

A Project Vision

Our goal is to stimulate new discoveries by providing scientists with *effective and dependable access to an unprecedented national distributed computational facility*: the Open Science Grid (OSG) [OSG00]. We propose to achieve this through the work of the Open Science Grid Consortium: a unique hands-on multi-disciplinary collaboration of scientists, software developers and providers of computing resources. Together the stakeholders in this consortium sustain and use a shared distributed computing environment that transforms simulation and experimental science in the US. The OSG consortium is an *open* collaboration that actively engages new research communities. We operate an *open* facility that brings together a broad spectrum of compute storage and networking resources and interfaces to other cyberinfrastructures, including the US TeraGrid [TER00], the European Grids for EScience (EGEE) [EGE00], as well as campus and regional grids. We leverage middleware provided by computer science groups, facility IT support organizations, and computing programs of application communities. We also provide opportunities for individual contributions of merit.

In this proposal we present a five-year program of work to maintain and operate the OSG facility, to provide education and training opportunities in its use, and to expand its reach and capacity. We propose a program of joint projects spread across more than 15 institutions, many of which are collaborations with external partners. This program builds on the achievements since 2003 of Grid3 [OSG01] and the initial work of the OSG Consortium [OSG02].

We propose to build a cyberinfrastructure that can grow to provide thousands of users effective access to 100,000 CPUs, 10s of PB of storage, located at hundreds of sites and interconnected by multiple 10Gb/s network links. A unique feature of the OSG facility is support for the dynamic integration of new resources and applications and the harnessing of all available resources, thus extending the return on investments of our computing infrastructure and easing the inclusion of new communities. The OSG infrastructure relies on the resources and expertise of large scale computing facilities, builds on the foundational cyberinfrastructure developed and deployed by the GriPhyN [GRI00], iVDGL [IVD00] and PPDG [PPD00] projects, and depends on a broad range of computer science and IT development groups, university facilities and network providers.

The active engagement of computer scientists, information technology engineers, biologists, astrophysicists and researchers from other domains in the OSG consortium ensures that we continue to deploy, operate and evolve a generic cyberinfrastructure and thus can facilitate the needs of new scientific communities.

The requirements in scale of resources, users, capacity and performance of the OSG distributed facility are driven by our user communities, in particular the physics communities that are committed to the use of OSG to meet their massive computational, storage and networking needs. These are the ATLAS [ATL00] and CMS [CMS00] collaborations of the Large Hadron Collider (LHC) at CERN [LHC00], the LIGO [LIG00] Scientific Collaboration, the STAR [STA00] RHIC nuclear physics, and CDF [CDF00] and D0 [D000] Tevatron Run 2 experiments. Additionally, the usability, generality and multi-dimensional scale of the OSG infrastructure is driven by the active engagement with scientists from astrophysics Sloan Digital Sky Survey (SDSS) [AST00] and Dark Energy Survey (DES) [AST02], the multi-disciplinary GRASE [GRA00], as well as bioinformatics and genetics application communities such as GADU [GAD00].

Technical activities, which engage, train and include new researchers and organizations both in the US and overseas are integral parts of the OSG program of work. We support, and are expanding, the successful grid summer schools, the I2U2 eLabs [I2U00] initiative, student contributions to OSG projects have specific outreach projects with collaborators in Africa, South America [OUT00] and Asia.

Technical collaboration with TeraGrid, EGEE, and other regional, national and international cyberinfrastructure projects are a cornerstone of the OSG vision. The OSG facility is an integral element of the worldwide computing infrastructures of its stakeholders: the Worldwide LHC Grid (WLCG) [LHC01], the LIGO Data Grid [LIG01], and the SAMGrid [TEV00]. These collaborations promote interoperability and commonality which increases the overall effectiveness of the computing infrastructure.

Deploying and sustaining a *dependable and effective cyberinfrastructure* of the complexity, heterogeneity and scale of the Open Science Grid facility requires an experienced and well-organized team as well as an ongoing effort to extend the functionality and robustness of the software stack. Our experience has taught us not to underestimate the ingenuity and discipline that are needed to transform a loosely-coupled collection of autonomous sites into a production quality facility that is capable of serving a diverse group of Virtual Organizations (VOs). The focus of the VOs is their science. They expect the resources to be there when they need them and to have these resources allocated according to VO specific and site defined policies. Access to these nationally distributed resources should not require significant modifications to applications and should preserve the users' local computing environment. Domain scientists expect a transparent, 100% reliable, laptop level ease-of-use interface to serve as the gateway to the unbounded resources of the national infrastructure, while computer science and IT developers want to focus on developing and evaluating novel capabilities without interference. As we support an increasingly diverse community of resource providers and users it will take a significant amount of detailed support, deep technical understanding and operational commitment to maintain a transparent, easy-to-use distributed computing environment.

We will actively engage new sites and scientific communities and incorporate them into the live infrastructure. We will add to the capacity of the facility to meet the needs of large-scale, data-intensive scientific research and small-scale individual investigations. We will make available to a wide community new developments coming from the DOE SciDAC-2 and NSF computer science and application research programs.

We fully expect the OSG to become a home for an increasingly diverse set of science applications including: molecular biology, genetics, protein chemistry, nanotechnology, climate, geophysics of the earth and the hydrodynamics of the ocean, human reasoning, economics, natural language processing, behavioral psychology, geographic information science and more. Additionally, our engagement, education and training programs will ensure the inclusion of students, educators and next generation researchers not only in science, but also in IT.

B Science Benefits

OSG provides a common fabric for simulation and experimental scientists who run small (CPU days) or large (CPU centuries) scale scientific applications, with special utility for high throughput computing applications. These are large ensembles of loosely coupled parallel applications for which the overhead in placing the application and data on a remote resource is a fraction of the overall processing time, and for which the computations are sufficiently loosely coupled to be able to take advantage of opportunistic resources. The OSG facility is naturally matched to the needs of large-scale parameter sweep tasks such as simulations, BLAST [BLA00] searches, feature extraction from collections of images, and event processing that are common in the life-cycle of most experimental and observational sciences. Another common feature of this class of sciences is the need to interleave analysis of experimental data with large-scale simulations to determine acceptance of data collection instruments, develop new analysis techniques, and compare measurements with theoretical predictions. This requires support for turnaround times that meet the expectation of interactive users. The science reach of such experimental work is significantly impacted by the overall capacity and ease-of-use of the computing infrastructure.

