SUSTAINING AND EXTENDING THE OPEN SCIENCE GRID:
SCIENCE INNOVATION ON A PETASCALE NATIONWIDE FACILITY

1-Year Extension for US LHC Deliverables

1 Abstract/Summary

Project Title: SUSTAINING AND EXTENDING THE OPEN SCIENCE GRID: SCIENCE INNOVATION ON A PETASCALE NATIONWIDE FACILITY (1-Year Extension for US LHC Deliverables)

Applicant: The Board of Regents of the University of Wisconsin System

Principal Investigators: Miron Livny, Ruth Pordes

This proposal document gives a brief summary of the current status and main accomplishments of the Open Science Grid (OSG) project to date, an outline of the future goals of OSG, responsibilities and work towards future capabilities for the US LHC for the one-year extension in the current project to the DOE Offices of High Energy Physics (OHEP) and Nuclear Physics (NP), and a brief outline of the status of planning for the OSG for 2012 to 2016.

Four years into the OSG project the US ATLAS and US CMS experiments have expressed their reliance on the OSG to support and evolve their underlying distributed facility in the US as well as provide them with many different services that enable their end-to-end systems to be effective. ALICE USA has recently joined the Consortium and communicated their needs for project deliverables. A significant part of the OSG contributions are those towards the World Wide LHC Computing Grid project on behalf of the US LHC software and computing projects.

Other HEP and NP experiments rely on OSG resources, software and services and contribute to the Consortium: STAR continues to ramp up its use of Virtualized and Cloud resources, Glue-X is prototyping the use of the infrastructure for simulation and data movement. D0 relies heavily on sharing resources across OSG for event simulations, and CDF relies on sharing distributed resources for data analysis. The Intensity Frontier experiments are reusing the infrastructure as their simulation and developments ramp up.

We propose this 1-Year Extension of the program of work to enable seamless continuation of support of the US-LHC experiments that rely on the OSG.

- This support will be used to provide part of the base services (agreed to using the “Assessment of Core Services provided to U.S. ATLAS and U.S. CMS by OSG” as a basis) to US-LHC.
- We will plan the new structure of OSG for the longer-term future program to be put into place by September 2011. In FY11 we will use OSG savings to cover the 25% OHEP funds missing for the institutions listed.
- We will work with all our sponsoring program offices to develop a framework starting in FY12. We will work with ASCR and OCI to understand the DOE and NSF scope for OSG as part of the advanced computing and CI visions.

OSG intends to fully support the US LHC communities in services and software for their distributed facilities, with leadership contributions to the LHC analysis needs, WLCG project, operations, software, security and Tier-3 support.

2 Key Proposal Data

Applicant/Institution: University of Wisconsin-Madison, Research & Sponsored Programs
Street Address/City/State/Zip: 21 N Park Street, Suite 6401, Madison, WI 53715-1218

Principal Investigator (PI): Miron Livny
PI Postal Address: 1210 West Dayton Street, Room 4367, Madison, WI 53706
PI Telephone Number: (608) 262-0856
PI Email: miron@cs.wisc.edu

Funding Opportunity Announcement Number: DE-FOA-0000412
DOE/Office of Science Program Office: High Energy Physics (HEP)
DOE/Office of Science Program Manager Contact: Alan Stone
DOE Grant Number: DE-FC02-06ER41436

Extension to Award #MSN106761 from DOE, CHICAGO OPERATIONS OFFICE
Sponsor Award Reference #: DE-FC02-06ER41436
Start Date: SEP 30, 2006 End Date: MAR 31, 2011

Sub-contract Institutions:

UCSD	PI: Frank Wuerthwein
U-Chicago	PI: Robert Gardner

Budget for these institutions is included in this proposal.

Collaborating Institutions:

BNL	PI: Torre Wenaus
FNAL	PI: Ruth Pordes
LBNL	PI: Doug Olson

These three national labs have submitted FWPs (and budget) thru their respective processes.

3 OSG Current Status and Accomplishments

Four years into the Open Science Grid project the US ATLAS and US CMS experiments have expressed their reliance on the OSG to support and evolve their underlying distributed facility in the US as well as provide them with many different services that enable their end-to-end systems to be effective. A significant part of the OSG contributions are those towards the World Wide LHC Computing Grid project on behalf of the US LHC software and computing projects.

Other HEP and NP experiments rely on OSG resources, software and services and contribute to the Consortium: STAR continues to ramp up its use of Virtualized and Cloud resources, Glue-X is prototyping the use of the infrastructure for simulation and data movement. D0 relies heavily on sharing resources across OSG for event simulations, and CDF relies on sharing distributed resources for data analysis. The Intensity Frontier experiments are reusing the infrastructure as their simulation and developments ramp up.

LIGO has expressed its continuing reliance on OSG provided software distribution and support. LIGO applications are increasingly able to access resources through the OSG infrastructure.

Applications from other science domains – most notably molecular dynamics, structural biology, and protein matching, in addition to increasing use by with chemistry, mathematical modeling and additional bioinformatics applications – rely on sustained production runs across the OSG.

The OSG communities and teams continue to be energetic and engaged, acting as a vibrant, well established virtual organization, sharing knowledge and expertise in addition to: operational services including production support on dedicated and shared resources available to the OSG user communities, problem reports and solution, cyber security, hosted services shared across multiple research communities; providing software collections with common packaging, distribution and support; consulting with existing and new user communities to provide planning and support for application preparation and production running. The OSG project contributes services to the mission and needs of the OSG Consortium. A summary of the OSG project accomplishments:

- Met LHC deliverables and needs for data taking.
- Met LIGO deliverables and application needs.
- Enabled science for multiple diverse research groups and benefited publications of results from other physics and non-physics science.
- Proved the (Virtual) Society organizational framework, principles, technologies and methods to include new entrants. Demonstrated that the whole is greater than the sum of the parts for existing members, new entrants and new science.
- Successfully moved single-community software to multi-community adoption (e.g. Bestman, Glidein-WMS).
- Success in federated operation with EGEE (now EGI) and WLCG.
- Accredited by DOE and NSF program offices and projects (e.g.ESNET, I2 and Teragrid, letters of commitment to both XD teams) as part of the ongoing distributed computing landscape.
- Advanced the national interests in services, support, and software for distributed systems physics and beyond.
- Improved project management processes

4 Planning for the Future

Participants in and contributors to the OSG have been actively discussing our vision and planning for the future beyond the current round of project funding. The following are our strategic aims:

- Sustain and expand a broad shared secure cyber-infrastructure in the US serving researchers spacing from the campus to the international arena.
- Enable end-to-end innovative research using distributed computing.
- Support and evolve usable, robust, secure software technologies for the distributed computing communities.
- Foster an excellent workforce through providing practical training and hands-on education.

