


State of the Open Science Grid

LATBauerdick/Fermilab





Management Team



◆ OSG council elected OSG Executive Director

- ★ transition from Ruth starting on April 1 w/ 2 month overlap period

◆ Executive Team


- ★ Lothar Bauerdick as Executive Director
- ★ Miron Livny (PI) as Technical Director
- ★ Dan Fraser as Production Coordinator
- ★ Chander Sehgal as Project Manager
- ★ Frank Würthwein as Resources Manager
- ★ Michael Ernst, Frank Würthwein as Applications Coordinators

◆ Area Coordinators

- ★ Production: Dan Fraser ---> Operations: Rob Quick, Campus: Dan Fraser
- ★ Security: Mine Altunay
- ★ Software: Alain Roy
- ★ Technology: Brian Bockelman
- ★ User Support: Chander Sehgal



OSG continuing for another 5 years!




Open Science Grid

- ◆ Besides focus on physics and the huge momentum of the LHC, there is a fascinatingly broad spectrum of different science applications making great use of the OSG — see Frank Würthwein's talk
- ◆ We successfully made the case for continuing the OSG
- ◆ Very strong support from funding agencies to continue
- ◆ Continuation proposal was very favorably reviewed
- ◆ Now we need to deliver on our promises!




OSG Project Continuation



Open Science Grid

- ◆ OSG project will continue to receive funding across DOE and NSF for 5 years starting ~1 April 2012, with ~\$5.4M/year
 - ◆ DOE HEP \$1.6M, DOE NP at \$50k, NSF OCI \$1.0 M, NSF PHY \$2.75M
 - ◆ Funding was 25% lower than requested and lower than 2006-2011 levels
 - ★ NSF largest contribution to Universities, both from physics and from OCI
 - ★ DOE funding to labs in particular LHC related efforts BNL, Fermilab
 - ◆ Can only succeed in the vision with continued contributions from stakeholders and success of satellite projects — continue to engage pushing satellites!
 - ★ Will have a significant review before 3rd years beginning of 2015
- ◆ Very significant NSF OCI funding means stronger emphasis on Campus Infrastructures, partnerships with other NSF funded projects
 - ★ while keeping fully committed to continue the successful support of LHC
- ◆ OSG to expand with reduced overall effort — need to be very efficient!
 - ★ Recent LHC review demonstrated big focus on cost/benefits
 - ★ work with project manager and area coordinators improving accountability


Continue the successful OSG program!



- ◆ The model of OSG being a Consortium – with contributors, stakeholders, opportunity for loosely coupled satellites – and a Core Project – staff, deliverables – remains strong
- ◆ Continued focus on LHC – support for ATLAS, CMS, ALICE distributed computing in the US and contributions on behalf of US LHC to WLCG.
- ◆ Extend support to more stakeholders – communities (e.g. astrophysics and nuclear physics) and scientists local to the campuses:
 - ★ Partnership/Participation in NSF as an XSEDE Service Provider (see next slides)
 - ★ Emergence of NSF Institutes – centers of excellence on particular cyberinfrastructure topics across all science domains
 - ★ Evolution of DOE ASCR SciDAC program – collaboratories report submitted, nature of future solicitations unknown (came out yesterday?)



New opportunities with OSG and XSEDE working together



◆ OSG and XSEDE have very similar goals

★ XSEDE combines “high-end” services, in particular HPC centers

★ and works with OSG to give XSEDE users access to the OSG HTC services

★ effort within XSEDE to support users to run on OSG


Science requires diverse digital capabilities

- XSEDE is a comprehensive, expertly managed and evolving set of advanced heterogeneous high-end digital services, integrated into a general-purpose infrastructure.
- XSEDE is about increased user productivity
 - increased productivity leads to more science
 - increased productivity is sometimes the difference between a feasible project and an impractical one





OSG as XSEDE Service Provider




Open Science Grid

- ◆ As an XSEDE Service Provider support users given allocations by the (NSF supported) XSEDE request and allocation process
 - ★ OSG now part of the XSEDE resource allocation committee XRAC
- ◆ First users in summer 2012 through dedicated Submit interface.
 - ★ OSG-XSEDE Submit host ready for friendly user testing
 - ★ 2,000,000 CPU hours on OSG allocated, 3 XSEDE allocation requests re-directed toward OSG at last XRAC meeting
 - ◆ in the areas of HEP theory simulations and beam transport
 - ★ Probably need to negotiate acceptance of VO=osg at more sites and enable gfactory for those sites
- ◆ Expect this to go in full production and to support multi-stage workflows across OSG and other XSEDE Service Providers
- ◆ Transparency across OSG and the rest of the XSEDE sites should give OSG users added opportunity to gain access to more resources when needed.




