heartbeat use and/or throughput through "clicking" on the status. (6) Provide web interface to configuration to start/stop GDMP heartbeat at certain sites (e.g. if it is interfering with production), change the frequency of the update etc. Examples of maps from the European projects can be seen is in the References[57].

11.5.2 Other Proposals

NISBAS Center: SKC administration has submitted a proposal for the creation of a NISBAS Center for Excellence in Teaching. The focus of the proposal is to develop a training center for K-12 teachers to improve science, math, and technology education in National Indian School Board Association Schools (NISBAS). If funded, SKC will be able to hold summer workshops and national educational meetings at SKC (starting next year), where GriPhyN/iVDGL researchers could participate by giving talks and tutorials.

COSM Physics Frontier Center: Hampton University was recently awarded a Physics Frontiers Center by the NSF: the Center for the Study of the Origin and Structure of Matter (COSM[58]). The research includes experimental particle and nuclear physics research at the energy frontier (ATLAS at the LHC) and at the precision frontier (Jefferson Lab). As COSM gets fully integrated into the research and development infrastructure at the University, more resources will be devoted to iVDGL activities. A principal goal of the Center is to build a network of HBCU's to carry out research at the cutting edge of nuclear and particle physics. Two other HBCU's are a part of the initial network of Universities in COSM, Norfolk State University in Norfolk, VA and North Carolina A&T State University in Greensboro, NC. Additional HBCU's will be added during the next year. COSM officially begins on July 1, 2002 at the participating institutions.

11.5.3 Meetings and Workshops

It was decided then that E/O meetings will be held at least once or twice a year, and will rotate locations between UTB, HU, SKC, and possibly other Tier3 centers. HU has expressed interest in hosting the next meeting (possibly in Fall 2002) while high school teachers participating QuarkNet activities are still on campus.

Manuela Campanelli proposed holding an E/O workshop in conjunction with one of the "All-hands GriPhyN/iVDGL Meetings." This would allow a more direct and larger participation of minority students at such a meeting, without needing additional travel money for students. UTB has volunteered to host such a meeting in 2003.

11.5.4 Dissemination

As one of the major goal for year 2002-2003, Campanelli and several GriPhyN, iVDGL and PPDG members will be working on the creation of a joint newsletter to inform the grid computing community about the project activities. The model we currently considering is the Gravitation APS newsletter *Matters of Gravity*[56]. An ideal goal would to launch the newsletter in November 2002 in occasion of the SC2002 conference.

11.5.5 Construction of Linux clusters

Tim Olson (SKC) will develop a detailed design of the SKC Linux cluster during summer 2002, with construction planned to begin in January 2003 with Year 2 funding. Two undergraduate students will be hired in January 2003 to aid in the construction and operation of the cluster. The SKC cluster will be used for LIGO data analysis, and for instruction in cluster and grid computing in the new B.S. in Information Technology degree program.

HU faculty and students are using a 50-node Intel-Linux computing facility at the US-ATLAS host laboratory, Brookhaven National Laboratory (BNL) for their simulation work. HU has also an Internet2 connection to BNL for the purpose of such studies. The Group has begun simulations using ATHENA/ATLFAST++ on the US-ATLAS cluster at BNL. These simulations have focused on signatures of extra dimensions in ATLAS thus far. Additionally, there are GEANT4 simulations ongoing using these same computing facilities at BNL. It is expected that this effort will provide a testbed for iVDGL software.

References

1 IVDGL home page, http://www.ivdgl.org/.

2 GriPhyN home page, http://www.griphyn.org/.

3 Virtual Data Toolkit home page, http://www.lsc-group.phys.uwm.edu/vdt/.

4 PPDG home page, http://www.ppdg.net/.

5 UK e-Science Programme home page, http://www.research-councils.ac.uk/escience/.

6 EPSRC home page, http://www.epsrc.ac.uk/.

7 UK Fellowship Call, http://www.research-councils.ac.uk/escience/documents/ivdglcall.pdf.

8 TeraGrid home page, http://www.teragrid.org/.

9 European DataGrid Project home page, http://www.eu-datagrid.org/.

10 DataTAG Project home page, http://www.datatag.org/.

11 High energy physics Intergrid Coordination Board (HICB) home page, http://www.hicb.org/.

12 GridPP Project home page, http://www.gridpp.ac.uk/.

13 LHC Computing Grid Project home page, http://lhcgrid.web.cern.ch/LHCgrid/.

14 Japanese Data Grid Projects, http://ninf.apgrid.org/welcome.shtml, http://www.apgrid.org/, http://datafarm.apgrid.org/.