The science communities contributing to and using the OSG facility see the possibility to extend the scope and reach of their science through the ability to scale up their computing while maintaining a constant allocation of human resources to their computing operations. For many of them, the effort involved in harnessing the power of large collections of resources limits the scope of their computing activities. Thus the users and stakeholders in the OSG Consortium are not only depending on the maintenance of the distributed facility and improvements in its capabilities but are also continuing their active contributions to OSG activities.

Physical Sciences: The physics and astrophysics collaborations have committed to the use of OSG as an integral part of their ongoing and expanding distributed computing systems.

LIGO: LIGO currently operates the LIGO Data Grid (LDG) [LIG01] and is hierarchically adapting this to the OSG infrastructure while maintaining a robust operating facility. The binary inspiral search analysis is the first to be adapted to the OSG infrastructure and further analyses will be adapted as the Virtual Data Toolkit (VDT) [VDT00] grows to support additional services for advanced workflow. Also, LIGO plans to make opportunistic use of compute cycles on OSG sites other than LIGO by reserving storage space for input data for epochs longer than the typical data pipeline run time. With an annual science run of data collected at roughly a terabyte of raw data per day, this will be critical to the goal of transparently carrying out LIGO data analysis on the opportunistic cycles available on other VO hardware. By 2009 LIGO will have completed one or more years of continuous coincident observation. Demands for compute cycles needed to exploit the full scientific value of the data will be very high with new ideas and analyses highly likely during this period. Identifying and monitoring of the opportunistic compute cycles and the advanced workflow management of jobs will be needed to provide low latency data processing across the OSG. Important services will be needed to identify available storage and compute cycles and publish this information intelligently to the LIGO workflow management environment.

The LHC: The US ATLAS and US CMS collaborations are fully depending on the OSG distributed facility to meet their data distribution and analysis needs. They are making their Tier-1 and Tier-2 resources accessible to the OSG and developing their data analysis systems to use both their own and opportunistically available resources. The US LHC software and computing programs continually use, test, validate and stress the OSG infrastructure to ensure needed performance and capability. The NSF-funded DISUN [DIS00] project activities are synergistic with the OSG mission. Extensions in the data, storage, security and VO service capacities of the OSG are essential for the LHC experiments, which together expect to accumulate up to 20PB of data during 2008, and to serve that data via 30PB of disk space to close to 100MSpecInt2000 worth of CPU power across 50-100 computing centers worldwide. The US LHC software and computing programs will thus continue their contributions to the OSG program of work. The factor of 7 increase in beam energy of the LHC over the Tevatron leads to a two orders of magnitude increase in production cross section for the top quark and similarly heavy particles. Coupled with a factor of 10 increase in instantaneous luminosity, the LHC experiments expect to have three orders of magnitude better sensitivity to a wide range of new phenomena, such as supersymmetric particles and TeV scale resonances decaying to leptons, within the first year of data taking [LHC02]. To fully exploit these scientific opportunities, functionality and scale of the computing required for the LHC in the US must be fully commissioned. Additionally, OSG must interoperate with the EGEE to be a significant contributor to the Worldwide LHC Computing Grid used by LHC physicists.

Tevatron Run II: D0 and CDF have existing and expanding distributed systems and software—SAMGrid and dCAF [TEV01]—that support their worldwide data processing and analysis needs. They are gradually making their resources accessible to the OSG Facility and adapting their software to integrate with the OSG and VDT software stacks. As the Tevatron increases its luminosity, to maximize the physics sensitivity for new phenomena as well as precision measurements on known particles like the top quark and W boson, the experiments need to scale their computing systems commensurately. As

analysis of the Tevatron data will continue for several years after the end of data taking, they will transition to the shared cyberinfrastructure in order to minimize operational costs.

Nuclear physics: For the past several years the STAR nuclear physics experiment has been heavily relying on Grid technologies to distribute its data between BNL (Tier-0) and LBNL (Tier-1). Using technologies now deployed as part of the VDT, this approach has enabled a doubling of the analysis throughput for several hundred nuclear physicists [STA01]. Driven by its rich long-term physics program [STA02], the STAR raw data rates will grow by an order of magnitude over the next two years. STAR will make its university based resources in the US and Brazil accessible to the OSG and will access its resources, as well as those owned by others, in an opportunistic and optimal manner for data intensive processing. Ramping from a few users to large scale daily use, STAR will reach the peak of its statistically data challenged program by 2009 when it will depend on the OSG to provide the software for virtualization of access to resources for its' distributed data analysis jobs.

The **Sloan Digital Sky Survey (SDSS)** makes opportunistic use of resources on the OSG for QSO fitting [AST01], near earth asteroid searches, and southern galactic hemisphere co-addition of images. SDSS needs the extensions in data and storage management to support access to datasets by analysis jobs. Simulations for the **Dark Energy Survey** will rely heavily on the OSG Facility and will make increasing use of available compute cycles. SDSS resources are being made accessible to OSG as part of the FermiGrid campus facility.

Multi-Disciplinary Sciences: The OSG currently supports a few multi-disciplinary communities that have made their local resources accessible to the OSG Facility and enabled their users to use the distributed infrastructure. The two most active communities to date are: Grid Resources for Advanced Science and Engineering and the Genome Databases and Update project.

Grid Resources for Advanced Science and Engineering (GRASE) has interfaced its regional grid resources in New York State to the OSG and provides a portal to enable researchers from a diverse range of disciplines access to both local and OSG resources. The applications include the Molecular Structure determination (*Shake-and-Bake*) and PHASES, a convenient, efficient pathway from diffraction data to protein mapping, computational chemistry, ecology, earthquake and geology applications, as well as protein sequence searching (SledgeHMMER).

The **Genome Databases and Update (GADU)** project based at Argonne National Laboratory uses local, OSG and TeraGrid resources for periodic searches through DNA and protein databases to compute and publish new and updated genomes using the BLAST, PFAM, and BLOCKS bioinformatics tools.

Computer Science Research: Computer Science groups, notably the Condor [CON00] and Globus [GLB00] projects, use the OSG to validate and evaluate novel distributed computing technologies. The value includes the availability of an at-scale facility which enables measurement of the effectiveness of advances in Computer Science. Students and researchers use resources on the OSG to analyze the use of and develop new distributed algorithms and methodologies.

Other collaborations are in progress, for example preparing the raw data of functional magnetic resonance imaging (fMRI) studies for analysis by the Dartmouth Brain Imaging Center [FMR00], and working with the Computational Chemistry Grid (CCG) through the Texas Advanced Computing Center.