Through implementation of the following goals:

- Fully support the US LHC communities in services and software for their distributed facilities, with leadership contributions to the LHC analysis needs, WLCG project, operations, software, security and Tier-3 support.
- Fully support LIGO in services and software for their distributed facilities, with leadership contributions to Advanced LIGO's evolving software, security and application needs, and better transparency between LDG and OSG in support of production analyses.
- Continue to come to agreement on and then support current communities production needs across the distributed infrastructure (STAR, Run II, Glue-X, ICECUBE, SBGrid etc.) – and new communities as they come to the table.
- Extend current "society" to include other communities, projects and partners.
- Fully integrate the campuses as part of the national cyberinfrastructure.
- Stronger partnerships with DOE distributed computing facilities and NSF XD, CF21, and future programs, advances and innovations in this area.
- Extend the current collaborations to interface to and benefit from satellites.
- Distributed computing laboratory for study and evolution.
- Enable science and engineering – i.e. both pure and applied research.
- Be agile with respect to new capabilities needed by the stakeholders.
- Help with the design and planning, requirements and constraints on external development.
- Help with hardening, transition to production, integration and support.

Specific areas where US ATLAS and US CMS have expressed interest in OSG contributions for future capabilities and capacities are bringing the following services to production use: Configuration Management; Efficient use of multi-core processor technologies; new storage and data management technologies and frameworks; Use of commercial and scientific clouds; and Virtualization. The references list several documents and requirements gathered from the US LHC as well as other communities relying on the OSG that we are using to plan for the future.

We are currently discussing a model that allows for a core project with targeted deliverables matched to the goals and mission of different sponsors. We have extended the mission statement to express this:

"*The Open Science Grid (OSG) advances science through open distributed computing. The OSG is a multi-disciplinary partnership to federate local, regional, community and national cyberinfrastructures to meet the needs of research and academic communities at all scales.*"

5 Program of Work for 1 year Extension for US LHC Deliverables

We propose this program of work under the following assumptions. Note that if our efforts to secure future funding from ASCR and MPS fail we will need to go back to the drawing board for funding after September 2011.

- The funds will be used to provide part of the base services (agreed to using the “Assessment of Core Services provided to U.S. ATLAS and U.S. CMS by OSG” as a basis) to US-LHC.
- We will plan the new structure of OSG for the longer-term future program to be put into place by September 2011. In FY11 we will use OSG savings to cover the 25% OHEP funds missing for the institutions listed.
- We will work with all our sponsoring program offices to develop a framework starting in FY12. We will work with ASCR and OCI to understand the DOE and NSF scope for OSG as part of the advanced computing and CI visions.
- By March 2011 we will develop more detailed plans for an agreed upon short list of additional focused activities of common interest to the NP and HEP communities that will advance the effectiveness of their computing infrastructure.

5.1 Deliverables and Budget Justification

The budget justification covers partial support for activities from existing OHEP and NP funded institutions. Other institutions in OSG will be continuing significant levels of support, especially in the Operations and Software activities[1] during FY11 and the first half of FY12 covered by existing program funds, member contributions and satellites. The detailed deliverables are defined in the detailed WBS and Statements of Work for each of the institutions for FY11, and later for FY12.

		One-year Extension Responsibilities
BNL	VO Layer (include WMS and Security). Forward Looking (Technology Group).	Operations of the ATLAS Panda system. Support for GUMS software. Leader of the Technology Group.
FNAL	WLCG Compliance and Interoperations. WLCG representing US interests and forward looking. Security Operations. Tier-3 and Production Support. Program & Project Management.	Operational Security, Representation and support for the WLCG accounting and information reporting and support. Representation in planning for WLCG. Tier-3 and US LHC Production storage support. Program Management and Administration
LBNL	Security Operations. Nuclear Physics.	Support for the OSG RA. Support for ALICE USA WLCG contributions (NP).
U Chicago	Tier-3 and Production Support. Integration testbed. Site support. Operations documentation.	Tier-3 support. Production coordination. Integration Testbed support. Release documentation.
UCSD	Security Operations. VO Layer. Forward Looking. Tier-3 and Production Support.	GlideInWMS Factory operations and support. Scalability testing and characterization, including support of US LHC VO specific targeted campaigns. Support for operational Security. Support for US LHC Tier-3 storage Hadoop implementations.

Table 1: Institutional Responsibilities for 1-year DOE Extension

Table 2 below maps the effort in the extension to the needs described in the “Assessment of Core Services” document. All FTE represent an approximate percentage. The FTE cost due to overheads and average cost at each institution vary. Note that in the future the profile of responsibilities and effort may move between institutions as part of the restructuring planning discussed in later sections and the supplement to this proposal.

[1] Note that these are explicitly Not covered by the 1 year extension responsibilities listed below.

