Open Science Grid

Answers to Reviewer's Questions

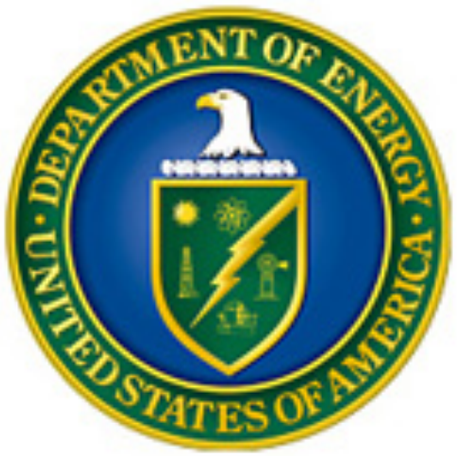

OSG Agency Review

Arlington, VA, August 20, 2014

Open Science Grid

LHC Continues to Depend Critically on OSG

- OSG enabled the US LHC to deliver the pledged computing resources
  - ★ for CMS, the US delivered more than 40% of the Tier-1 resources worldwide and 35% of the physics data analysis capacity
  - ★ over the past year, US provided more than 0.5 Billion CPU hours to WLCG
- OSG enabled USLHC scientists to be leaders and active participants in the discovery process, by delivering data to the Tier-2 and Tier-3 center
  - ★ for CMS > 5 PB to Tier-2 and > 1 PB to Tier-3 centers over the past year
- OSG enabled US LHC to contribute more than 2.5B simulated events to CMS over the past year (Atlas ?)
- OSG enabled US LHC to provide first-in-class Tier-1 and Tier-2 facilities
  - ★ measured in availability and reliability metrics that OSG collected and reported to WLCG
  - ★ 6 of the 7 USCMS Tier-2 sites were in the top-10 (of 52 worldwide) in terms of availability and reliability, demonstrating the reliability of OSG services
- OSG ran and processed more than 6,000 probes per day to generate reports to WLCG
- OSG is positioned to support the Run 2 data challenges in the same way we supported the data challenges Run 1

Q1: OSG interact with other grids

- How does OSG interact with other grids or infrastructure stacks worldwide, including WLCG, EMI, and similar efforts?
  - Members of the Executive Team have roles and interactions with other CIs including and as examples:
    - WLCG – many visits to CERN, participation in workshops such as the Federated Data Storage etc. Membership in the WLCG Management Board.
    - XD – OSG PI/Technical Director is a member of the XD Service Provider Forum.
    - XSEDE – OSG Resources Manager/Application Coordinator a member of the XSEDE Users Advisory Committee, and PI/Technical Director is one of the two reps from the XD forum to the external advisor board of XSEDE.
    - XRAC – Head of User Support/Project Manager attends XRAC meetings to review allocation requests and discuss “directing appropriate users to OSG” as well as participates in the XSEDE quarterly meetings on request.
    - ESnet – project lead is member of the OSG Council

Q1: cnt'd

- Area Coordinators take on the responsibility for particular interactions as part of the annual planning:
  - Technology/Software Area Coordinator is a member of the WLCG Management Board, attends many/most of the technical workshops, and oversees a regular – monthly – phone meeting on technical items of interoperation across between WLCG/EMI/EGI and OSG
  - Security Officer is a leading/active member of the WLCG security working groups, the XSEDE security operations group, and in ongoing contact with the CILOGON and IDM infrastructure lead at NCSA (Jim Basney - previously on OSG staff).
  - Operations Area Coordinator attends the regular WLCG Operations meetings, semi-annual EGI technical and user forums and has ongoing communication with peers in WLCG, EGI, EMI and XSEDE. He works to insure interoperability of VOs that span the EGI/OSG infrastructures. There are many example VOs for this, like Auger, ENMR, GEANT4, CTA, and XENON.
  - The OSG Operations lead is also on the External Advisory Board for EGI and work with the MoBrain (includes ENMR interoperability) and ELIXIR (interested in Galaxy submission to grid resources) EGI Competence Centers to provide interoperability for international partners. The actual work here is simply interoperability, which we are already doing, and sharing ideas for technology to service the Bioinformatics and Medical communities.
- Members of the Executive Team are well regarded and invited regularly to Grid and other Distributed High Throughput Computing conferences, workshops and visits throughout the world including GridPP in the UK, China, South and Central America etc. Members of the Council have helped coordinate a Grid day in a physics school in Africa for the past few years and members of the OSG Staff have helped teach