B.1 Achievements To Date

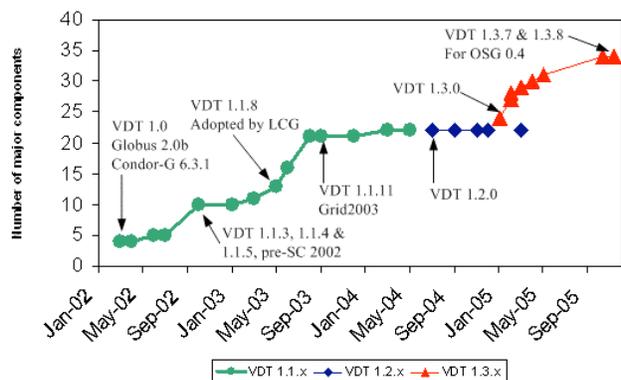
In the fall of 2003 the three projects that form the Trillium partnership, GriPhyN, iVDGL and PPDG, established Grid3: a national-scale cyberinfrastructure consisting of more than 3500 CPUs located at 30 universities and DOE laboratories. The software stack of Grid3 was



based on the VDT. Building and operating Grid3 has provided not only an immense amount of experience in maintaining a production cyberinfrastructure, but also has enabled substantial scientific contributions. These contributions have been described in the GriPhyN and iVDGL annual reports [GRI01, IVD01], PPDG news updates and quarterly reports [PPD01], the newsletter “Science Grid This Week” [SGT00], as well as many individual project reviews and papers. Crucially, the members of Trillium learned how to incorporate a facility mindset into management and operations, including delegating responsibilities, creating clear points of contact and utilizing a development infrastructure (Grid3Dev/OSG-ITB) [OSG03] to test and validate new VDT releases without disturbing production.

After almost 1.5 years of operation, Grid3 was replaced in October of 2005 by the Open Science Grid facility. Earlier in the year, the members of the Trillium partnership formed the Open Science Grid consortium as the governing body of the newly formed cyberinfrastructure. More than thirty University and Laboratory groups, as shown on the map, now make their resources and services accessible to the OSG. Condor, Globus, SDM [SDM00] and other computer science and application development groups provide and support the core middleware.

The methods and achievements from Grid3 and the three Trillium projects have not only directly benefited the participants, but influenced and informed other e-science efforts in the US and abroad. Our close and ongoing collaboration with European grid efforts have led to significant mutual benefits. The deployed software stacks of both OSG and EGEE are based on the VDT, which includes middleware contributions from all four efforts. The neighboring figure shows the evolution of the VDT since its inception in the winter of 2002 .



Our partnership with the Grid Integration Group (GIG) of the TeraGrid [TER01] started about one year ago with the GIG review. Since that time we have had increasingly beneficial contacts and exchanges. In the multi-grid interoperability initiative recently started within the context of GGF, members of the OSG consortium have been instrumental in the progress toward security/authentication, data movement and management, job management, and information schema.

C Organization

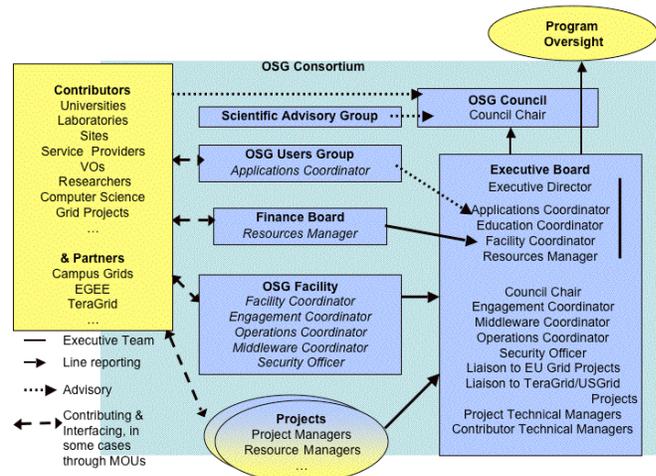
Our proposed program of work consists of three main thrusts: (1) the OSG Facility; (2) Education, Outreach and Training (EOT); and (3) Science Driven Extensions. The table summarizes the Full Time Equivalents (FTE) allocated to each of these thrusts and to the executive staff (comprising the executive director, resources manager and communicator). The management of all OSG activities is based on scientific engagement and oversight, as well as structured management and execution. The OSG Council—with the guidance of a Scientific Advisory Group—provides the scientific coordination. The council elects an Executive Director to manage programmatic activities with the help of an Executive Team. The Executive Director also appoints an Executive Board to direct the OSG program of work, draw up policies and represent the OSG Consortium in dealing

FTE	Facility					EOT D3	EXT	Staff	Total
	D2.1	D2.2	D2.3	D2.4	D2.5				
Boston			0.5						0.5
BNL		0.5							3.5
CalTech				0.5			3.0	0.5	4.0
Columbia						0.5			0.5
Cornell							1.0		1.0
FermiLab		1.0	1.5				2.0	2.0	6.5
Indiana U	3.0								3.0
LBNL	1.0	0.5	0.5						2.0
RENCI				1.5					1.5
UCSD							2.0		2.0
U at Buffalo	1.0		0.5	0.5					2.0
U of Chicago	1.0	0.5				2.0	1.0		4.5
U of Florida						0.5		0.5	1.0
U of Iowa		1.0							1.0
Wisconsin		1.0	4.0	0.5	1.0				6.5
Totals	6.0	4.5	7.0	3.0	1.0	3.0	12.0	3.0	39.5
			Total Facility: 21.5						

with other organizations and committees All appointments to the Executive Board are subject to council approval.

Ruth Pordes was recently elected for a two-year term as the Executive Director of the OSG. Her Executive Team consists of: Facility Coordinator (Miron Livny), Resources Co-Managers (Paul Avery and Albert Lazzarini), Applications Co-Coordinators (Torre Wenaus and Frank Würthwein), and Education and Training Coordinator (Mike Wilde). Together this team is responsible for all aspects of the program of work (deliverables, milestones and activities) and finances. The Executive Board includes the members of the Executive Team (cited above) as well as Security Officer (Don Petravick), Engagement Coordinator (Alan Blatecky), Operations Coordinator (Leigh Grundhoefer), Middleware Coordinator (Alain Roy), liaison to European grid projects (John Huth), liaison to TeraGrid and US grid projects (Mark Green), deputies to the Executive Director (Rob Gardner and Doug Olson), and the identified external project managers (currently Ian Foster, CDIGS/distributed systems technologies, and Harvey Newman, Ultralight/advanced networks).