	Sub Areas	Part of OHEP/NP Extension?	Needs	OHEP/NP Extension[2]
1	WLCG Interoperability and Integration	Partial	0.5	0.25
2	WLCG, Compliance with MOU (e.g. accounting, facility capacity reporting)	Partial	1.5	0.75
3	WLCG, Representing US Interests	Partial	0.75	0.25
	Operations, Grid Operations Center	None	3	0
4	Operations, VDT (Middleware Distribution), Integration Testbed, Documentation, Development Support	Little. WLCG specific - no core.	9	1.0
5	Operations, Cyber Security	Partial	2	1.25
6	VO Layer, Workload Management System Support	Partial	2.5	2.0
7	Forward Looking, Design and Scalability	Partial	1.75	1.1
8	Program Management & Administration	Partial	0.5	0.25
9	Tier-3 & Production Support	Partial	1.5	1.25
	Total		23	8.1

Table 2: Mapping of Effort in the 1 year extension proposal to the US LHC needs

5.2 Responsibilities and Deliverables

This section details the responsibilities of each institution to the areas defined above. These are in line with the current FY11 project WBS. Each activity works across the deliverables for all VOs supported by OSG. In addition, each work area addresses Engagement and User Support needs.

5.2.1 WLCG Interoperability and Integration

From the Assessment document: "OSG provides software components that allow interoperability with European grid sites, including selected components from the gLite middleware stack such as LCG client utilities (e.g. for file movement, supporting space tokens etc), and file catalogs (server and client). It is vital to the LHC program that the present level of service continues uninterrupted for the foreseeable future, and that all of the services and support structures upon which the LHC program relies today have a clear transition or continuation strategy."

Interoperability with the European infrastructures is currently needed by the LHC experiments and WLCG and not at present by other OSG stakeholders.

FNAL: Work with the relevant technical people in the WLCG, EGI, EMI, US LHC and OSG staff to plan, integrate and test interoperation of new services and software depended on by the US LHC experiments (CE, SE, Security, Information etc) being deployed on the different infrastructures and NGIs.

5.2.2 WLCG Compliance with MOU and Representing US interests

From the Assessment document: "U.S. ATLAS and U.S. CMS greatly benefit from OSG's Gratia accounting services, as well as the information services and probes that provide statistical data

[2] Based on the following distribution of effort: BNL 1.2, FNAL 2.4, LBNL 0.5, UCSD 2, U-Chicago 2

about facility resource usage and site information passed to the application layer and to WLCG for review of compliance with MoU agreements."

FNAL: Planning, management, tracking, communication, auditing, reporting and oversight:

- Work with the WLCG Management Board, US ATLAS and US CMS management to report, audit, and follow-up on missing and incorrect information for: accounting, availability and reliability, installed capacity, storage installations. Ensure the completion of the installed capacity reporting.
- Work with the WLCG Information Officer, US ATLAS US CMS, and OSG staff to ensure the central information reporting service is evolved to meet the scaling needs for LHC data analysis (increase in number of Tier-3 sites), deployment of a "OSG Top Level BDII" reducing the critical service reliance on the CERN BDII and the continuing issues with "glitches" that result in non-reporting of US sites as available for US LHC jobs.
- Maintain the work planning and reporting for US LHC (and other) deliverables in the OSG project area that contribute to the WLCG (WBS, Area Coordinators meetings, US ATLAS major account management etc).

LBNL: Coordinate between ALICE USA and OSG to ensure integration of extensions and reporting to add Alice USA to the OSG reporting and interoperation activities are met.

5.2.3 Integration Testbed, Documentation, Development Support

The contributions here do not cover any of the core effort of the Virtual Data Toolkit at this time. The deliverables include support for the Integration Testbed, documentation and testing of new releases of the OSG software, and some contributions to storage software support.

From the Assessment Document: "Middleware Deployment Support: Middleware deployment support is an essential and complex function that the U.S. LHC facilities are fully dependent upon. The need for continued support for testing, certifying and building a middleware as a solid foundation our production and distributed analysis activities runs on was served very well so far and will continue to exist in the future, as does the need for coordination of the roll out, deployment, debugging and support for the middleware services. In addition the need for some level of preproduction deployment testing has been shown to be indispensable and must be maintained. This is currently supported through the OSG Integration Test Bed (ITB) providing the underlying grid infrastructure at several sites with dedicated test instances of VO-specific services like PanDA, the ATLAS Production and Distributed Analysis system, running on top of it. This implements the essential function of validation processes that accompany incorporation of new grid middleware services and new versions thereof into the VDT, the coherent OSG software component repository.

Storage Services: U.S. ATLAS and U.S. CMS profit from OSG's support of storage services. The investment by OSG in support of BestMan SRM is an excellent example where OSG provided software services requested by other stakeholders that now have proven to be very beneficial to LHC sites. CMS and ATLAS now deploy BestMan as part of the preferred Storage Element solution at Tier3 sites, as well as a high performance Storage Element option for the Tier2 sites. OSG has been instrumental to improve the understanding of the complex issue of I/O characteristic and performance, and is providing advice and expertise to measure the I/O capability of sites."

LBNL: Support and testing of Bestman storage implementation.

U Chicago: Be responsible for the Integration testbed, the distributed testbed on which new releases of the VDT are tests, release documentation for new Site installations, revise ITB testing to speedup release cycle.

- Keep sites up to date with incremental updates.

- Provide support for on-demand and scheduled VO testing.
- Support sites through chat support
- Update documentation based on recommendations of documentation group
- Review and prune existing documentation to remove outdated materials
- Work with Content Management and Training to develop technical materials relevant to OSG site administrators.
- Contribute technical content and documentation relevant for OSG site administrators.
- Contribute technical content and documentation relevant for the persistent ITB.

5.2.4 Operational Cyber Security

From the Assessment Document: "One of the essential parts of grid operations is that of operational security coordination. The coordinator is provided by OSG today, and relies on good contacts to security representatives at the U.S. LHC Tier-1 center and Tier-2 sites. Thanks to activities initiated and coordinated by OSG (e.g. defining a security framework) a strong operational security community has grown up in the U.S. in the past few years, driven by the needs of ensuring that security problems are well coordinated across the distributed infrastructure. Part of this important activity is risk definition and assessment, security audits of facility components and training of the facility personnel. Policy development: Appropriate security policies are a mandatory foundation for sustainable operations of the world-wide computing facilities. The OSG security coordinator needs to participate in the work of WLCG's Joint Security Policy Group (JSPG). "

FNAL: Perform all functions as the OSG Security Officer, acting on behalf of the US LHC collaborations with the WLCG as the USA based Security Officer.

