Open Science Grid

UCSD

OSG-LHC Deep-dive at IRIS-HEP weekly EB meeting

Frank Würthwein
OSG Executive Director
Professor of Physics
UCSD/SDSC

June 3rd 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.

OSG Effort Portfolio

A	B	C	D	E	F	G	H
	∑	Technology	Operations	Support	Networking	Security	Management
∑	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

This is both a challenge and an opportunity:
Challenge = different projects add demands on shared services.
Opportunity = different projects drive the same deliverables.

A visual of OSG's Ecosystem

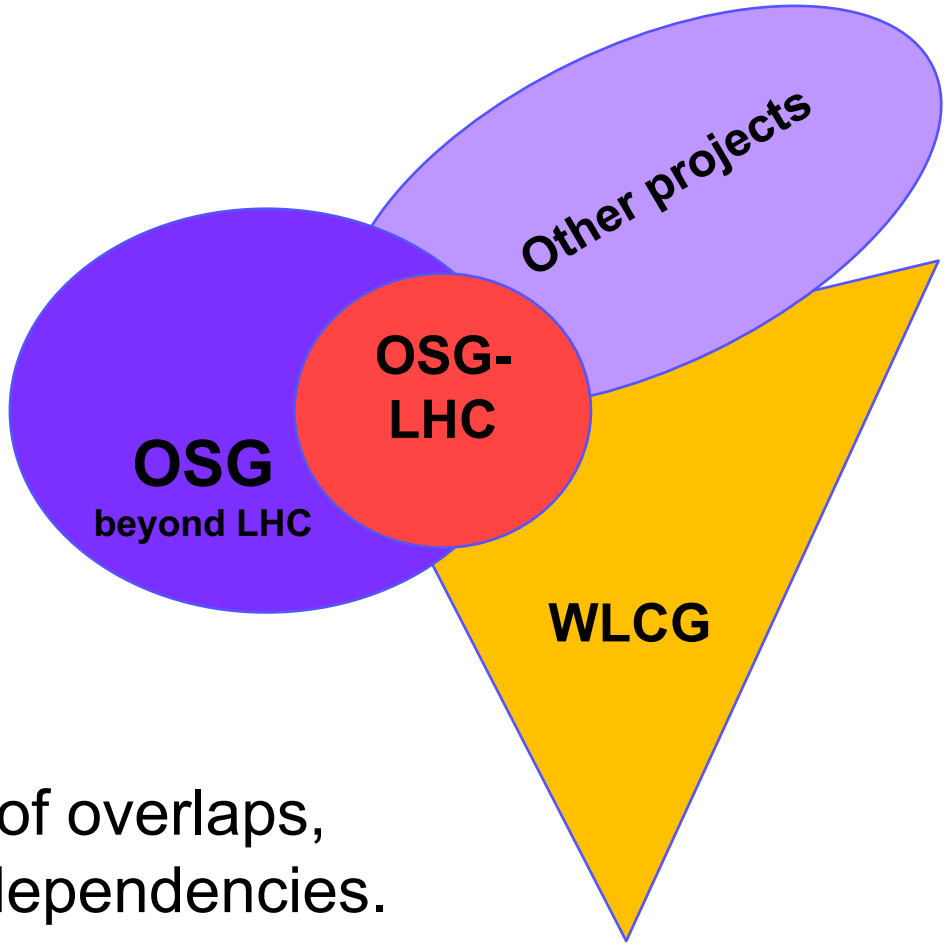

A complicated set of overlaps, collaborations, and dependencies.

Oversimplified List of Objectives

- OSG N5Y
  - Focused on individuals, small groups, and campus research support organizations
- SAND
  - Focused on networking and network performance
- OSG-NP
  - Focused on NERSC integration and NP collaboration building
- TNRP
  - Focused on containerization, devOps, and the networking engagement
  - Measuring & debugging network connectivity to/from commercial cloud
- ”CESER”
  - Focused on LIGO, IceCube, and data caches in the Internet2 backbone.