The OSG Council is the governing body of the Consortium, and the Council Chair (Bill Kramer) is a member of the Executive Board. The Scientific Advisory Council includes leaders of scientific projects using the OSG and leaders in the area of distributed computing who advise the Council on the scientific direction and benefits of the Consortium's progress. The Resources Co-Managers are responsible for all financial matters and reporting advised by the Finance Board, which includes representatives from the contributing projects. The Users Group (initial members being Kent Blackburn (LIGO), Mark Green (GRASE), Jerome Lauret (STAR), Igor Sfiligoi (CDF), Sebastian Goasguen (NanoHUB)) works closely with the Applications and Engagement Coordinators to ensure the needs and schedule of the stakeholder organizations are reflected in the priorities and deliverables of the development and integration efforts. The Education and Engagement Coordinators work with domain specific advisory committees to help with application and training decisions and feedback.



C.1 Reporting and Reviews

Our program of work is project based. The Resources Managers oversee the agreements and invoicing of the deliverables of each project. The Executive Board meets every six weeks to review requirements and milestones. The Resources Managers provide financial and accounting reports to the Executive Board and Council at the time of these meetings. People supported by the project will submit monthly status and effort reports to the Resources Managers for review by the Finance Board.

We plan program reviews (to include external reviewers) every 18 months. The Scientific Advisory Group will review and report on the scientific benefits and value from the OSG program. The reviews will address usability, performance (including improvements) of and extensions to the OSG Facility, as well as the status and experiences in including new communities, individual researchers and resource providers. These reviews will also address the achievements of education, outreach and training, and the deliverables, schedule and effectiveness of the extension projects.

D Program of Work

Operating and maintaining the OSG Facility requires dedicated effort throughout its lifetime. The facility thrust of the program of work sustains a robust, usable infrastructure with increasing quality of service through operational and support activities; ongoing attention to security and troubleshooting; packaging, maintenance and testing of new software stacks as technologies evolve; and engagement with new communities of users.

In addition to the efforts needed to sustain a flightworthy cyberinfrastructure, the program of work attends to the needs of our users for end-to-end capabilities. Developing and deploying such capabilities requires close interactions with the domain-specific applications as well as advances in the functionality of the OSG software stack. The science driven extensions thrust of the program is structured as a set of well-defined projects that include external partners. When ready for production, the software tools developed by these projects are integrated to the VDT, tested by the integration team and deployed within the OSG Facility. Thus the program of work includes:

- *Management, operation and evolution of the distributed facility* including software integration configuration and deployment and well defined and documented middleware releases. Facility activities include system wide performance and availability monitoring and analysis, comprehensive functional testing, and in depth engagement in support of users and system administration.
- *Engagement and training programs* to actively help new entrants make their computing and storage accessible to the common infrastructure, to help new researchers use the common environment, and provide hands-on workshops and published materials for training and dissemination.
- *Provision and maintenance of an at-scale integration and validation testbed* providing a heterogeneous platform for vertical and horizontal system testing of new releases, new technologies, new capabilities, and new applications, to facilitate smooth transition into the production environment.
- *Development and integration of extensions* to develop job and workflow management, security and policy services, integration of advanced network fabric capabilities needed to meet the scientific requirements and schedules of the stakeholders.
- *Interoperation* of the OSG infrastructure with other distributed environments from local campus infrastructures to the national and international grids that are forming the transparent worldwide cyberinfrastructure.

Moreover, the success and progress of the Facility also depends on continuous growth in and contributions to compute, storage and networking capacity as well as continued enhancement in the cyberinfrastructure capabilities in the US, including:

- A comprehensive program for *Grid security research and development* to ensure the integrity and defense of the open infrastructure on which our science depends.
- External funding of *local facilities*, for the purchase and support of the compute, storage and network hardware of the facilities made accessible by the OSG infrastructure.
- An aggressive program of *network and middleware research* to ensure the necessary increases in scale and performance of the infrastructure.
- A sustained program of *middleware software development, packaging and support* to ensure the continued availability, and further development, of the core software on which OSG builds.
- The development of *applications* that can be adapted to run on a distributed, loosely coupled infrastructure.
- A sufficient *cyber security infrastructure* supporting single sign-on to access grid resources, engagement of existing computing facility cyber security staff, risk based mitigation measures.
- *Commitments from the stakeholders* to enable policy driven sharing of resources and when possible the use of common services.

D.1 The Facility

The proposed program of work presents a five-year effort to maintain and operate an effective distributed facility while expanding its capacity and reach. The OSG Facility is structured to facilitate inclusion of new sites, expansion in software capabilities, increases in the number and diversity of the communities it serves, improvement in availability, reliability, security, and performance, and reduction in per-resource operations costs. To meet the challenges of sustaining a national facility of excellent and improving quality requires a significant effort. This effort must be well-managed and dedicated to the mission. The OSG Facility is structured as a coherent team of 21.5 FTEs who are organized in four activities: *operations and control, security and troubleshooting, software release and support, and engagement with new user communities*. A dedicated activity coordinator manages each activity. Together with the Facility Coordinator, these four coordinators form the Facility Coordination Team and are responsible for the facility portion of the OSG program of work.

The OSG Facility operates and evolves to meet the requirements, deliverables and milestones of the physics communities that are engaged in the OSG consortium and are committed to the use of OSG to meet their compute, storage and networking needs, in particular the US-LHC collaborations, the LIGO [LIG00] Scientific Collaboration, STAR [STA00] and the Tevatron Run 2 experiments. At the same time the OSG Facility operates and evolves as a generally usable and effective infrastructure for all stakeholders and partners.

An important aspect of the facility program of work is strong interaction with other US cyberinfrastructure efforts. Under the direction of a dedicated liaison, we will continue our partnership with the TeraGrid-GIG project to ensure interoperability between the two infrastructures and to ensure the leverage of common software capabilities, services and support infrastructure.

D.1.1 Facility Operations

The Facility operations activity is responsible for the daily operation and transition of the facility to new releases of the OSG software stack. It follows a distributed support model that includes site administrators, user support organizations and technology providers. It can thus accommodate the expected scaling without a significant increase in effort. This distributed structure enables leveraging of operation services provided by other e-science efforts – in particular, NMI Grids Center [NMI00] and EGEE.

The activities involved in providing Operations services to the Facility include: maintaining and publishing system wide monitoring analysis and diagnostic information as well as working with site administrators to ensure effective and appropriate use of the resources; deployment, triaging and tracking the resolution of users problems and questions; and ongoing testing and publication of the function, availability and performance of resources and services. Operations include: maintenance of registration and agreement repositories; support for an integration infrastructure; provisioning, configuration and documentation of new releases of the OSG common software stack; definition and documentation of operational procedures and practices; and operation of the central catalogs and discovery services. Ongoing work will introduce redundancy and fault tolerance into facility wide services, publish and analyze accounting information and collaborate with operations organizations of peer infrastructures, including especially EGEE and TeraGrid.