- Lead and deliver operational security for incidents involving US LHC VOs and sites, WLCG VOs and sites where there may be OSG impact.
- Perform risk assessments and mitigations.
- Perform vulnerability watch, assess impact on OSG assets and collaborators, and lead the efforts to mitigate them; and provide the dedicated effort required to address security alerts as they happen.
- Conduct the US portion of WLCG security drills.
- Participate in certificate and security fora where needed to provide the US input and tasks in parallel with the WLCG European arms.
- Respond to tickets related to Security coming directly to the OSG operations center or through the WLCG and European operations centers.
- Provide operational security support for the US LHC Tier-3s, help with training and documentation.
- Continue the work with the US LHC collaborations to deploy a more usable ID management infrastructure and services. Complete the deployment of the new certificate request and renewal frontends; in FY12 investigate and plan a more integrated ID management system for web and grid based services – working with such activities like the US CMS Wiki/Web services team deploying OpenID, Shibboleth based University ID federations, InCommon and the US Tier-1 and Tier-2 DOE Laboratories.
- Understand and discuss security needs of virtualization technologies and clouds.

UCSD: Provide support for the OSG Security Officer.

- FY11: Prepare how-to guides for sites and users including: How to install and configure the VDT; How to operate a site securely (for site admins, day-day tasks, how to register with OIM, read security advisories, check security team's signatures, etc); Security aspects of how to access OSG for end user (get cert, import into a browser, renew, register with VOMS, etc).
- FY12: Help with evaluation of ID management technologies that integrate web and grid access for LHC and WLCG.

5.2.5 VO Layer Workload Management System Support

This provides support for the VO workload management systems in use by US ATLAS and US CMS, specifically PANDA and GlideinWMS.

From the Assessment Document: "The U.S. LHC program requires that OSG personnel will continue to be involved in the process of developing and applying experiment-specific services for LHC data analysis on top of the OSG middleware stack. Examples include scalable workload management systems, like glideinWMS and PanDA and high performance systems for data storage and data access. While the development of such services resides in the experiments, the OSG provides support for integrating the services into the OSG and global Grid infrastructure. In PanDA, for example, the OSG provides for the integration of security infrastructure mandated and deployed by WLCG and OSG to provide secure and traceable operation of pilot-based multi-user workload management systems."

BNL: Support PANDA for US ATLAS and make extensions to enable a more general use of the services for US LHC and other communities.

- Support of LHC experiments
- Provide operational support of Panda Pilot interface with site authentication/authorization systems
- Use WMS to enhance OSG site testing
- Provide integration of Panda Pilot with the standard OSG Test Suite
- Create a light-weight Panda Data Mover
- Create a DAG-capable metascheduler for Panda
- Improve Panda monitoring capabilities
- Migration of Autopilot and Job Monitoring sections in Panda to the new monitoring framework
- Integration of Panda Monitoring with the Dashboard and other LHC information systems

UCSD: Support for GlideinWMS operations and extensions for US CMS and make administrative additions to make the operation more efficient and robust.

- Operate OSG GlideinWMS pilot factory according to Operations SLA
- Maintain and upgrade instances of glideinWMS pilot factory
- Monitor and help troubleshoot problems with using the glideinWMS

5.2.6 Forward Looking, Design and Scalability

As LHC data taking continues and the pace of user analysis expands it is essential that the services and software OSG provides to US ATLAS and US CMS keep pace with the increase in scalability and technology needs. Effort is needed to ensure that these new capabilities are robust, integrated, and usable in the OSG environment.

A specific deliverable of this responsibility during the 1-year extension period will be to define and plan medium to longer tern activities needed for the evolution of the US LHC and WLCG distributed systems – at present including configuration management, production use of multi-core processor resources, data management services, usability of clouds and production deployment of virtualization techniques. In order to meet this scope we will plan these serially with a deliverable of one plan every quarter.

From the Assessment document: "A technology working group under the leadership of the Technical Director of OSG comprising participation from U.S. ATLAS, U.S. CMS, LIGO and OSG is investigating, researching, and clarifying design issues, resolving questions directly, and summarizing technical design trade-offs such that the component project teams can make informed decisions. In order to achieve the goals OSG needs an explicit, documented system design or architecture so that component developers can make compatible design decisions, and virtual organizations (VOs) such as U.S. ATLAS and U.S. CMS can develop their own applications based on the OSG middleware stack as a platform. As a short-term goal the creation of a design roadmap is in progress."

BNL: Technology Group

- Coordinate OSG Blueprint Meetings - about 3 times a year
- Conduct technically planning in response to feedback and requests from major stakeholders; develop technology and architecture recommendations
- Publish quarterly reports to the OSG-EB & OSG-ET on issues, actions, and recommendations resulting from the work of the Technology group

FNAL: Support the Technology Group and related taskforces in planning future capabilities and needs.

UCSD: Scalability, Reliability and Usability

- Test scalability and reliability of storage
- Test scalability, reliability and usability of CEs - GRAM5 and CREAM as provided by VDT – and multi-core access
- Test scalability, reliability and usability of glideinWMS.
- Quarterly report on Condor scalability - evaluate new Condor and GlideinWMS releases
- Monitor and facilitate glexec deployment on OSG
- Test scalability and reliability of centrally operated services
- Perform a scalability and reliability test of WLCG services – GRATIA, BDII instance
- Investigate and develop usability guidelines
- Develop and improve monitoring tools

5.2.7 Program Management and Administration

It is essential that the contributions to the US LHC program be in good coordination and collaboration across both the specific activities as well as the rest of the OSG program.

FNAL: Partial contributions to project management, reporting, tracking, and Executive Director led activities.

5.2.8 Tier-3 & Production Support

OSG provides technical, procedural and training support for the Tier-3 administrators. The production coordinator ensures that the production and operational needs of the communities are met, are planned for, and are tracked on a weekly basis.

FNAL:

- Support storage implementations, documentation and operational planning for T3s.

