Assessment of Core Services provided to U.S. ATLAS and U.S. CMS by OSG

February 25, 2010

Introduction

Four years into the Open Science Grid project the LHC experiments, together with a number of other sciences, can well rely on the OSG to provide vital services for their scientific computing infrastructure. With the recent start of LHC data taking and the start of physics data analysis U.S. ATLAS and U.S. CMS open a new chapter in their use of

OSG. In the coming years the LHC science will require both production quality services, open to many hundreds of simultaneous users, while simultaneously needing to scale up the computing environment almost exponentially over the coming years, to keep up with the ever-increasing data samples expected from the LHC.

Strategic Importance of the OSG Consortium

The computing systems for both ATLAS and CMS consist of a grid of more than a hundred distributed computing centers. The Grid approach helps to distribute the huge and complex multi-PetaByte data sets and sophisticated analysis and data simulation needs across a large number of sites. Distributing the computing also addresses the complex funding structure of a truly distributed worldwide collaboration of hundreds of universities and labs. Different from a company or a single institutional computing center with top-down management and coordination covering all required services, the LHC experiments are fundamentally a group of researchers that need to rely on a loosely affiliated set of computing centers.

The Open Science Grid, working with the Worldwide LHC Computing Grid project (WLCG), provides the binding glue across centers in the U.S. Each U.S. center is a member of the OSG consortium and participates directly in the OSG. When new sites join, in particular individual universities with local computing for their local science group, they become members of the OSG and participate in and profit from the services provided.

OSG develops and maintains a center-of-expertise in High Throughput Computing (HTC) and Grid Computing which is leveraged by U.S. ATLAS and U.S. CMS as new needs become apparent.

OSG is viewed as the right forum for building a community for establishing best-practices and knowledge sharing about building and evolving the computing infrastructure within the campus of institutions in the U.S. participating in LHC data analysis. To be specific this should include more efficient and professional system administration, which will lead to reduced downtimes, less manual configuration, improved system tuning, reduced costs, and more efficient use of the scientist’s time. Conversely, in terms of connecting campuses to national infrastructure we require OSG to develop a strategy across all existing work areas to support these connections.

In summary, OSG plays a key role regarding our participation in a national community which may help leverage local investments, increase revenue and diversify funding resources.

In such a distributed environment the LHC science program cannot just rely on bilateral agreements with individual sites, even if some of them like the participating national labs have the know-how and IT infrastructure to provide a large set of the required services. The OSG consortium addresses the need for a homogenous approach across all sites on important issues. This includes not only the middleware and interfaces to site services, but also the approaches and procedures concerning computer security, incident response, user registration etc. In addition the OSG provides an important forum between diverse set of IT providers, systems and applications developers, and science users, to address the technology advances required to scale the computing resources as required by the evolving scale of LHC computing.

Domain and computational scientists work together to solve the end-to-end application needs based on advancing the principles, methodologies and frameworks of large-scale computing as it is needed in the worldwide distributed LHC computing facility. OSG teams in close collaboration with U.S. LHC facility staff apply the conceptual insights gained on the distributed infrastructure. U.S. LHC computing is benefiting from advances OSG has made in job management allowing transparent use of heterogeneous clusters as a uniform distributed facility; OSG maintaining and evolving a secure high-availability production infrastructure; OSG provisioning remotely accessible storage caches for middleware code, and from use and management of independent Cyber Infrastructures in a federated model which allows for scaling of the whole system.

OSG thus has become a major strategic component for the US LHC scientific programs addressing critical needs for LHC computing. OSG also benefits university computing centers and national laboratories that are providing computing for science. It allows them to provide and manage their facilities across their broader program and to capitalize on economies made possible by sharing expertise and support.

OSG Services and Software

US LHC relies quite extensively on services and software provided by OSG, as well as on processes and support systems that have been produced or evolved by OSG. Over the course of the past years US ATLAS and US CMS have invested heavily in OSG in many aspects – human and computing resources, operational coherence and more.