15 GDMP home page, http://cmsdoc.cern.ch/cms/grid/.

16 Globus MDS page, http://www.globus.org/mds/.

17 Condor Hawkeye page, http://www.cs.wisc.edu/condor/hawkeye/.

18 FLAMES monitoring system web page, http://cil.cern.ch:8080/FLAMES/.

19 MOP home page, http://grid.fnal.gov/test_beds/MOP_install.htm.

20 DAGman home page, http://www.cs.wisc.edu/condor/dagman/.

21 IMPALA is described in the online document, http://computing.fnal.gov/cms/Monitor/docs/cms_documents/ProductionSoftware/CMS-PSOFT-002/cms-psoft-002.html.

22 SC2001 physics analysis demo, http://pcbunn.cithep.caltech.edu/sc2001/SC2001%20demo%20Bandwidth%20greedy%20physics%20analysis.htm.

23 ORCA home page, http://cmsdoc.cern.ch/orca/.

24 Clarens home page, http://clarens.sourceforge.net/.

25 Root home page, http://root.cern.ch/.

26 Magda home page, http://atlassw1.phy.bnl.gov/magda/info.

[27] ATLFAST home page, http://www.hep.ucl.ac.uk/atlas/atlfast/.

[28] GRAPPA home page, http://iuatlas.physics.indiana.edu/grappa/.

[29] Ganglia monitoring tool home page, http://sourceforge.net/projects/ganglia.

[30] Internet2 home page, http://www.internet2.org/.

[31] Led by Caltech on the US side and CERN on the European side.

[32] Internet2 HENP Working Group page, http://www.internet2.edu/henp/.

[33] Internet2 end-to-end Working Group page, http://e2epi.internet2.edu/index.shtml.

[34] TransPAC home page, http://www.transpac.org/.

[35] PPDG Certificate Authority work, http://www.ppdg.net/pa/ppdg-pa/cara/ppdg_cara_issues.htm.

[36] DOE Grids Certificate Authority home page, http://www.doegrids.org/.

[37] DOE Project Management Authority (PMA) page, http://www.doegrids.org/pages/doesgpma.htm.

[38] Monitoring web site, http://www-unix.mcs.anl.gov/~schopf/pg-monitoring/.

[39] Monitoring talk, http://www-unix.mcs.anl.gov/~schopf/Talks/internet2_pgmon_2_2002.ppt.

[40] Talk on Grid monitoring use cases, http://www-unix.mcs.anl.gov/~schopf/Talks/info_services_lcg_3_2002.ppt.

[41] Grid monitoring requirements document, http://www-unix.mcs.anl.gov/~schopf/pg-monitoring/ReqDoc/mon_req.v3.doc.

[42] PPDG/GriPhyN/iVDGL monitoring charter, http://www-unix.mcs.anl.gov/~schopf/pg-monitoring/charter.html.

[43] GLUE schema, http://www.hicb.org/glue/glue-schema/schema.htm.

[44] GriPhyN-iVDGL document server page, http://www.griphyn.org/documents/.

[45] EAC meeting, January 9, 2002, http://www.griphyn.org/meetings/2002Jan7/eac_agenda.html.

[46] HICB home page, http://www.hicb.org/.

[47] Grid Laboratory Uniform Environment (GLUE) home page, http://www.hicb.org/glue/glue.htm.

[48] Global Grid Forum home page, http://www.gridforum.org/.

[49] GLUE document, http://www.hicb.org/glue/GLUE-v0.1.doc.

[50] Main parts of diagram come from Cristina Vistoli of DataTAG.

[51] Facilities workshop http://www.acf.bnl.gov/UserInfo/Events/iVDGL-Workshop.shtml.

[52] LIGO code to compute gravitational radiation, Physical Review D, v 63, 042003.

[53] Quarknet home page, http://quarknet.fnal.gov/.

[54] Home page for GriPhyN/iVDGL Outreach activities, http://www.phys.utb.edu/griphyn.

[55] Interview with Brownsville Herald, http://arXiv.org/abs/gr-qc/0202044.

[56] Matters of Gravity, article in Spring 2002, http://arXiv.org/abs/gr-qc/0202044.

[57] Grid maps, http://ccwp7.in2p3.fr/mapcenter-bin/cmd?port=3001&cmd=hm&param=55.

[58] COSM Physics Frontier Center home page, http://www.hamptonu.edu/science/COSM/COSM.htm.