U Chicago:

- Lead weekly Production calls, highlight important issues, and follow up on action items.
- Work with the VO, Sites, and Operation Area coordinators to coordinate activities related to production.
- Work with the Metrics Coordinator to monitor the performance of the OSG infrastructure and develop measures and metrics that assess the state of the facility.
- Serve as the OSG Tier-3 Liaison to the US-ATLAS and US-CMS communities and establish regular, ongoing, point-to-point contact with the US ATLAS and US CMS Tier-3 support staff to enable cross-VO communication and collaboration, where appropriate.
- Work with Atlas and CMS to define an operational plan for supporting T3s that includes appropriate VO and OSG combined efforts.
- Work with the Security Coordinator and Atlas/CMS T3 liasons to provide a practical "hands-on" security plan for T3s.
- Work with Production area to provide support for LHC Tier 3 sites.
- Create, maintain and use a virtual Tier 3 cluster to test different configurations

Test software solution targeting simple deployment for small sites.

6 Future Capabilities and Capacities

The US LHC experiments have expressed interest in OSG contributions for future capabilities and capacities. We are already planning these activities to be able to include the goals and deliverables in the future proposal for OSG. Initial evaluation and prototyping activities are in place for all these areas. We will be developing more detailed plans over the next six months; *the information presented below provides status and direction but these are NOT deliverables for the 1-year Extension.*

Work in these areas is lower priority tha the core support deliverables above. At the moment the work is being done as part of the Forward Looking area above, the core software areas, focused task forces with specific work goals (e.g. WLCG data management prototypes).

6.1 Configuration Management

The need is to make the configuration, both initial and updates, of individual grid resources and entire grid sites easier, more robust and more usable without rewriting the existing more than 20,000 “low level” lines of configuration scripts across the OSG software components.

Work to Date: A Blueprint discussion was held at BNL in May 2010, followed by design sketches and discussions within the VDT team in Madison.

Owners: OSG Software area: Tim Cartwright, US ATLAS: John Hover

Current status and thinking: The problems to be addressed are: (a) multiple configuration files and tools that must be used to configure a system, (b) too many configuration settings with little distinction between critical, must-be-changed settings and advanced settings with good defaults that only a few sites will ever touch, and (c) many opportunities to create invalid settings and to cause settings to drift out of synchronization with each other.

The proposed new configuration system adds a cohesive, usable framework on top of the existing OSG scripts that allows them to be then used underneath more general-purpose configuration tools in use at sites (e.g., Puppet, Bcfg2, cfengine, YAIM). The OSG configurations typically do a set of actions rather than single settings of these more general tools. For example, in the script that “Enables a new VO on a site” a username and password for the VO's database is created, the username is generated algorithmically from the VO name and other bits, and the password is randomly generated, with both then fed into a VOMS configuration command.

We goal is to provide a uniform interface to an installation’s configuration settings addressing the issues listed above. The main elements of the new system are: a central configuration settings repository for important user-visible and shared settings; a uniform and usable API for managing settings, both from outside the system and from within configuration scripts and a master configuration tool that applies all user-controlled settings to individual software components and a consistent and repeatable way. The API will be designed with site configuration in mind, so that integration at the site level with existing tools will be easy to implement.

Short term plans: The initial goal is to create the new configuration system to support VDT native packages for the OSG. While this work can start in FY11 we know that the deployment will take a long time. The LHC shutdown in FY13 will be a natural time for completing the development, integration and deployment. One reason is that a new configuration system means change for existing site administrators. We plan to make the transition as easy as possible, via documentation and automated tools to convert existing configurations to the new system. Our short-term plans are to complete a prototype of the configuration system, and begin using the system as needed to support our native packages. At the same time, we will share our plans with the broader OSG community and get feedback, both at the high level and at the detailed technical level from early testers of our native packages. We will iteratively improve the configuration system along with the supported packages.

Issues: The main issues are that (a) more discussion is needed within the OSG community to validate and improve our design and to get buy-in from site administrators and other stakeholders, and (b) a technical prototype is needed to verify that the on-paper design will work and have the desired properties. (c) The need for sustained effort to work on the system.

Longer Term Thinking: Once the native-packaging version of the configuration system is more established, we want to consider back-porting the system to Pacman VDT based on stakeholder demand. There are significant technical challenges as we start to affect the existing Pacman installations. Long-term, we need to address site-configuration challenges as well. We will need to create a set of templates for common grid set-ups (e.g., Tier3 CE, SE, worker nodes) to show how to integrate the existing site-management tools with the improved machine configuration system.

6.2 Efficient use of multi-core processor technologies;

The need is to provide end-to-end support for US LHC applications that use multiple-cores on a single CPU. The modifications to the experiment frameworks are in testing and the experiments have told the WLCG that they are ready for prototype testing of "all-cores on a CPU" job execution on selected sites. OSG needs to provide full support in all services involved.

Work to Date: An OSG High Throughput Parallel Computing Satellite was funded in 2010. The experiments are adapting their frameworks and algorithms.

Owners: OSG Production and Operations: Dan Fraser, HTPC: Miron Livny, US ATLAS: Torre Wenaus, US CMS: Liz Sexton-Kennedy.

Current status and thinking: Effective use of multi-core technologies is a crucial aspect of meeting future LHC computing requirements, especially as 16 and 32+ processor cores are rapidly becoming main-stream, and code developers have already begun to exploit this capability with codes such as AthenaMP and CMSSW (that also runs in multicore mode). The HTPC satellite is working toward making parallel multi-core jobs run just like regular high throughput computing jobs across the grid. Another aspect is the engagement of the all sites, but in particular the US LHC Tier-2 and Tier-3s, across the OSG to test and begin deploying this capability.

Short term plans: Toward this end, the HTPC team is currently focused on several different areas: 1) Working with both the GlideinWMS and Atlas teams on enabling pilot factories to submit, manage, and monitor HTPC jobs on the grid. 2) Upgrading the Gratia accounting service to correctly support accounting for individual cores in a multi-core scenario. 3) Enabling dynamic discovery of HTPC resources. 4) Documenting both how to set up and configure sites as well as how to submit HTPC jobs across the grid so that users can readily take advantage of this capability.