Open Science Grid

Q2: Charge Element I

- OSG has multiple components designed to address the needs of several communities and collaborative projects. Are the program's priorities and allocation of resources reasonable in meeting those needs? Has the program been successful in meeting diverse requirements?
  - Yes, Yes, and Yes
    - Answer to Q3 and 6 outline how we are addressing the LHC needs
    - Adding the IF needs following stakeholder requests and discussion with DOE helped to unify some of the technical approaches and principles
    - the goal to become a XD service provider was part of the proposal, and this work allowed us to put in place a dependable OSG Open resource addressing the needs of a large number of stakeholders from campuses
    - the new OSG Connect service addresses the cost of entry issues and eases access to OSG enabling more sites to become part of OSG
    - New functionalities like adding the OSG-CA and network monitoring, again driven by stakeholder requests, further help to address stakeholder needs.

Q2: Charge Element II

- II. Evaluate the degree to which OSG has been successful in meeting its proposed goals in its first two years.
  - ★ OSG has been exceedingly successful, in meeting the growing and changing computing needs of the U.S. LHC program, and in facilitating the increasing demand for DHTC.
  - ★ For evidence w/r to the LHC goals see answer to Q3.
  - ★ Among the evidence for success w/r to the larger community we offer the following:
    - We provided usage metrics for the OSG Open Facility that shows diversity of science, diversity of scale from single PI to sizable communities, and significant scale of use outside the LHC community. The impact of OSG is broadly acknowledged by citations to our main research paper in major peer reviewed journals across about a dozen scientific disciplines.
    - Established a “software attic”, and cleaned up this attic leading to an overall reduction of the software stack by 20%. This benefits the entire OSG Consortium, including the OSG project as it reduced the complexity of the software stack we maintain.
    - Turning the “overlay job manager” (glideinWMS) into an OSG service to support more than a dozen communities, including Intensity and Cosmic Frontiers.
    - Adoption and delivery of the DOE CA to mitigate the closedown of the DOE Grids CA
    - etc — there are many ore examples

Q2: Charge Element III

- Evaluate the goals and program of work for years 4 and 5, including their appropriateness and value to the communities served.
- We organize the workplan along three focus areas, highlighted below.
  - Provide value to VOs and resource owners by enabling DHTC and resource sharing for their set of applications:
    - This includes continuing to the support the large VOs
    - In addition to sustaining services, extension of LHC VOs for Run 2.
    - Complete the transition to new HTCondor-CE.
    - Continue to reduce the complexity of the software stack by another 25% (clean the software attic) resulting in 40% reduction. The value includes a simplified stack with better support, and making it easier to install.
    - Continue strong program of work in stable operations
  - Provide opportunities for the Long Tail of Science:
    - Full usability of OASIS, simplifying data management.
    - Serve new campus communities through OSG-Connect.
    - Ramp-up intensity frontier experiments.
    - Decrease the need for user-unfriendly X509 certificates.
    - Research and apply understanding of our opportunistic resources.
  - Promote use of DHTC for science:
    - Run the OSG User Schools.
    - Publish OSG research highlights monthly.
    - Evolve our ability to do flexible resource provisioning.

Q2: Charge Element IV

- IV. Evaluate the level of funding, a) in the light of work done during years 1 and 2; and b) with respect to the program of work for future years.
- program of work was sized according to the level of funding we succeeded to secure
  - At start of project the project went through a process to down-scale the deliverables and scope due to not receiving the full funding requested
- sized our promises to stakeholders/ expectations of stakeholders to the budget
  - feel that we delivered what we promised, in terms of the services, sized to the budget
- we are not aware of any deliverables that LHC expected and that we are not delivering on
- to address the broader goals form NSF and in particular ACI, we invented/ developed a new set of services with OSG Direct and OSG Connect
- these efforts put OSG on a very successful path to support science through DHTC, used and explored broadly throughout DOE and NSF science