Functional testing of the sites and the services of OSG is an integral part of the Facility's daily routine. These tests are also instrumental during the integration phase of a new software stack when new software and extended capabilities are tested. Following the integration phase the new release is documented and configured into a production distribution.

The Operations Coordinator provides regular reports of metrics of use and activities of the Facility. These metrics are accessible via web interfaces [ACD00, MON00] and include utilization of the accessible resources, percentage success in functional tests, number of registered resources and VOs, number of

support tickets and their time to resolve, etc. We will deploy and maintain policy tools and procedures to improve the allocation and efficiency of resource usage.

Operation services are provided by 6 FTEs, including management by a dedicated operations coordinator. These FTEs are distributed as follows: 3 for operations, 0.5 FTE for the dashboard, 0.5 FTE for functional testing, 1 FTE for software integration, and 1 FTE support for the integration infrastructure.

D.1.2 Security and Troubleshooting

By its very nature, the Facility presents many security challenges. The proposed program of work includes a security team that is devoted to the protection and security of the Facility. The Security Officer and security team provide dedicated effort for a continuing program of operational and support activities. They work with the operations teams, resource, site and VO administrators, users and software providers. They perform activities to provide and use log files, inventories and catalogs; they ensure timely and appropriate notification of events and vulnerabilities; and they deploy and operate auditing and analysis tools in support of security and protection.

We are developing a *Security Management Plan*, which includes procedures to protect the assets of the OSG, ensure timely response to incidents, and provide for the inventory of services. The plan calls for periodic probes and audits to measure the effectiveness of our infrastructure and procedures. We will evaluate, integrate, and use new tools and technologies as they become ready for deployment and operation. All OSG users agree to an Appropriate Use Policy and all resource and service providers register a Service Agreement. We will administer the OSG Registration Authority (to replace the existing iVDGL and PPDG RAs) and OSG VO.

Closely related to security is the end-to-end troubleshooting and resolution of any fault or problem that occurs in the Facility. The number of components, services and interfaces used by any application and the non-deterministic, asynchronous, parallel nature of distributed, multi-step operations over a global computing infrastructure make analysis and resolution of problems a challenging, technically deep and time-consuming activity. We include effort in the Facility directed to such activities. In some cases, this team is augmented by effort temporarily assigned from other activities in order to guarantee timely resolution of faults or problems in the infrastructure.

4.5 FTEs are allocated to this activity, including management by the OSG security officer. 2 FTEs are assigned to the security team and 2.5 FTE are assigned to the troubleshooting team.

D.1.3 Software Release and Support

The OSG Facility includes a released and well-defined software stack. The base software stack supports the minimal set of functionality needed for resources to participate in and users to use the OSG. In addition the software stack includes common services used by several, or all, Virtual Organizations.

The OSG software stack is based on the Virtual Data Toolkit (VDT) with minimal configuration and OSG specific additions. The VDT relies on the Condor and Globus middleware releases, and the NMI release, build and test infrastructure [NMI00]. Support for the maintenance and continued development of the software, including Condor and Globus, is outside the scope of the OSG program of work. VDT is released for many hardware platforms and OS versions and provides the foundation for OSG heterogeneity in processor and storage. The current VDT will be expanded to include existing groups with additional expertise in storage and provide increased support for storage and data management components. We depend on the Pacman [PAC00] tool for the packaging of versioned sets of software using a distributed set of software caches.

Software provisioning includes VDT packaging, validation, component testing, distribution and installation, and OSG software configuration, releases, and functional validation software. It encompasses multi-version software management and support, improvements in the functional testing of components and integrated software sets, improvements in ease and robustness of installation and configurations,

deployment, integration and configuration of additional storage and data management components and services, especially in support of the wide-area and local high-throughput data I/O needs of the data intensive science applications. The scope of the work includes expansion in the number of platforms and OS releases needed by the user community, and integration and deployment of OSG extended capabilities and new technologies. We will improve the software provisioning efficiency and response by making support for incremental updates of the VDT and for quick releases to provide immediate response to identified vulnerabilities or critical patches.

We will continue to work with the EGEE and WLCG in support of a common software stack with the VDT being the vehicle for bi-directional adoption of common components. In particular, we will continue to work on the robustness and compatibility of the EGEE VO management and workload management software. We will also continue to work with the TeraGrid-GIG on common interfaces and interoperable software components. We are exploring ways the support needs of our communities can be leveraged by complementary efforts and are collaborating to help ensure a seamless environment for user applications.

7 FTEs are allocated to this activity, including management by the OSG Software Coordinator. 1 FTE is allocated to management and liaison work, 0.5 FTE is allocated to the software distribution framework, 1.5 FTE to collaborate with the EGEE and TeraGrid, 1.5 FTE to developing functional and regression tests, and 2.5 FTEs assigned to software packaging and release support.

D.1.4 Engagement of New Communities

Expanding our capacity and reach by including new resource and user communities is a core mission of the OSG. Engaging scientific domains outside of physics and the few currently collaborating broader groups to use OSG software for research and applications is a difficult undertaking for several reasons:

- Physicists have a several year lead over other disciplines because the OSG infrastructure and software has largely been developed through direct collaboration with and adoption by physics applications.
- Other disciplines have traditionally been able to afford little effort at common community software or cooperative approaches, which can take advantage of Grid technologies and portals. This is now changing as both NSF and DOE are encouraging biologists, chemists, environmental scientists and so forth to more actively explore the use of cyberinfrastructure.
- Lastly, there are few incentives for domain scientists to explore new technologies and approaches, as they cannot afford to take a lot of time to learn how to use and master these technologies. New technologies such as portals must not only be easy-to-use and intuitive, they must provide some competitive advantage to the user being more productive or being able to address problems and research areas that are not being adequately addressed today.

The engagement team provides focused effort, with responsibility and authority, to engage with each new discipline in series. The engagement team ensures the needed technical support to achieve the interface changes in the Facility needed to enable these new disciplines to use the infrastructure effectively and easily. This team is located inside the Facility organization to ensure continued and close cooperation and timely response and attention of the management, operations, software, and security personnel. The Education and Engagement Coordinators have a special relationship and are ex-officio members of each other's teams.