In addition the OSG efforts have aided the integration with WLCG partners in Europe and Asia, which is essential to the operation of the worldwide distributed ATLAS and CMS computing facilities. OSG has been crucial to ensure U.S. interests are addressed in the WLCG. The U.S. is a large fraction of the collaboration both in terms of participants and capacity, but a relatively smaller number of larger sites within the WLCG collaboration.

The components and procedures developed in the process have become the basis for support and operation covering the interoperation between OSG, EGEE, and other grid sites relevant to ATLAS and CMS data analysis. 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 continue 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.

Middleware Architecture

OSG provides its middleware distribution as a heterogeneous software system consisting of components contributed by a wide range of projects. Based on experience and observations U.S. ATLAS suggested to start working on a coherent middleware architecture. One of the difficulties we ran into several times was due to inter-component functional dependencies that can only be avoided if there is good communication and coordination between component development teams. 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.

OSG Support Services

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. In fact US LHC relies on the VDT and OSG packaging, installation, and configuration processes that lead to a well-documented and easily deployable OSG software stack. The OSG software team provides a comprehensive repository as well as the associated caches, and resolves cross component issues that naturally develop with the evolution of multiple independent software products.

Support Infrastructure

Regarding support services the OSG Grid Operations Center (GOC) infrastructure at Indiana University is at the heart of the operations and user support procedures. It is also integrated with the GGUS infrastructure in Europe making the GOC a globally connected system for worldwide ATLAS and CMS computing operation.

Operations and Service coordination

Operations and service coordination is currently CERN’s responsibility, in close cooperation with the OSG GOC, and the Regional Operation Centers (ROC) in other regions. While it is likely that the WLCG service coordination will remain at CERN after EGEE ends, in Europe it is anticipated that ROC responsibilities will move to the Tier-1 centers (unless the European National Grid Infrastructures are in place and able to do this). In the U.S. we rely on the OSG GOC to maintain the function of the ROC.

Network Coordination

Operational connections between grid operations and the network service providers in the U.S., prominently ESnet, USLHCNet and Internet2, and GEANT/NRENs in Europe are needed. Tools to support workflows and a clearing house for network related tickets should be provided. This allows the coordination of grid issues and a single point of contact to the network community, in the U.S. and with connections to global instances (e.g. ENOC).

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.
We expect that future storage developments — which we know will be required to scale up the storage systems to unprecedented sizes in the coming 5 years — will profit from a strong collaborative approach between OSG stakeholders, and that OSG will be ready to provide integration, deployment and operational support for the next generation of storage solutions.

Operations Tools

Essential tools have been developed and are still being developed that are required for the daily operation of OSG, as well as the overall management and reporting including accounting, configuration management, and operations support. U.S. ATLAS and U.S. CMS require that maintenance and development remains part of OSG's responsibilities.

Accounting

Accounting services are provided and maintained through OSG's Gratia data gathering infrastructure. This infrastructure is well integrated with the ones provided in other regions. The accounting data is published into the APEL database and thus available to the WLCG.

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

Configuration Management

U.S. ATLAS and U.S. CMS would benefit from a configuration data base such as the GOCDB. This is considered an important service that could be used to define all the services and sites within the distributed OSG facilities, as well as contact and management data for the services. Monitoring and reporting tools would use configuration data provided by the configuration database. With the OSG Information Management System (OIM) OSG provides a place where sites can advertise their downtimes, which is used by our facilities, but it is felt that full configuration management capabilities are needed.

We understand configuration management to mean a distributed tool that makes system (OS and software) deployment and configuration simple and reliable. If a particular server needs to be restored, re-created, or duplicated the system updates itself according to a centrally defined profile.
In the context of OSG, the idea would be a tool that allows administrators to deploy, change, and backup and restore (via versioning) their entire site software configuration.