Other Projects

- SciToken
  - The transition from person to capability authentication.
- SLATE
  - Towards a federated devOps model of Edge Services.
- Several Security related projects
  - Effort at the T1s that we integrate with
  - Other NSF funded projects

The People in OSG-LHC


Deputy Director

Tim Cartwright


Security

Zalak Shah

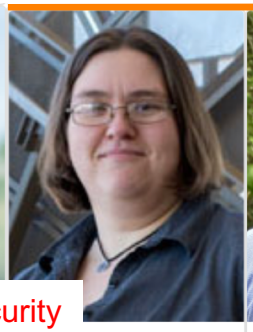

Susan Sons


operations

John Thiltges

Marian Zvada

software

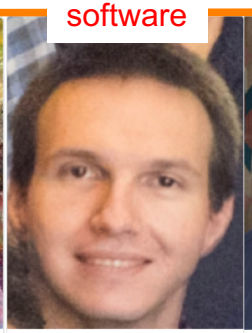

Mátyás (Mat) Selmeci


releases

Tim Theisen


networking

Shawn McKee


software

Brian Lin

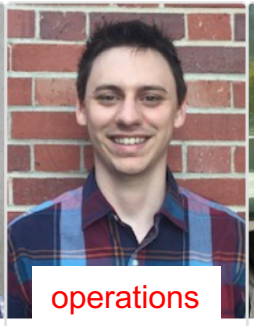
operations

Derek Weitzel

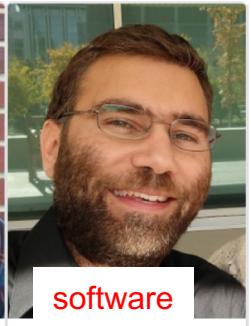
software

Carl Edquist

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

Open Science Grid

UCSD

Walk Through of Annual Report

Intro

The Open Science Grid for the Large Hadron Collider (OSG-LHC) area is organized into the subareas Operations, Software and Release, Security, and Networking. IRIS-HEP provides funding for effort in the area at Indiana University, University of Michigan, University of Nebraska–Lincoln, and University of Wisconsin–Madison (UW–Madison). Area leads of OSG-LHC are Würthwein (UCSD) and Cartwright (UW–Madison), the former in his capacity of Executive Director of the broader OSG collaboration of projects. OSG-LHC has been fully staffed from day one of IRIS-HEP. It comprises roughly a third of the total OSG effort in year 1 of IRIS-HEP, and collaborates very actively with the broader OSG ecosystem of funded projects.

Operations

- The OSG Operations team is responsible for deploying, configuring, and running the OSG-owned services that contribute to the overall OSG fabric of services. Also, the team provides operational advice on services that are run by others and that contribute to the OSG. The U.S. ATLAS and U.S. CMS projects depend on OSG services to connect sites, to provide site metadata, to operate data repositories, and to account for site usage, among other things.
- During the 9-month reporting period, the OSG helped provide just over one billion hours of computing time to the U.S. LHC projects. Further, there were no major OSG service outages during this time. Thus, the OSG continued to provide critical and reliable production infrastructure for U.S. LHC computing. Production operations do not simply happen automatically — it takes a dedicated and focused team continuously monitoring the health of a wide variety of systems, identifying and fixing problems that arise, in addition to provisioning new systems and adapting to the changes in the overall technology landscape. On top of those ongoing activities, the team also worked on several projects as described below.

Software & release (I)

- The OSG Software and Release team produces integrated software systems that comprise the OSG fabric of production services. The software comes from a variety of sources: from other teams in the Institute (i.e., DOMA); from outside of the OSG-LHC, unchanged, or with OSG-LHC patches or packaging contributions for better integration with production services; or in cases where we identify small gaps in the software stack, written in-house. In addition, OSG-LHC maintains responsibility for infrastructure software between the time the original developers discontinue support, and the time an orderly replacement strategy can be developed and executed. This is often referred to as the **"software orphanage service".** At present, an open-source fork of the Globus Toolkit is maintained by OSG-LHC as part of the orphanage.
- Over half of the team's effort goes toward routine software updates, tests of the integrated software stack, support of site and service issues, and maintenance of documentation.

Software & Release (II)

- In general, the team remains flexible to accommodate changing needs of the U.S. LHC experiments. For example, to more closely coordinate with U.S. ATLAS as well as ATLAS globally, the OSG Software team manager attended the U.S. ATLAS facility meeting at Argonne National Lab (with the OSG Deputy Executive Director joining for one day, as well) and the ATLAS Software and Computing Week at CERN (with the OSG Executive Director joining for one day as well). Participation in these events is in addition to attending the regular U.S. ATLAS facility phone conferences, which has been ongoing for quite some time.