Extending OSG across the Campuses



Open Science Grid

- ◆ Enable local cross-campus sharing of resources.
- ◆ Currently ~8 campus infrastructures – mostly LHC Tier-2s.
 - ★ US LHC Tier-3s are “beachheads” for their campus.
- ◆ Goal to increase access to these by more scientists and increase the number
- ◆ OSG should be well positioned to support emerging scientific revolutions, like in biology where e.g. every desktop will have a sequencer and locally 10s independent researchers each generate TBs of data/day for analysis, storage, retrieval – sounds like a great DHTC use case for OSG!





OSG services and infrastructure will fully support application parallelism to use many (hopefully all) cores on a compute-node (HTPC).

- ◆ Current status of HTPC – beta testing successful, as reported yesterday
 - ★ Production services need updating (information, testing, accounting)
 - ★ New releases of Condor with needed functionality are in development.
 - ★ Existing sites/resources on OSG need to be encouraged to support this use-case by having the infrastructure proven to work well and not negatively impact other work.


OSG services and infrastructure will support access to purchased and open Cloud resources. Which ones depend on user community requests.

- ◆ Current status of VM and Clouds
 - ★ Development work under the ExTENCI satellite to extend Condor s/w to support VM build, deployment, access.
 - ★ Production services need updating (information, testing, accounting)
 - ★ GlideinWMS releases support access to Amazon/EC2.

- ◆ The Job Overlay infrastructure (GlideinWMS) has been a tremendous success and boost to OSG usability, usage and effectiveness.
 - ★ ~50% of all cycles on OSG were used through glideinWM, and additional 25% through Atlas pilot-based WMS
- ◆ However, the current OSG services that support data – transport, access, archive, curation, provenance – across the end-to-end infrastructure are significantly deficient compared to the support for jobs!
 - ★ in my mind, this potentially will be a big obstacle to scientific productivity
- ◆ The large communities have their own implementations/deployments of Data Services, New and small communities lack these.
 - ★ the AAA concept is a huge step forward, potentially also for non-LHC users: w/ AAA scientists have their data accessible everywhere they want to run!
- ◆ We should work out a strategy
 - ★ the AAA project could show the way towards generalizable Data Access/ Remote I/O services for use by other OSG communities
 - ★ other Data Services – IRODS, Globus Online, Wide Area Lustre – will be usable on OSG and other resources to address particular needs.




Continue Federated Infrastructures



Open Science Grid

- ◆ Continue interfaces to peer infrastructures in Europe and elsewhere.
- ◆ In Europe collaborate with European Grid Infrastructure and its follow on, as well as national infrastructures in UK, France, Italy, Germany etc.
- ◆ Worldwide continue partnerships in South America, Asia and Africa – compatible with them also collaborating with Europe.





◆ The OSG project is organized in 5 Areas

★ Production: Dan Fraser


◆ Operations: Rob Quick, Campus: Dan Fraser

★ Security: Mine Altunay

★ Software: Alain Roy

★ Technology: Brian Bockelman

★ User Support: Chander Sehgal




Production Area

- ◆ See Dan Fraser's presentation later this morning!
 - ★ hugely promising progress with Campus Grids
 - ★ this opens new ways for scientists to use OSG
 - ◆ great opportunity for expanding OSG

- ★ OSG across Campuses is a major focus of this meeting!




OSG Operations Highlights



Open Science Grid

- ◆ Excellent operational performance
 - ★ Critical Services Availability BDII - 99.95%, MyOSG – 99.91%
 - ◆ just a few hours of downtime over the last year!
 - ★ Non-Critical Service Availability 99.87% Availability
- ◆ Strong support of OSG community
 - ★ ~150 Support Tickets per month
 - ★ No SLA Exceptions Reported
 - ★ Weekly OSG Operations Meetings and daily WLCG Operations Meetings
- ◆ New development like putting new services in operations
 - ★ like GlideIn Factories at IU and CERN, RPM repositories, OSG JIRA, the new XSEDE front end, Gratia Web etc etc
 - ★ first to develop a Balanced Scorecard with Assessment Activity
 - ★ also worked with WLCG developers to include test-environment for application-specific site availability tests, end-to-end RSV/SAM Testing