Operations Portal

OSG currently operates the OIM and MyOSG systems to collect and provide key data pertaining to the VO and Sites; further improvements to this portal could be useful

- To provide contact information for sites and services
- To host a broadcast service as a mechanism for OSG-wide publication of service changes etc.
- To support automated reporting of daily and weekly operations issues

Availability monitoring tools

The collection and reporting of facility services availability and reliability data is an important functionality provided by the OSG, which is part of the contractual MoU between the U.S. LHC operations program and the WLCG project. RSV probes are installed on all relevant facility service components at the Tier-1 and Tier-2 sites. The data is centrally aggregated and published to the WLCG repository.

Security and Policy

Operational Security Coordination

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).

Grid Certificates

As grid certificates for users, hosts and services are a mandatory prerequisite we need OSG to operate and further develop the associated services, i.e. the DOE Grid Registration Authority and interfaces to the ESnet DOE Grids CA for documentation, monitoring and validation.

Application Support

Utilizing a Grid-based computing approach, computational resources are distributed over many independent sites with only a thin layer of Grid middleware shared between them. This deployment model is very convenient for computational resource providers at Tier-1 and Tier-2 centers around the globe. They can continue to operate their local distributed resources according to local preferences and expertise, integrating them easily with other, non-Grid resources.

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.

Virtualization

From a technology viewpoint, the number of cores per machine will continue to increase in the near future and the challenge becomes implementing software in ways that can efficiently exploit them. The increased number of cores per machine has helped to drive the rapid adoption of virtualization. In addition to its benefits for resource consolidation, virtualization creates opportunities for a more flexible approach to offering computing services. Both technologies are rapidly maturing, particularly in terms of performance and management tools. Physics applications can benefit from these advances but computing services need to adapt to support them. U.S. ATLAS and U.S. CMS have collected application requirements and experience gained so far using multicore and virtualization together indicate a need for support beyond test environments. With ATLAS and CMS have worked to run production in a virtual environment, and issues of packaging and releases are being addressed.

Based on current understanding we propose the following actions that are meant to be carried out at U.S. LHC facility sites in close collaboration with OSG and similar activities in other regions as well as forums like HEPiX:

- Provide infrastructure at the centers for the preparation of virtual machine (VM) images, in particular CernVM and the Virtual Organization's application software delivery to them. VM images can be generated by tools offered by the CernVM project (http://cern.ch/CernVM), which provides a virtual software appliance for developing and running LHC data analysis.
- Include the capability to run VM images at Tier-1 and Tier-2 virtualized batch systems.

- Develop and test scheduling options for parallel jobs in mixed workload environments.

Actions requiring Grid-wide collaboration:
- Establish procedures for creating images that can be trusted and run at Grid sites. This is needed for Virtual Organizations like U.S. ATLAS and U.S. CMS to be able to run their images on their own facilities and on opportunistic resources.
- Investigate scenarios for reducing the need for public IPv4 addresses on Worker Nodes. Virtualization is increasing IP address usage and given the IPv4 address limitations (www.ipv6actnow.org) public IPv4 addresses need to be used wisely.
- Deploy multicore performance and monitoring tools (e.g. KSM, PERFMON) at U.S. LHC facility sites.
- Provide input to initiatives for running multicore jobs Grid-wide, e.g. MPI (Message Passing Interface) Working Group recommendations.

Grid interoperability with clouds:
- Prototype a solution to run Grid jobs on academic and commercial cloud resources.

Effort

Current OSG LHC services and the associated effort

Major Area	Sub Area	Effort [FTE]	U.S. LHC Priority
WLCG	Interoperability and Integration	0.5	High
WLCG	Compliance with MOU (e.g. accounting, facility capacity reporting)	1.5	High
WLCG	Representing US Interests	0.75	Medium
WLCG Subtotal		2.75	
Operations	Grid Operations Center	3	Medium
Operations	VDT (Middleware Distribution), Integration Testbed, Documentation, Development Support	9	Medium
Operations	Cyber Security	2	High
Operations Subtotal		14	
VO Layer	Workload Management System Support	2.5	Medium
Forward Looking	Design, Scalability	1.75	High
Program Management & Administration		0.5	Medium
Tier-3 & Production Support		1.5	Medium
Total		23	

