

OSG

IRIS-HEP Advisory Board Meeting

Frank Würthwein
OSG Executive Director
Professor of Physics
UCSD/SDSC

September 9th 2019

What is OSG ?

- OSG is a consortium dedicated to the advancement of all of open science via the practice of Distributed High Throughput Computing, and the advancement of its state of the art.
- It is a collaboration between IT, software, and science organizations.
- It is governed by the OSG Council, maintaining its by-laws, and electing an executive director for 2 year renewable terms to coordinate a program of work.



Open Science Grid

OSG Effort Portfolio



A	B	C	D	E	F	G	H
	Σ	Technology	Operations	Support	Networking	Security	Management
Σ	17.56	7.46	4.05	2.82	1.24	1.00	0.98
OSG N5Y	6.86	2.48	2.05	1.82	0.00	0.20	0.32
IRIS-HEP	5.94	3.08	1.00	0.00	0.39	0.80	0.67
SAND	0.75	0.00	0.00	0.00	0.75	0.00	0.00
OSG-NP	1.00	0.50	0.50	0.00	0.00	0.00	0.00
TNRP	1.00	0.70	0.20	0.00	0.10	0.00	0.00
CESER	2.00	0.70	0.30	1.00	0.00	0.00	0.00

IRIS-HEP is roughly 1/3 of total effort in OSG

OSG Council Composition

Institution or VO	Representative Name	Affiliation
University of Nebraska - Holland Computing Center	David Swanson	OSG Council Chair, campus users, and resource provider
University of Florida	Paul Avery	Campus resource provider
Brookhaven National Laboratory	Eric Lancon	Campus users and resource provider
HTCondor Project	Miron Livny	Software provider
DOSAR	Horst Severini	Education
Fermi National Accelerator Laboratory	Elizabeth Sexton-Kennedy, Alternate: Burt Holzman	Campus users and resource provider
University of Chicago / Globus	Rob Gardner	Campus resource provider, users, and software provider
Indiana University	Craig Stewart, Alternate: Rob Quick	Campus users and resource provider

SBGRID	Piotr Sliz	User community
Solenoid Tracker at RHIC - STAR	Jerome Lauret	User community
Stanford Linear Accelerator Center	Stefan Hoeche	Campus users, resource provider, and theory users
University of California San Diego	Michael Norman, Alternate: Frank Würthwein	Campus users and resource provider
U of Pittsburgh Brain Trauma Research Center	Don Krieger	Research user
US CMS	Ken Bloom	User community
US ATLAS	Kaushik De	User community
IceCube	Kael Hanson	User community
LIGO	Peter Couvares	User community
Jefferson Lab	Amber Boehnlein	Campus users (GlueX) and resource provider
IRIS-HEP Project	Peter Elmer	User community

Pending member: Inder Monger, ESnet

Council selects its own members with the idea that this allows for maximum flexibility over time.

Most members represent organizations, and are as such “fungible”.

The People in OSG-LHC



Operations = UNL
Security = Indiana University
Software = U. Wisconsin – Madison
Networking = U. Michigan

A total of 6 FTE across 11 people.
 These people have worked together
 and with the LHC program for years.

For full team see: <https://opensciencegrid.org/about/team>

Key Goals in Plain English

- Do business as usual to **support US LHC Ops program** in software & computing
 - Accounting, CVMFS, Security, software releases, network performance collection ... the things the US LHC S&C leadership agreed to with each other.
- **Engage Intellectually** with the LHC software R&D program in the **US and at CERN**
 - OSG staff routinely attend US ATLAS & ATLAS meetings in US and at CERN, in person and via phone/video.
 - Engage with WLCG via technology, deployment, security, and management groups.
 - Working closely with DOMA R&D group in IRIS-HEP.
- **Bring new ideas & capabilities from R&D to operations**
 - Gridftp & GSI end of life replacement
 - Transition from person to capability authentication
 - Re-engineer OSG to provide DevOps deployment paradigm in addition to integrated software stack as RPMs.
 - Working closely with US ATLAS, SSL@IRIS-HEP, OSG outside LHC on containers as first class citizens for OSG services deployments.



Running the OSG

- Heavy use of slack, skype, phone, sms, zoom, vidyo, ...
 - We rely on people “knowing each other”, in many cases for many years.
- Each area has its own weekly meeting.
- Tim Cartwright runs a weekly staff meeting.
 - Area coordinators are scheduled for presentations.
 - policy & architecture discussions with short-medium term timelines happen here.
- Fkw runs a weekly Executive Team meeting.
- Marian Zvada runs a weekly Xcache meeting that includes the Xroot developers, and is open to people globally.
- And a few other weekly meetings.
- Quarterly meetings:
 - with Internet2 management
 - “Blueprint” to discuss long term high impact architecture and design issues.
- Annual planning meeting.
- Annual OSG All Hands Meeting

When we run OSG we are “funding origin agnostic” as much as possible.

Description	
Document and test integration of current LHC uses of XRootD in the OSG	Done
Design a process that allows site administrators to provide feedback on testing-grade software to expedite its release	Done
Develop web form for LHC site administrators to register service downtimes	Done
Complete the transition from Globus Toolkit to Grid Community Toolkit packages in EPEL and OSG repositories	Done
With DOMA, define a replacement path and schedule for GridFTP and GSI	Done
Coordinate with WLCG management to contribute as appropriate to the first LHCC review of WLCG in early 2019	“Done”

First 3 months

2nd 3 months



Open Science Grid

OSG PEP

Months 6-12



Complete the overhaul of the networking performance data pipeline, including tests of data restoration from tape and starting the transition of data collection from pull to push

Done

Align the OSG Cybersecurity program with the Open Science Cybersecurity Program Framework, the US-LHC Ops program, the US-LHC Tier-1s, and WLCG, and distribute responsibilities among teams to improve effectiveness and reduce unnecessary duplication of effort

Expanded scope thus completion delayed

Complete the OSG Operations transition by cataloging all OSG services (owned or operated) and updating OSG Service Level Agreements (SLAs) in collaboration with key stakeholders including the LHC experiments

Done

Summarize past year of the monthly accounting reporting to WLCG, including all issues with reporting and their resolutions

“Done”

Release OSG 3.5 with major additions to and deletions from OSG software for OSG-LHC

Done

Implement the process that allows LHC site administrators to provide feedback on testing-grade software to expedite their release into production

Done

Evaluate use of data federation software and configuration by US ATLAS and US CMS for common use

Done

My Worries

- Impedance mismatches between various different entities
 - e.g. US-LHC ops programs, global experiments, WLCG, OSG outside LHC.
 - This is going to make it hard to exploit more disruptive opportunities and/or stay on schedule with more disruptive goals.
 - Joint caching software
 - devOps
 - Globus transition
 - Analysis systems deployment
 - The first AS prototype I know that is multi-user was assembled by a grad student in their spare time.
- IRIS-HEP budgeting was done in a hurry and too tight.
 - reporting requirements were not budgeted.
 - Salary escalation not included appropriately thus people are not retained on budget.
 - In Y3 we did not budget management for OSG in IRIS-HEP to make 5 year budget work.
 - Many places in IRIS-HEP depend on "synergies" with other projects to achieve goals.

- We are on track.
- The IRIS-HEP part of OSG benefits from the various other efforts across the other 5 contributing NSF awards.
 - A strong team overall that is coming together.
 - Successfully avoiding impedance losses at boundaries between awards.
- Overall, I think we are doing more new things than I expected ...
- ... and are having a larger impact on the global LHC community than I expected.