◆ Focus on Core OSG


- ★ Reaching a security maturity at core services
- ★ Reviewed and strengthened core infrastructure services (GOC, VDT)
- ★ Established policy, procedure and process to core security operations: software vulnerability assessment, security of middleware stack, incident response, cross-collaboration with XSEDE, EGI, WLCG

◆ Future Focus: Moving into campuses

- ★ Site and VO security:
 - ◆ Goal: joining OSG strengthens site security, not weakens
 - ◆ Makes campus Identities an integral part of OSG identities
- ★ OSG CA is the first step in transitioning ID landscape
- ★ Next steps: Streamline the process for obtaining, renewing, and using identities in OSG




Software Area



Open Science Grid


- ◆ See talk this afternoon by Alain Roy
 - ★ major upgrade to OSG 3.0 with move to RPM-based distribution
 - ★ much more “user friendly”
 - OSG installation can be done on 2 slides... :-)



- ◆ New area led by Brian Bockelman with two sub-areas:
 - ★ *Architecture*: Maintain and evolve the OSG Blueprint Document of architecture and principles, and run the OSG Blueprint meetings.
 - ★ *Investigations*: Strategically investigate individual technologies that may be potentially disruptive in the medium-term on the OSG
 - ◆ these are typically few months targeted “campaigns” to address specific issues
- ◆ Last year’s activities:
 - ★ Improving OSG batch system accounting
 - ★ Tackle some glexec usability issues
 - ★ Exposing VMs to VOs through OpenStack and Condor VM Universe
- ◆ This year:
 - ★ Data management with iRODS (collaboration with user support)
 - ★ “Cleanup” of info services (collab. with Operations)
 - ★ Follow-ups and continuation of 2011’s accounting and VM work
 - ★ engagement with WLCG TEGs: OSG can help with common solutions




User Support Area



Open Science Grid

- ◆ Main Focus is to help new communities use OSG effectively
- ◆ many examples of getting involved with new OSG user communities!
 - ★ SLAC Users communities
 - ◆ Additional sites for SuperB VO; Establish Ohio Supercomputer Center as an OSG site
 - ◆ SLAC Theory Group to run phenomenology codes on OSG as a proof-of-principle
 - ★ NEES — moving users from Engage incubator to NEES VO,
 - ◆ Oregon State University as an OSG site under NEES VO
 - ★ Electron Ion Collider — production completed
 - ★ SBGrid/eNMR jobs running on OSG, need to follow up
 - ★ DES — help adapt image pipeline to run on OSG
 - ★ LSST — moving users from Engage incubator to LSST VO
 - ★ coming up: Belle VO and PNNL integration into OSG
- ◆ Testimony of impact of User Support is newsletter articles from the new communities:
 - ★ EIC: <https://twiki.grid.iu.edu/bin/view/Management/NovDec2011>
 - ★ NEES: <https://twiki.grid.iu.edu/bin/view/Management/Oct2011Newsletter>
 - ★ GlueX: <https://twiki.grid.iu.edu/bin/view/Management/Sept2011Newsletter>
 - ★ OSG as XSEDE SP: <https://twiki.grid.iu.edu/bin/view/Management/Aug2011Newsletter>
 - ★ ALICE: <https://twiki.grid.iu.edu/bin/view/Management/Jun2011Newsletter>

Goal for each of us: bring in new OSG users through the Campus!



- ◆ Extending the science communities that can use and will profit from OSG will be a recurring theme in upcoming presentation and in the “lightning session” this afternoon!
- ◆ Let us play to the unique strength of OSG:
 - our presence in science and on the campuses and at the labs
 - ★ the triumphant success of DHTC ideas in the physics and other domains proves we have something outstanding to offer to science communities!
 - ★ only OSG has this broad presence on almost 100 campuses and the close scientific collaboration of domain and computing experts
 - ★ so, go out and find scientists on your campus with science that can benefit from Distributed High Throughput Computing!
- ◆ let us bring in new OSG users at **each of the OSG campuses**
 - ★ and let’s make sure that we continue to provide an excellent environment that allows great science to happen across the OSG



Open Science Grid

June 25–28, 2012

University of Wisconsin–Madison

2012 OSG USER SCHOOL

Harness the power of distributed computing

- ◆ lectures, discussions, and hands-on activities for new users of HTC and OSG
- ◆ **help spreading the word** and finding good candidates (mostly grad students)
- ◆ forward announcement emails and put up posters around your home institutions
- ◆ Students accepted to this program will receive financial support

More info + Financial support + Application (due March 30):

opensciencegrid.org/UserSchool