Non-LHC Areas supported by OSG and the associated effort

Other Areas	Effort [FTE]
LIGO	2
Engage	1.5
Communication, management, administration	1.5
Education, training	1.65
VO support	1.5
VDT	0.7
Operations	0.2
Integration Testbed	1.1

Cyber Security	0.8
SBGrid	0.8

Additional future OSG LHC services and the associated effort

Major Area	Sub Area	Effort [FTE]	U.S. LHC Priority
Configuration Management	• Across services on different hosts • Local configuration versioning • Software repositories	1	High
Integration of Commercial and Scientific Clouds	• Transparency at application layer • Cloud interfaces at sites • Data handling & Data provisioning	2	Medium
Usability for collaborative analysis	• User access to shared storage (incl. change management) • Grid-level diagnostics • "dynamic collaborations within large VOs • Overall scalability of services	2	Medium
Active management of shared capacity, utilization planning, accounting and reporting, and change		1	Medium
End-to-End Data Management challenges in light of advanced networks	• Dynamic circuit reservation • Dynamic data Placement	2	Medium
Total		8	
Total U.S. LHC related effort (present and proposed) supported by OSG		**31**	
Total OSG supported effort (includes non-LHC related effort)		**42.75**	

Resources for EGI.org *

Costs for	EGI.org Costs (in FTE/a)
Operations	17
Middleware interfaces and final certification	8
User Community Services	11
External Liaison Functions	4
EGI.org Management and Administration	11
Total	**51**

* as presented by Bob Jones – EGEE Project Director – at the EGEE III Review in June 2009
http://indico.cern.ch/getFile.py/access?contribId=2&resId=1&materialId=slides&confId=53198
Note the list above does not include effort devoted to middleware development as defined in the proposal of the European Middleware Initiative (EMI). More information about EMI is available at
https://register.nordu.net/speakers/files/AlbertoDiMeglio.pdf

Existing and foreseen Service Level Agreements (SLA) for OSG Core Services

Critical Priority

- BDII
- MyOSG

High Priority

- CA Distribution includes OSG Software Cache
- RSV including WLCG RSV to SAM Reporting

Normal Priority

- OSG Display
- OIM
- MIS VOMS
- GRATIA
- WLCG Comparison Reports
- GOC Ticketing
- Ticketing Web Interface
- Notification Services

SLAs are documented at
https://twiki.grid.iu.edu/bin/view/Operations/ServiceLevelAgreements

Potential issues associated with transitioning of the OSG core services to the LHC Operations Program

Risk	Impact	Occurrence Probability	Mitigation
Work in progress cannot be completed prior to end of OSG funding	This might cause parts of the OSG program of work not to be delivered	High	This will have to be addressed by a strong and clear agreement governing the priorities, roles and responsibilities between OSG management and the stakeholders, by the regular progress monitoring through the Executive Team and Executive Board and the Council, as well as a review process whereby all activities will be reviewed to track their progress.
Loss of technical expertise leading to failure to provide the functionality needed by existing and planned for applications	This would cause applications not to be able to use the Grid infrastructure	Medium	This needs to be addressed pro-actively by an application driven evolution of the foreseen OSG infrastructure governed by the Executive Team. Close collaboration with OSG management and technical staff will ensure U.S. ATLAS and U.S. CMS requirements are appropriately taken into account.

Risk	Impact	Occurrence Probability	Mitigation
Failing to implement changes necessary for the transition to a sustainable infrastructure while continuing to provide a stable service.	This would affect U.S. ATLAS and U.S. CMS reaching their goal of a smooth transition towards a sustainable infrastructure.	High	The procedures, processes and governance of a future sustainable Grid Infrastructure in the U.S. needs to be elaborated by OSG and the stakeholders. Close links with OSG will be set up and specific deliverables and milestones will be identified to monitor progress. Yet there is still the risk that a smooth transition will not be possible in the lifetime of OSG because the final structures are not specified well enough. In this case OSG will have to continue to provide its service unchanged and prepare for future transitions that will occur after OSG's lifetime. Funding must be secured to allow operations in this transition period since the Grid infrastructure is now mission-critical for our operations.