Q2: Charge Element V

- Discuss the appropriateness and effectiveness of plans for collaboration, personnel, and management organization, and for meeting demands on computing and data processing, including, for all, flexibility in response to unanticipated risks.
  - The OSG organization has proven to be rather effective
  - The organization setup of having a consortium of stakeholders with a council driving the goals and overseeing the project, and a set of by-laws have allowed the project to deliver on the highest priority items and provides regular and immediate feedback for eventual corrective actions
  - The project has demonstrated that we are able to address unforeseen needs
    - through stakeholder requests, examples is e.g. the glideinWMS services etc
  - project was repeatedly re-organized and the project structure adapted to changing needs and prioritys, e.g. by merging the technology and software areas, pulling out the release process making it deliver to the operations area, adding areas like networking, adding and managing new sub-projects like the OSG-CA project
  - we have been successful to re-assign effort and funding between institutions, to re-prioritize tasks and assign manpower appropriately, across areas, to add effort to respond to changes, mitigate risks et
  - We believe we have an effective organizational structure both in the consortium and in the project to remain agile and responsive

Q2: Charge Element VI

- VI. Evaluate OSG in the context of similar and possibly competitive efforts worldwide
- Federation of Autonomous Cyberinfrastructures
  - One of OSG's fundamental principles is that of Federation of CIs. The peer European projects initially aimed to be a single homogeneous system. They subsequently evolved into the model pioneered by OSG: EGI to a federation based on National Grid Infrastructures (NGIs) ; the LCG to the Worldwide LCG made up of regional infrastructures. We are different from the other NGIs in being community organized and driven, rather than top-down managed provisioning.
- Driven and Controlled by the Stakeholders
  - OSG is stakeholder driven and does not develop software. This has allowed us to take appropriate decisions on the technology components based on the user needs;
  - An example is how OSG evolved technology smoothly to the job overlay paradigm while maintaining the operating production infrastructure without the disruptive and competitive transitions the CreamCE, and the gLite WMS experienced by EGI/ EMI. The WLCG has now also transitioned to the job overlay paradigm.

Q2: Charge Element VI cnt'd

- Partnering with peer CIs worldwide
  - OSG partners with local CIs in other countries to provide DHTC experience and knowledge, support services in the short-medium term and work to enable local autonomy, provisioning and support in the longer term.
  - For example, GridUNESP in Brazil initially used OSG central services such as the CA to demonstrate locally that DHTC can provide utility to their local researchers, and have gradually transitioned to supporting their own CA and other operations services, while maintaining an interoperable software stack that enables them to provide resources accessible to both OSG and EGI.
- Opportunistic Open Facility
  - US HPC centers support users based on allocations which are presented and reviewed. OSG is unique in being an infrastructure that spans DOE and NSF resources and providing "opportunistic" access to otherwise unused resources. This has provided 90Million CPU hours in the past year, which at 5¢ an hour it $4.5M of value this year alone. Our goal is to increase this in the future.
- Use Local Resources as well as Remote Resources
  - OSG promotes the use of local resources as part of a transparent cyberinfrastructure. This promotes the distribution and transfer of knowledge of creating, maintaining and using DHTC, naturally encouraging the next scientific and technical generation.

Q3: LHC

- What are the measures of success in meeting the needs of ATLAS and CMS? How have you met them?
  - OSG provides a wide range of services to ATLAS and CMS, each of which have metrics that OSG has either met or exceeded expectations.
  - Rather than making a complete list, we provide a flavor below for the three most important categories.
  - (i) provide a secure global single sign-on environment in collaboration with international partners.
- This includes software, processes, training to personnel in its deployment and operations, days-to-day security operations etc.
- Evidence of success are:
  - the existence of the global production infrastructure that OSG is part of.
  - the successful security audits & drills
  - the lack of security incidents

Q3: LHC cnt'd

- (ii) scalability of core software and instances of network facing services
  - Evidence of success are:
    - the work on scalability that we have documented in project reports and publications since OSG inception.
      - Over the years the experiments have defined the scale of a variety of software & services, and OSG then worked with its satellites and software providers on measuring the scalability, and ultimately improving it. This was and continues to be an iterative process that is triggered at times by the experiments, and at times by internal OSG project considerations, or software provider request. The former is in response to concerns about future scalability while the latter is in response to new releases of crucial pieces of software.
    - the successful operations of the experiments on OSG at the full Run 1 scale
- (iii) routine operations of the OSG facility
  - Evidence of success are:
    - the very fact that we have a detailed set of SLAs most of which were developed in collaboration with ATLAS and CMS.
    - the fact that these SLAs have been tracked, and goals have been met or exceeded.
    - public statements of success of "the Grid" from senior ATLAS, CMS, and CERN management.
      - E.g. the clear and unambiguous credit the CERN D.G. Rolf Heuer has given the global Grid infrastructure as being an essential contribution to the discovery of the Higgs Boson that lead to the Nobel Prize in Physics in 2013.