Software & Release (III)

Metric	Y1Q1	Y1Q2	Y1Q3
7.1. Packages updated	32	37	56
7.2. Defects discovered	1	1	2
7.7. Tickets opened	34 software 22 packaging 11 documentation 39 support	50 software 22 packaging 15 documentation 24 support	44 software 19 packaging 14 documentation 40 support (6 CMS, 8 ATLAS)
7.7. Tickets closed	30 software 27 packaging 5 documentation 33 support	42 software 10 packaging 26 documentation 29 support	54 software 26 packaging 20 documentation 41 support (8 CMS, 5 ATLAS)

Security

- Cybersecurity is crucial to distributed computing; without trust, the entire fabric of services breaks apart, and without monitoring and verification, vulnerabilities can more easily turn into actual security incidents. The OSG Security team works on supporting certificate distributions, software security, incident response, and coordination with OSG's stakeholders in cybersecurity.
- During the entire reporting period, OSG had no operational security incidents.
- The team continuously monitored a variety of third-party sources for **announcements of software security vulnerabilities** (in, e.g., common operating systems, common software, and OSG-specific software). If there is an announcement that is found to affect OSG infrastructure, the team sends their own version of the announcement to a broad list of stakeholders, including service and site security contacts. Such announcements include an analysis of the vulnerability that is specific to the OSG architecture, configuration, and use of the affected software, plus known mitigation options, thereby making it easier for OSG service and site operators act on the vulnerabilities.

Networking

- Networks are fundamental to the operation and use of the LHC computing infrastructure. The OSG Networking team defines, develops, operates, and promotes the use of a suite of network monitoring tools so that networks can be treated as a managed component of the OSG. As a result, OSG has become the source of critical network metrics for OSG sites and some of our partners (like WLCG). Further, the team has designed, built, and deployed a variety of monitoring and visualization tools on top of this data, so that they and others can watch for, troubleshoot, and suggest fixes for network problems that slow or prevent production work.
- As with the other OSG-LHC areas, a significant part of the team's effort went toward routine activities; in this case, supporting and optimizing the production network monitoring infrastructure, promoting the use and maintenance of site monitoring tools, monitoring the collected data, operating the prototype alarm notification framework and target alerts on problems at sites, and troubleshooting specific network issues.

Networking (II)

Highlights from routine work and smaller projects include:

- Augmented the data collection pipeline to better locate the sources of networking problems (the Autonomous System number of routers); done in collaboration with the SAND project.
- Drove an upgrade campaign, with WLCG and LHC, to have sites upgrade their perfSONAR toolkit instances to the newest release and ensure the underlying hardware is sufficient.
- Gave four presentations to diverse audiences, highlighting the usefulness of network monitoring for science throughput and laying out plans for upcoming work.
- Created a new website (https://toolkitinfo.opensciencegrid.org/toolkitinfo/) as an entry point to various network monitoring, documentation, and dashboard sites; done with the SAND project.

Forward looking statement

- In addition, we are starting to see the first small adjustments to goals and re-focusing of effort. To date, this relates primarily to software and security issues with respect to supporting a new software distribution, deployment, and operations model …
- …. then some history … followed by …
- this has led to a devOps model of deploying services in containers deployed via Kubernetes on other people's hardware. This is the mode OSG uses today for most of its XRootD-based Data Federation in support of projects including Laser Interferometer Gravitational-wave Observatory (LIGO), Deep Underground Neutrino Experiment (DUNE), MINERvA, Dark Energy Survey (DES), and NOvA, plus a number of individual scientists in bioinformatics and other sciences. It is how OSG manages XRootD caches in the network backbone of Internet2 and other collaborators, including a cache in Amsterdam in support of DUNE and LIGO, especially.