Issue: As the capabilities and needs of the users progress, and the management issues are solved, more sites across the OSG will need to enable their sites to accept HTPC jobs. There needs to be a greater awareness of the ease of adopting and using this technology.

Longer Term Thinking: HTPC jobs will soon "feel" just like HTC jobs across the grid both in job submission as well as job management, monitoring, and accounting. Sites will be readily able to accept and manage a mix of both HTPC and HTC jobs according to their policies without requiring HTPC users to utilize site-specific code in their job submission scripts.

6.3 New storage and data management technologies and frameworks;

The need is to provide end-to-end support for US LHC distributed systems data movement and access infrastrutures as they evolve. The WLCG is coordinating a set of demonstrators. A decision by the experiments and project on which ones become production services will be made in early 2012 and efforts to integrate the efforts as needed are starting. This is an area where we can expect significant evolution over the next few years.

Work to Date: WLCG data management jamboree. US based US ATLAS and US CMS demonstrators based on XrootD proxies and caching.

Owners: OSG: Brian Bockelman; US ATLAS: Doug Benjamin, US CMS: Brian Bockelman.

Current status and thinking: Over the last 5 years, grid storage has focused on VOs building a coherent system out of independent storage systems at sites. Each site provided access via an SRM endpoint, and the VOs were left to organize them into a coherent system. The WLCG adopted FTS for transfer management and VOs adopted or developed variety of different catalogues for tracking the location of files. The catalog and explicit transfer management systems have proven to be expensive to support. The catalogs try to present a “correct” view of the current state of storage at about 50 sites; LHC VOs have struggled to keep this functional. Additionally, the explicit transfer management systems like FTS require an experiment-specific transfer layer to drive transfers. Due to the difficulty of these grid systems, small sites often cannot participate and human-resource-constrained VOs find it difficult to use the grid.

In response to these issues, we have been exploring cache-based systems utilizing the Scalla software suite from SLAC. The architecture places large disk caches throughout the grid and provides a common (though possibly not central) endpoint users can contact for data; the requests are redirected to the local cache. The prototype architectures built allows either streaming of remote data or pulling it to a local cache. Since it is cache-based, the VO does not have to build data-movement systems – it is done transparently in the middleware. As the contents of the caches are transient themselves, there’s no need for persistent “correct” catalogs to track file location.

These caches are a promising tool for distributing data out to the grid and remove the need to track all files in the OSG, but do not remove the need to keep a catalog of the files at archival sites. These archival sites are assumed to be more reliable and smaller in number. We having been experimenting with the iRODS system to see if this can fill the catalog and data management role for small VOs in the OSG.

Short term plans: The xrootd demonstrator projects will continue to expand and rollout as they are adopted by LHC T3s and, hopefully, the experiments themselves. As the technology has been demonstrated, we believe the OSG will take an increasing role in packaging and deploying the technology in its grid. Once packaged by the OSG, we will start to focus on deploying xrootd out to all sites. If available opportunistically at many large sites, we believe it can start to be adopted by non-LHC VOs.

Issues: As xrootd picks up inertia in the LHC community, OSG needs to be well-positioned to help sites transition to this new architecture. OSG already has experience packaging the software as a self-contained storage element, which we will need to leverage into native packaging and new configurations for the caching architecture. OSG will also need to play a support role for system administrators in deploying xrootd and VOs who want to adopt it.

Longer Term Thinking: The combination of iRODS for management of archival storage and Xrootd for distributing data on the grid would be the first time the OSG has offered a top-to-bottom solution for data access and management. As with glideinWMS, this could be a major technological “selling point” for the OSG, in addition to our current personnel expertise in the area. Xrootd has built-in mechanisms for load-balancing and redundancy. It can support multiple security protocols, has a bittorrent-like multi-source mode, and is trivial to use with firewalls and NATs, unlike GridFTP.

6.4 Use of commercial and scientific clouds;

The need is to ensure usability of available cost-effective resources for the experiments end-to-end systems. The US LHC collaborations are currently investigating the usability of commercial and scientific clouds. The WLCG related working groups are moving ahead to provide the software infrastructure for provisioning and deploying the software and needed services. OSG needs to provide an interoperable environment for these capabilities.

Work to Date: WLCG investigations of tools and interfaces. US ATLAS use of NERSC Magellan and EC2, US CMS extensions of GlideWMS and use of ANL Magellan and Amazon EC2, ExTENCI joint work on use of WISPY at Purdue for US CMS and Clemson Cloud for US CMS.

Owners: OSG Technology Investigations: Brian Bockelman, US CMS: Burt Holzman, US ATLAS: Paulo Calafiura.

Current status and thinking: On OSG the LHC experiments use pilot-based systems where the VO requests a certain number of batch slots from a (possibly centralized) factory or manager, which submits jobs using Globus GRAM to OSG sites. The jobs start up and join the VO's batch system. The VO requests resources and the manager is responsible for allocating them from the existing grid sites via grid jobs.

The expectation is that this model can be adopted onto the cloud. VOs will request an allocation of resources from the factory, which will be able to service the allocation through directly starting virtual machines on clouds in addition to the grids. CMS and ATLAS have shown it is possible to run their software on the cloud; CMS has demonstrated the ability to integrate the clouds into the existing workflow systems (CRAB and ProdAgent), in addition to remotely streaming analysis data from locations on the grid.

While clouds don't bring any revolutions to data management, the Amazon S3 storage cloud has shown users find it easier to deal with data exposed through a single endpoint (and managed transparently by the system underneath), as opposed to managing many endpoints. We believe this to be complementary to our current thoughts on data management.

There are three primary challenges to adopting clouds as an OSG resource:

1. **Lack of a standard protocol**. Nearly every cloud middleware has its own communication protocol; as the implementations are immature, they sometimes implement the protocols inconsistently.
2. **Virtual Machine image contextualization**. Different providers may require different kernels, image formats, and require specific customizations to the site. For example, CMS needs a mechanism for the site to specify the closest HTTP proxy server. If there
3. **Security**. Operational security in a cloud environment will need to be understood, as well as the integration of authorization and access control policies.