Q4: Trouble Tickets

Ticket Status	Number of Tickets	Certificate Requests	Other Requests	Percent Certificate Requests
Open	110	70	40	64%
Closed	174	127	47	73%
TOTAL	**284**	**197**	**87**	**69%**

- 4. Provide information on number, time-to-resolution, etc., of trouble tickets omitting cert requests and other “rote” topics.
  - Non-certificate tickets are a mix of:
    - user job problems,
    - resource issues reported from monitoring or staff observations,
    - installation or software issues reported by resource providers
  - Time to resolution is not tracked as a single number due to the vast difference in ticket types (simple questions about where a document might be found that take minutes to software bug that takes months to develop a patch)
  - We looked at the time to first response for tickets opened during the month of July, which was 196.6 minutes. This number includes tickets opened after hours and on weekends (including a holiday weekend in July). The vast majority of tickets submitted during business hours were touched in much less and an hour.

Example Ticket Topics

- Over the past 24 hours we have seen pilots running on the order of minutes and/or returning no logs, which to us in factory side suggests you might be having a black hole worker node.
- For the past couple of weeks, it appears glideins submitted to Puerto Rico have been having difficulties.
- I got errors from some jobs running on US.ANALY_SLAC. I attach a list with times as they appear in panda logs.
- In a recent software release, we've discovered that the gratia probe deletes certinfo files and treats all records as local.
- BDII people ask us to remove unnecessary publication. For CMS supporting sites with cvmfs, they can remove VO-cms-CMSSW* from $OSG_APP/etc/grid3-locations.txt.
- Could somebody schedule the Phedex load test from T1_US_FNAL to T2_US_Purdue ? We are interested in measuring the disk to disk transfer rate between these two sites in both the directions.
- As presented and discussed during the 2 latest WLCG Ops coordination meetings, during August CERN is going to upgrade all CVMFS installation servers to version 2.1. From the Experiment Dashboards it looks like your site is still running an old version of the client < 2.1.19. Could you please plan the upgrade of your installation?
- While doing some unrelated Gratia queries, I noticed that probe condor:fermigridosg1.fnal.gov mapped to SiteName "None". I assume this should be mapping to FNAL_FERMIGRID.
- Thank you for registering the NET2-BU Resource. By completing the resource registration form you have started the approval process to have your resource added to the OSG Resource Structure.
- File transfers are failing at MWT_UC both as a source and a destination with communication error.

Q5: Prioritization

- Describe the prioritization process of last June specifically. Step through the specific input from stakeholders, from OSG Council, points from the proposal, etc.; and the decision process to arrive at the year 4 plan. What requests were deferred or dropped?
  - The stakeholder request system key inputs were:
    - 1) Run 2 startup items (scaling & robustness)
    - 2) Various software maintenance items from LHC
    - 3) IF requests for GUMS, OASIS features
    - 4) GPU and access to diverse resource types
  - Key inputs from the Council were
    - 1) Run 2 preparedness
    - 2) IF support
    - 3) A more generalized provisioning system
    - 4) Support interfaces to HPC systems for stakeholders

Q5: Prioritization ctn'd

- These inputs plus from proposal were evaluated with respect to available resources and the priorities of stakeholders and we converged on priorities (and a portfolio) for year3 as follows:
  - 1. HTCondor-CE deployment to USLHC sites before data taking restarts (Run 2 startup - Robustness)
  - 2. Testing for operation of key systems at 2X scale (Run 2 startup - scaling)
  - 3. Software distribution via OASIS for IF and investigate read-only data access via HTTP (Council)
  - 4. Growing Campus access to DHTC; more opportunistic harvesting (ACI - Proposal)
  - 5. Progress on the network monitoring data store and support for PerfSonar packaging (Strategic Initiative & LHC)
  - 6. Reduce dependence on certs (ID management - proposal)
- Had to continue to defer work due to effort/expertise on
  - 1. HTPC deployment (low resource support)
  - 2. Job Level Monitoring items (only one VO requesting)
  - 3. Support OSG-level interfaces to HPC systems (viewed as part of Provisioning system)
  - 4. Generalized provisioning system – coordination needed to develop plan with glideinWMS and Panda
  - 5. GPU access (low demand)