The activities of the engagement team will include identification of one or two scientific domains with a disciplinary champion for the engagement with OSG and a collaborative work agreement with a few key domain researchers to serve as early users and avatars for the new community. There will be agreed upon metrics to manage expectations and deliverables. We will liaise between the OSG applications and extension development projects and the new community and participate in the User Group. We will develop methods and processes for integration and/or interoperation of new users and approaches with the existing teams and activities as well as hold up to week-long working visits on-site with researchers in the new domain.

We will work through the Facility to expand the number of sites, users and VOs using the OSG, specifically working with resource providers and applications that are grid ready, and which can benefit from the OSG without full retooling. We will continue and increase collaboration and synergies with TeraGrid-GIG and EGEE, which have equivalent goals in providing transparent access to distributed resources across administrative domains. We will interface and support bi-directional access to their resources as well as cooperate with campus and regional grid organizations, initially including FermiGrid [FER00], GLOW [GLO00], GRASE, TACC [TAC00] and the Harvard Crimson Grid [HAR00].

3 FTEs are allocated to this activity, including management by the OSG Engagement Coordinator. 0.5 FTE for management/liaison role, 1 FTE as a domain software specialist, 0.5 FTE for software development, and 1 FTE for documentation.

D.2 EOT—Education, Outreach and Training

Our EOT program, described below, seeks to enable the effective use and growth of OSG through training and to use OSG to provide computational science education for high school and college students, with an emphasis on under-resourced communities both domestically and abroad.

The effort assigned to EOT is 3 FTEs, including management by a dedicated Education Coordinator. 1 FTE for training, technical and documentation support, 0.5 FTE for outreach to Africa and 0.5 FTE for outreach to under-resourced communities, as well as funding for students.

D.2.1 Training for OSG Users and Administrators

We will develop, present and maintain courses on the effective use of the OSG for individual users, site and VO administrators, support center personnel and new application communities. We will enhance the training material and hands-on laboratory used in our 2004-05 Summer Grid Workshops [IVD02], which covers the breadth of OSG usage, from obtaining certificates through running scientific workflows and creating new services. We will integrate this material with the evolving OSG component and procedure documentation.

In addition to a face-to-face workshop format, the courseware will be adapted for remote seminar delivery through technologies ranging from the Access Grid [ACC00] to basic teleconferences with remote desktop visuals. A self-paced Web format will also be created from the base material. Lab facilities that provide round-the-clock, hands-on training will be developed to support this wide range of delivery modes.

We will develop targeted training material and workshops to help site and VO administrators as well as support personnel. To help new user communities leverage the OSG, we will develop a *Grid Cookbook* in collaboration with the Southern Universities Research Association (SURA) [SUR00]. This will introduce Grid technologies in short, easy to digest chapters that describe basic Grid concepts, the major middleware technologies and standards, and how to apply them in science, engineering and biomedical activities.

D.2.2 Educational Outreach to Students

OSG EOT will promote the participation of students in science and technology by creating compelling computational science programs which publicize opportunities and teach techniques for student use of the OSG. One activity in this area is to increase the frequency, scope, and reach of the Grid student workshop. Additionally, we will sponsor students to work with the OSG Facility and application teams to provide hands-on experience and knowledge to the next generation of researchers and technologists. We will create student virtual organizations within OSG to provide students with computing resources for these projects. OSG EOT staff will provide technical support and services for education projects such as QuarkNet-Grid [I2U01] and the I2U2 eLabs initiative, which utilize the OSG for hands-on data analysis in both classroom and informal (museum) settings.

We will also work with CS and science faculties to make the OSG and its training materials available in emerging CS and computational science curricula. Such classes are being taught at a growing number of universities. We will hold workshops at appropriate conferences and meetings to pursue and promote this goal.

D.2.3 Serving Under-represented Communities

We will prepare and deliver training workshops tailored to meet the needs of under-represented students and minority serving institutions. We will seek advice from experts and participants on, and experiment with, various approaches for specific audiences. We will assemble training materials from existing sources to teach prerequisite UNIX, systems and networking skills, and provide training materials and OSG access to the Minority-Serving Institutions Cyberinfrastructure Institute, in collaboration with SURA.

The initial focus of the student research program will be Hispanic students in the Southwest and Florida, where programs are underway [CHE00]. We will use our collaboration with UTexas Brownsville, the largest Hispanic-serving school in the US, located in a technologically under-resourced region, as a pattern for meeting the needs of similar schools around the country.

D.2.4 International Outreach

OSG will serve a vital US role in eliminating global digital divides. We will support the South African HENP outreach program of the ATLAS group at Columbia [COL00], helping the University of the Witwatersrand in Johannesburg to create an OSG site and utilize the OSG for research. This will promote both African network connectivity and participation in the international science community. We will also work aggressively to create and extend OSG facilities at universities in Brazil [BRA00] as well as established facilities in Taiwan and Korea. Former students of our summer workshop from Argentina, working closely with LSU, represent further potential inroads in the Americas.

In all of the OSG EOT efforts described above we will strive to maximize effectiveness through both collaboration and management by metrics. We will collaborate with and leverage the work of other projects, and make the results of OSG work visible and freely available to these programs through multiple channels of dissemination. OSG EOT will work closely with the Engagement team to bring new communities into OSG, and with the TeraGrid EOT effort to identify common needs. The OSG and TeraGrid EOT coordinators (M. Wilde and S. Lathrop) are co-located at Argonne, which will help facilitate this cooperation. In all of our EOT efforts, we will measure and evaluate effectiveness and use participant feedback to tailor and improve future deliveries and set development priorities for courseware.

D.3 Science Driven Extensions

We propose to specify, develop, integrate and test new functionalities and include them into the OSG Facility through well-defined projects in collaboration with external partners. The evolving needs and schedule of the science drivers, the OSG Facility, and other OSG stakeholders drive these projects. The Applications Coordinators, together with the Users Group define, prioritize, deploy and test, while the primary responsibility for development remains with the external partners. The Applications Coordinators allocate effort (in consultation with the Executive Board) from the twelve-member Applications Team, and manage these projects.

Deliverables of the extension projects are specified such as to be broadly applicable, and thus generally useful and beneficial. Deliverables from external partners, as well as deployment and testing schedules will be agreed to through MOUs and included in the OSG tracking and reporting mechanisms. This provides benefits to both parties: the OSG stakeholders can expect more dependable schedules, as effort in the Applications Team can be reallocated to minimize schedule risk, and the external partners can gain the benefit from an experienced team committed to work through testing and deployment problems as they arise. We expect these partnerships to considerably shorten the overall time between requirements

specifications and deployed products, and to lead to overall improved software artifacts, which in turn will lead to scientific discoveries. The Applications Team will also help adapt existing systems and applications to use the OSG Facility. Short-term projects may be defined that include allocation of effort to specific organizations to achieve the needed adaptations.