Short term plans: We plan on integrating one or two cloud resources and exposing them to users of glideinWMS, while extending the glideinWMS capabilities through the satellite CorralWMS project. We are working with DOE Magellan and Purdue Wispy (via the ExTENCI project). We will work with the providers and the most interested VOs to provide a single image per site that is similar to an OSG worker node. GlideinWMS has a first implementation of cloud support in the factory, and we believe this code will make it into production over the next year. By limiting the number of sites and forcing VOs to share images, we will sidestep the lack of standard protocol and the contextualization issues. We are working with the Condor team to make sure Condor-G supports any important upcoming middleware protocols (such as OpenNebula) as our stakeholders need it. Additionally, we will work with the OSG security team to develop (or preferably adopt) a security policy for sites and VOs endorsing virtual machines.

Longer Term Thinking: It's possible that VM and Cloud technology will become increasingly popular for HTC, especially for VOs who have difficulty in deploying their software. This enables more potential "customers "for the OSG project.

It is difficult to imagine a scenario where the resources of the current OSG would consolidate under a single cloud provider, so we believe the current approach combining the pilot model with clouds will be the most effective. It is easy to imagine a case when a significant portion of our sites will have virtualized resources in addition to batch-system ones. The work proposed above will put us in the best position to support a variety of "gatekeeper" technology. By introducing heterogeneity through the integration of clouds, we will provide a roadmap for incorporating other gatekeeper technology such as Condor-C.

7 General Status of Planning for the OSG for 2012 to 2016

The status of the planning for FY12 through FY16 is as follows:

- We will call out the target community for each goal and deliverable. That is deliverables can be directed to US LHC, US ATLAS, US CMS, LIGO, ALICE USA, Glue-X, Biology, Campuses etc.
- We will restructure the project to provide more effective delivery to the users and communities with the following Areas: Program Staff[3], Security, Software, Operations and Production, User Support, Technology investigations, Campus infrastructures.
- We will roll up the deliverables, and associated effort, by stakeholder so that we will separately describe the full set of deliverables for the US LHC, WLCG, LIGO, etc. This will require more detailed breakdown of deliverables such as Virtual Data Toolkit release, security incident response etc, than we have had to date.
- We are progressing with discussions of “association” between OSG and the ESNET and NERSC facilities, in order to explore support for the general facility operations and software through existing DOE ASCR facilities. Given continued positive feedback, we will together plan a face-to-face meeting between the management groups within the next few months.
- We are in contact with NSF (MPS and OCI) in exploring opportunities for a cross-European-USA program for future technologies needed by the LHC and LIGO. Several areas of work might have some potential: Use of Cloud computing; Green Computing; Virtual Computational Environments (e.g. ROOT); Advanced Computational Frameworks (workflows, ontologies etc); Data access and preservation (global file systems?); Global ID management; Community structures that foster and support virtual communities.
- We have led a workshop to input recommendations for the NSF OCI Software Institute program and hope to continue a dialog with NSF as the program is defined and published. Additionally we have had members of the Campus Infrastructure NSF Task Force.

7.1 Organization

The program of work of the OSG project is established annually (starting in October) through a single WBS for the deliverables and staff allocations across all funding sources, established by the “area coordinators” and prioritized and managed by the project Executive Team (ET). Independent “Satellite” projects have been sponsored by the agencies that contribute other services to the OSG mission in close cooperation and collaboration with the core OSG project. This model is proving valuable in extending the benefits to OSG communities while ensuring coherent capability, technology and service evolution.

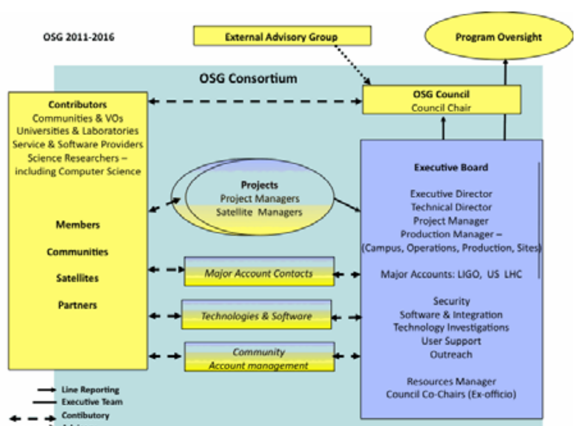

Figure 1: Proposed New OSG Organization Structure

The new organization structure of OSG will move to a smaller number of technical areas based on experience of management of the

---

[3] Includes WLCG interface and major stakeholder (US LHC, LIGO) coordination, Executive and Technical Directors; Program and production management; Assessment and metrics; Architecture and technology (including interoperation) oversight; Communication, external relations and administration.

deliverables to date. This will accompany an expansion of the satellite model. The program staff will ensure that the whole infrastructure and community stay coherent, harmonized and contribute to the overall mission. The proposed new organization chart is shown.

Table 3: Mapping of Area in "Assessment of Core Services Provided to US ATLAS and US CMS..." to New OSG Structure; Planned for at or before Sept 1 2011.

Sub-area in "Assessment of Core Services provided to U.S. ATLAS and U.S. CMS by OSG"	Area Responsible
WLCG Interoperability and Integration	Production and Operations
WLCG, Compliance with MOU (e.g. accounting, facility capacity reporting)	Production and Operations
WLCG, Representing US Interests	Program Office
Operations, Grid Operations Center	Production and Operations
Operations, VDT (Middleware Distribution), Integration Testbed, Documentation, Development Support	Software
Operations, Cyber Security	Security
VO Layer, Workload Management System Support	User Support Production and Operations
Forward Looking Design and Scalability	Technology Investigations Software
Program Management & Administration	Program Office
Tier-3 & Production Support	Production and Operations

Also, our ability in the future to make decisions quickly, clearly and with a common understanding of expectations, will remain an important part of our strategy and ability to be flexible and to make timely reactions to changes in or new user requirements, new opportunities and available technologies.