Open Science Grid

Q6: LHC Data Challenges

- Please describe and quantify the current data challenge and how it will address readiness for LHC Run 2. - Brian
- OSG and the LHC community have several work items to prepare for LHC Run 2. Through the OSG Council, the OSG project has received the following explicit requests:
  - (a) "Identify scalability limits for OSG services in preparation for LHC Run 2". We are currently preparing for to do a “scale test” of 200k running jobs using glideinWMS to a small number of test clusters. This will cover the scalability of the OSG CE (used by both CMS and ATLAS) as well as the glideinWMS/HTCondor stack itself (used by CMS). For our computing software, we believe this will either verify the systems are ready for Run 2 or identify a set of issues we will need to work on.
  - While the Run 2 data rates are significantly higher than Run 1, our understanding is the additional load will be on the storage systems, not directly affecting the grid interfaces on top. However, we will closely watch the USCMS data challenges (explained below) to see if any issues develop over the next few months.
  - (b) "Request project list the software components and prioritize the list for LHC Run 2". This work is currently ongoing. We do not predict this work will result in any new software in the stack nor will we be unable to support any grid software user by LHC. Rather, the primary focus will be to:
    - With the remaining time before Run 2, retire software (concrete example: remove redundant SRM clients).
    - Commission new versions of existing software and required capabilities which will (concrete example: validate new versions of HTCondor, validate IPv6 support).

Q6: LHC Data Challenges ctn'd

- The computing and physics organizations in the LHC are currently running internal data challenges. ATLAS's program of work is Data Challenge 2014 (DC14); CMS's is Computing, Software, and Analysis 2014 (CSA14). The CMS challenge currently focuses on commissioning new analyses and computing systems for Run 2 but not new scales for grid software.
- For 2014, the USCMS Tier-2 program includes several explicit data scalability goals. These include:
  - Expose storage and compute services via IPv6.
  - Demonstrate the ability to support 4,000 simultaneous off-site clients for remote I/O (using the Xrootd protocol).
  - Demonstrate the ability to transfer data offsite at greater than 20Gbps.
- The Atlas Data Challenge is focusing on the new analysis model, the derivation framework, analysis trains, etc - managed by the new production system JEDI
- The OSG pays close attention to this activity through the weekly Production meetings; the software team is committed to fix any IPv6 issues with OSG-supported software, but any work items from this effort have already been fixed.
- Finally, when LHC is closer to Run 2 startup, the currently twice per week WLCG operations meeting (with attendance by OSG Operations staff) will increase tempo to a daily meeting (Currently, the WLCG is not organizing the individual experiment's data challenges.) There is also a twice per month Operations Coordination meeting attended by Ops where non-immediate operations issues are discussed

Q7: Milestones

- Some year 1 and year 2 milestones are shown as delayed with no information about eventually being addressed. How have the milestones evolved? How are the milestones "re-baselined"? Since there is unallocated money, couldn't that be used to meet the delayed milestones? - Gabriele, Chander
  - Many of the delayed goals have been incorporated into the year 4-5 plan or are believed no longer to be priorities as standalone items; the major categories are
    - HTPC goals –> several sites and OSG services have demonstrated the HTPC capability, VOs are currently not demanding it at the scale of the milestone
    - Security/ID services –> planned as part of next major step associated with 2-year DigiCert contract expiry in June 2016; host certificates and certificates for gridFTP will need attention
    - Cloud Resources –> handled by ATLAS and CMS at their level making use of OSG infrastructure; OSG plans to include in year4 Resource provisioning plan
    - Job Level monitoring –> outsourced to pilot overlay projects (glideinWMS)
    - Dynamic Resource Allocation –> year4 Resource Provisioning plan
  - We baseline each year's work program which gives priority to stakeholder current needs (with high priority to US LHC operational needs).
  - Some of these goals are affected by partner schedules and priorities and we have limited control. Others require in-depth knowledge of OSG and new staff is unlikely to accelerate these milestones.