While the Application Team members are generally assigned to projects of defined duration, individual ongoing effort is assigned to the Globus project for integration and testing of new software (including the Virtual Data System/Pegasus), for collaboration on Security with external partners in the US and Europe, and within the Ultralight project for integration and testing of advanced network capabilities. We are working closely with the NSF DISUN project as well as the software & computing projects of the main stakeholders. We expect these to provide the primary infrastructure for the testing and integration efforts described here.

The following areas of work form the initial set of extensions needed for the OSG. To meet the deliverables we identify the initial list of external projects we plan to partner with—including those that depend on the success of proposals submitted to DOE SciDAC-2 and NSF PIF programs.

- **Data Storage Access and Management:** Specific extensions include support for space reservation and lifecycle management, role-based authorization, access, quota management and accounting and support for Petabyte sized storage resources and data management needs. We must support access to storage installations ranging from single file systems on small clusters to fully distributed systems of Petabytes of disk. These may be integrated with tape archives tens of Petabytes in size, and the data must be accessible to local and remote multiple TeraFlop computing clusters. While we are committed to the combination of SRM on top of gsiftp as the common interface for managed data transfer, satisfying the full dynamic range requires the use of more than one implementation. Success in this area depends on work on the SRM interface as well as the underlying storage element.

Partners: We will have joint projects with some or all of Condor, Globus/CDIGS, the Distributed Science, Scientific Data Management (SDM) and Storage Resource Management Centers for Enabling Technology (CET)s, and dCache and Objects On-Demand Scientific Application Partnerships (SAPs). We expect continued contributions from the SDM, STAR and US LHC S&C programs, BNL, Fermilab and LBNL, and collaboration with the EGEE/WLCG.

- **VO Specific Services Management and Use:** The capabilities needed by the complex and/or large-scale science applications of OSG scientists drive the distributed facility architecture to support community managed specific services deployed and dependent on the common infrastructure. These include such services as database caches, information and data catalogs, application pipelines and monitoring services. The requirements include lifetime management for persistent and short-term services, remote configuration and usability in diverse network and security environments of remote sites, and common auditing, accounting and diagnostic administrative infrastructures.

Partners: We plan joint projects with the Workspace/Edge Services activities in Condor, Globus/CDIGS, and the Distributed Science CET.

- **Advanced Workload and Workflow Management, Planning and Execution:** Requirements to support the effective and managed use of a common, shared distributed facility of tens to a hundred independent resource owners include a secure, robust, workload management system that supports late binding and opportunistic job scheduling. The scope includes intra-VO and inter-VO as allocation of resources as well as a dynamic policy definition and enforcement tools and management and advanced workflow administration and execution infrastructures.

Partners: We will have joint projects with Condor, Globus/CDIGS, and Distributed Science CET, including extending the Virtual Data System [VDS00] and the EGEE Workload Management System [EGE01].

- **Security Enhancements including Authentication, Authorization, Sandboxing, Auditing and Accounting:** Deploying and extending a secure and auditable, as well as performant and robust distributed facility will continue in an increasingly hostile cyberinfrastructure environment. We plan ongoing activities to: improve and advance our security infrastructure and procedures, provide tools to analyze and address vulnerabilities, and ensure an open defensible environment for science applications.

Partners: We have joint projects with Condor, Globus/CDIGS, Security for Open Science CET, EGEE gLITE and security middleware activity, and TeraGrid-GIG.

- **Advanced Networks Integration and Interfaces:** In the next three to five years, the needed scale and delivery of data access by the OSG scientific community may necessitate the integration and use of advanced network fabric management into the distributed facility. This will pose significant technical as well as intellectual challenges as we develop an understanding of the interfaces between CPU schedulers, storage systems, and networking.

Partners: We are planning joint projects with Condor, Globus/CDIGS, LambdaStation, UltraLight[ULT00], TeraGrid-GIG, Internet-2 [INT00] and PLaNetS PIF proposal.

E Milestones

We organize the OSG program of work into four Phases – two of eighteen months and two of one year. For the first two phases we define science goals and present deliverables in terms of resources, services, and effectiveness. These deliverables are based on the assumption that by 2008 up to 50% of all the ATLAS and CMS computing worldwide will be executed on the Open Science Grid. This implies that up to 15PB of disk space and 50MSpecint2000 of CPUs must be accessible by the OSG infrastructure. For these storage and compute resources to be effectively harnessed, 10Gbit/s networks will be used to move 10TBs within hours between sites. US-ATLAS and US-CMS have by far the most demanding workloads among all the OSG stakeholders. We expect that the requirements of all the other communities will not exceed the requirements of the two LHC experiments. Given these expected requirements, the OSG Facility must be capable of managing at least 100,000 jobs per day.

The focus of the Facility thrust is the routine operation of the infrastructure and the incremental evolution of the OSG software stack. The overall effectiveness of the OSG Facility depends on its ability to minimize the impact of intermediate failures or limitations in the functionality of the software. Throughput suffers when resources that can be matched to pending tasks remain idle or when work has to be redone. When a system fails to successfully complete the execution of a job due to reasons beyond the control of the user, turnaround suffers as jobs have to be resubmitted. Improving the effectiveness of a large distributed infrastructure like the OSG Facility is an endless process. Our milestones present a sustained and measurable effort to maximize the throughput and minimize the turnaround time of the OSG Facility.

The milestones for the Extensions clearly depend on the success and timeliness of the deliverables from the partner development projects. The specific expectations and schedules will be agreed to at the start of each technical activity.

E.1 Phase I: Months 1-18

Overarching goals are to establish the operational procedures of the OSG Facility, improve the dependability of the infrastructure, deploy common storage management software, meet the milestones of the physics collaborations with critical path dependencies on OSG and demonstrate the usability of the infrastructure by additional campus grids and domain sciences.

Science Goals:

- LIGO: Use OSG for binary inspiral analysis. Expand the user community and types of applications.
- LHC: Support for US LHC service challenges in preparation for the start of data taking.
- STAR: Migration of all (most) simulation to an OSG based operation, use of opportunistic resources with a combined software packaging and deployment and on-the-fly SRM deployment.
- CDF: Full use of OSG for software release 7 series Monte Carlo simulation.
- D0: Full use of OSG sites for reprocessing and Monte Carlo. Dynamic deployment of SAM services in the common Edge Services Framework.
- SDSS: For the QSO Fitting Project, fit all spectra beyond data release 5. For Near Earth Asteroids, discover NEO's in all cumulative SDSS data, SDSS Coadd and DES simulation applications.