7.2 Future Planning

We have asked for preliminary goals for each of the new areas of work. We continue to work internally on these plans. Here is a summary to date:

Production & Operations: *Production:* Identify, extract, and socialize best practices and coordinate these with the community to execute best practices in a manner that minimizes disruptions in the productivity. Reliability, as noted above, is the crucial element that allows scientists to depend on this capability, and improving the reliability will remain a primary focus. *Operations:* Continue to operate and improve operations for the hosted and support services. It is expected that the number of services will increase for OSG prime. *Site Coordination:* Maintain as well as develop new ways to communicate among the ever-diversifying number and scope of OSG sites. This effort also involves writing and organizing documentation to ensure that the best practices remain current.

Campus Infrastructure: The goal of this effort is to adapt the OSG capabilities to meet the needs of the campuses. This will require OSG components to be modified in some cases for flexibility in security, job submission, accounting, scheduling, monitoring, and user management. The goal is to support 2 additional campuses per year.

Security: Each year has the goal of continuing operational security[4]. The overarching goal for the next five years is to strengthen our abilities to detect, monitor and response to security

[4] Includes ensuring incident response, vulnerability detection, ST&E, secure communication channels with security contacts, attending to security tickets (security related software problems, cert problems) etc.

attacks proactively and reactively, including: FY'12: Complete resolving most-pressing problems in ID Management. Completion of work started in FY10. i) Complete integration of federation based id management technologies in OSG. ii) Identify the most important OSG assets. Identify ways to monitor and mitigate the risks. FY13: i) Develop a risk model and identify mitigation and monitoring methods. Identify/design or develop monitoring tools that can be provided to the site/VO security contacts. Develop an attack spread model that assesses how an attack spread across sites/users can affect overall grid security. ii) Create a software testbed for the purposes of security testing. This allows us to catch errors and also to incorporate new software into OSG quickly and securely. FY14: i) Continue with completing risk management solutions. Provide sites and the VO resources with developed solutions so that monitoring can be achieved in a distributed manner. ii) Complete the software testbed project. iii) Analyze and develop security models that can work with cloud based services, including understanding affects of virtualization on our risk model, access control, data storage, data access and privacy. FY'15, '16 i) Continue developing security models that can work with cloud based services. ii) Access needs of data storage systems. For all these we will engage the CS research community when useful to develop ideas/solutions.

Software[5]*:* Continue support for packages and integration needed by the user communities. Plan to adopt packaging for native Linux distributions instead of a single build description. There would then be a source package per platform. The VDT software cache could get much smaller and we could donate-to/maintain-within distributions. The VDT's main value would be less about the software cache and more about: High-quality configuration; Integration testing; Packaging expertise shared within our community.

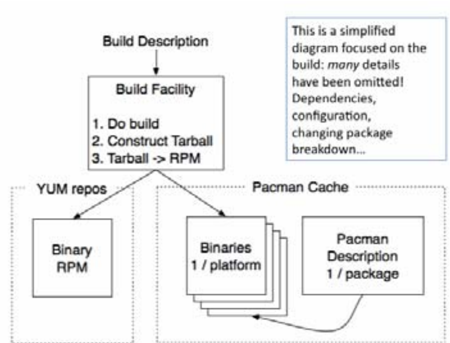

Figure 2: Software Build Process

User Support: Offer a one-stop mechanism to support the user communities as they plan for and use OSG services. Understand their needs, provide technical guidance on how to grid-adapt their applications, help resolve issues and debug technical problems, and create a collaborative community among VOs. Reduce the elapsed time for new science communities in achieving production (analysis, simulation, etc) on the OSG. Provide a single point of contact which enables: A technical review of a community's computational approach by a panel of experts within OSG providing guidance and technical recommendation to the VO on how best to adapt their applications for execution in OSG; Technical support for the VO as they adapt their software and begin to deploy and ramp-up to production; and Troubleshooting support during the initial stage of running on the OSG. The goal is to enable two new science communities per year.

Technology Investigations: Manage technological change throughout the next five years. The choice of investigations will be driven by the requirements and interest of OSG stakeholders and council. Focus on short, 6-9 month investigations of *existing* technologies that can be adopted by the OSG *and the stakeholders*. It is of course very difficult to predict five years down the line. Here is the current list of planned investigations: FY'12: Data Management, Cloud integration; FY'13: Beginning Job Monitoring, ubiquitous use of multi-core technologies; FY'14: Advanced Job Monitoring, Data management and preservation; FY'15: Integration of multiple Cloud resources.

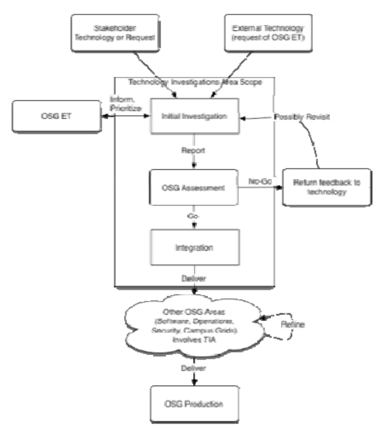

Figure 3: Process of Technology Investigation

Program office: Continue to: foster and extend the community

[5] From presentation at recent OSG meeting.

and team environment of the OSG; Work closely with and support the OSG Consortium Council; Together with the communities, plan technologies and services in support of the next running periods of LIGO and the LHC; Identify and work with new clients and contributors, identifying any new "major stakeholders"; Promulgate timely, agile project management; Incorporate satellite interfaces into coordination and reporting; Hold blueprint meetings as needed; Improve multi-media communications; Increase collaboration with the NSF XD and CF21 programs, DOE ESNET and NERSC facilities and other projects as useful.

8 References

1. Community Requirements
2. Assessment of Core Services provided to U.S. ATLAS and U.S. CMS by OSG
3. Software Hardening
4. Campus Communities
5. Workforce development
6. OSG-TG Principles of Collaboration
7. Current Architecture (Draft) Definition of Satellites
8. OSG At-Large Virtual Organizations FY 2010-11 Plans, Needs, Requirements; Consortium Stakeholder Input to the OSG Council, and to the OSG Executive Board