Facility:

- During the first six months we will develop the following plans: Establish procedures and infrastructure in support of agreements for the use of OSG.
- Establish OSG RA with one business day response to requests.
- Establish test procedure and periodic practice for cyber security incident response.
- Accommodate additional science domain and flow of jobs from GLOW campus grid to OSG.
- OSG-EGEE: Working operations and security incident response interfaces between the infrastructures. User level virtualization of OSG and EGEE for LHC applications.
- OSG-TeraGrid: Example applications running across OSG and TeraGrid. Coordinated security incident response. Simple coordinated operations. Interoperation of accounting and information systems.

Education, Outreach and Training:

- Review and revise courseware and create course calendars.
- Create template and cost structure for in-person delivery.
- Enhance questionnaire and create simple analyses and reporting tools.
- Setup of I2U2 VO and establishment of scheduling priorities for this VO.
- Create template for student-support funding applications.
- South African Grid site planned and added to ITB .
- Include OSG resources seamlessly in Brazil and Argentina.

Extensions:

- Support for role-based authorization and access control for storage to be included in the VDT. Make sure that at least ATLAS, CMS, LIGO, plus one of the other application communities can incorporate this functionality into their software stack.
- Add to VDT a storage element supporting SRM v2. Make sure that all current stakeholders, plus one of the other application communities can incorporate this functionality into their day-to-day operations on OSG.
- Specify, develop, and deliver an initial limited functionality implementation of a VO specific services management infrastructure ready for widespread deployment in the Facility. This first version supports dynamic deployment and remote configuration of services but not yet complete lifetime management. Make sure that at least current stakeholders, plus one other application community can incorporate this functionality into their day-to-day operations on OSG.
- Specify, develop, and deliver a first generation audit system ready for widespread deployment in the Facility.

- Deliver an implementation of a workload management system that supports late binding and opportunistic job scheduling. Make sure that at least ATLAS, CMS, plus one of the other application communities can incorporate this system into their day-to-day operations on the OSG.
- Deliver an enhanced version of VDS which is fully SRM aware, and capable of benefiting from the deployed storage infrastructure on the OSG. Make sure that at least LIGO, SDSS, DES, and GADU can incorporate this system into their day-to-day operations on the OSG.
- Deliver specifications and a first large-scale prototype system that allows integration of network management tools into the distributed facility.

E.2 Phase II: Months 19-36:

Overarching goals are to complete the extensions needed by the initial stakeholders, increase the capabilities and usability for new communities and establish operating procedures for effective use of the infrastructure.

Science Goals:

- LIGO: Expand the user community and types of applications such that the entire LIGO/LSC (LIGO Scientific Collaboration) user base derives scientific benefit from transparent operations across LDG and OSG. Stochastic and some aspects of the burst data analysis pipelines will be extended run on OSG.
- LHC: Support for low and then high luminosity LHC Physics analysis. The LHC data volume explodes from less than 1PB to more than 20PB during this phase! User activity changes dramatically both in number and intensity of use. Reliability and robustness is likely to be the overriding concern for the LHC application communities.
- STAR: Support for user batch analysis on the distributed facility, object based and interactive analysis for the STAR collaboration. Scaling of the order of 10k jobs/day and beyond as well as a robust infrastructure will need to be reached by the second 1/3rd of Phase II.
- CDF: Analysis CAF infrastructure and data analysis applications on the OSG-CAF.
- D0: Initial support for user analysis on some OSG sites. Continued and expanded use of OSG for all stages of the overall analysis chain.
- SDSS: Continued use with scaled resource needs.

Facility:

- Reduce the “in-effectiveness” metrics of the Facility by 50%.
- Deploy distributed logging infrastructure and establish periodic analyses of cyber security audit logs.
- Support transparent data movement between OSG & EGEE.
- End to end monitoring and problem determination across the Grid infrastructures.
- Support transparent movement of data and applications across OSG & TeraGrid.

Education, Outreach and Training:

- Support of I2U2 VO(s).
- Add new modules as needed e.g. SRM.
- Create tutorial material to provide necessary prerequisites in LINUX and networking.
- Setup of student VO for independent research projects.
- Modularize courseware for self-paced delivery. Test and extend existing material.
- South African Grid site launched for scientific analysis.

Extensions:

- Expand the use of new capabilities introduced in the previous phase towards additional application communities.

- Expect to spend significant effort on understanding robustness, reliability, efficiency, and ease of operations issues for the new capabilities as deployed on the production grid as a result of work in the previous phase. Work with computer science partners on new releases that address these issues.
- Deliver a first production system that incorporates new network management tools into the distributed facility.
- Deploy a high-level user language to consistently express user interactive analysis and batch-based workflow.
- Expect to work on specifying requirements for new capabilities with new communities that started initial operations on the OSG during the previous phase.
- Deliver the final auditing system to be used by OSG for the remainder of this 5-year funding period.

E.3 Phase III and IV: Months 37-48 and 49-60:

The needed extensions in capability for the initial primary science drivers will have been delivered. There will be a focus on operating and maintaining a robust infrastructure—keeping it flightworthy. We expect the new communities to take on the primary push for new capabilities and drive new extensions.

Education, Outreach and Training

- Create virtual laboratory for easy, safe, concurrent delivery to multiple groups and support of self-paced individual learning.
- Create model for remote virtual delivery and develop client-side toolkit for remote delivery.
- Computational Science curriculum workshop at a recognized venue such as SC, GGF, or HPDC.
- Create virtual-machine-based sandbox for safe testing on an extended Grid and easy install and teardown of lab resources.

F Conclusion

We propose to sustain and expand the Open Science Grid distributed facility as a common shared cyberinfrastructure to benefit science. The initial scale and capacity of the OSG Facility will be driven by our physics research stakeholders—US LHC, LIGO, STAR and Run II. The OSG Consortium will grow our grass-roots infrastructure into a dependable and effective nationwide distributed facility—a green field for innovation and discovery. These goals will be achieved through the three thrusts that constitute our program of work: the Facility, education, outreach and training, and extensions through partnerships with external development groups. The hands-on technical engagement and contributions of all OSG partners will move us forward to deliver an infrastructure that a broad and diverse range of scientific communities—from the single researcher to the 1000 strong physics experiments—can rely on. This will enable university and laboratory IT facilities to realize the promise of a universal cyberinfrastructure, will engage the software development community in the practical application of computer science in the service of challenging applications, will train a new generation in distributed information technologies, and will stimulate new and transformative scientific hypotheses for exploration and discovery.