Forward looking (II)

- More recently, U.S. ATLAS asked OSG-LHC to support this deployment and operations model also for the U.S. LHC community. U.S. CMS is still significantly behind both U.S. ATLAS and OSG in this regard. However, we expect that both experiments in the US, and possibly worldwide, will increasingly adopt deployment and operations models along those lines, and thus expect that this will continue to be a focus of OSG-LHC in Year 2. In the meantime, more and more services used by U.S. LHC are made available, and used, by the broader OSG community via containers on Kubernetes. Most of the effort that drives this is from outside OSG-LHC, for example the Pacific Research Platform (PRP), the NSF-funded Toward a National Research Platform project, and SLATE. **OSG-LHC and thus IRIS-HEP have an opportunity to leverage work and experience from a much broader community.**
-

OSG Data Federation

6 Data Origins
9 Data Caches

FNAL: HEP experiments
U.Chicago: OSG community
Caltech: Public LIGO Data
UNL: Private LIGO Data
SDSC: Simons Foundation
NCSA: DES & NASA Earth Science

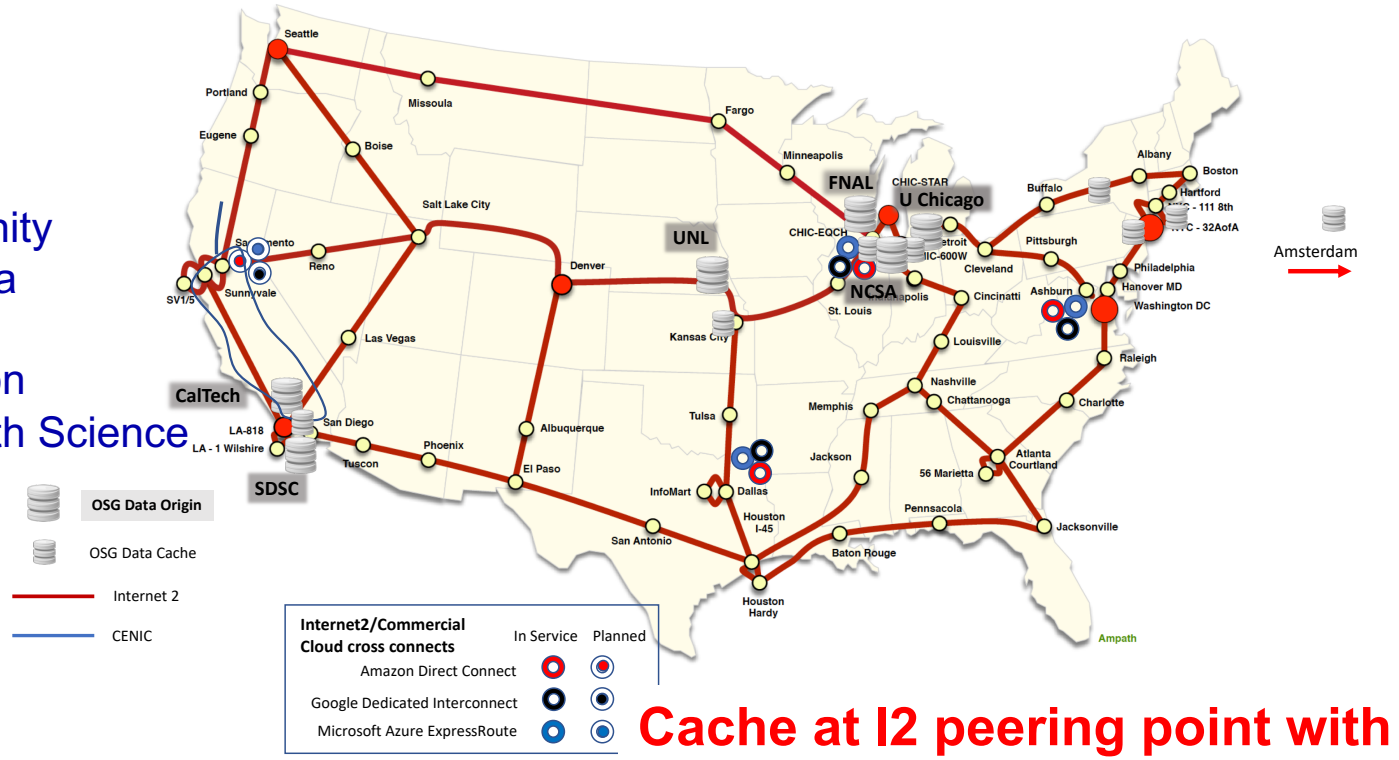

Directory	Working Set	Total Read
/pnfs/fnal.gov/usr/dune	13.107GB	395.537TB
/pnfs/fnal.gov/usr/minerva	255.266GB	270.994TB
/gwdata/O1	169.585GB	258.341TB
/pnfs/fnal.gov/usr/des	193.57GB	120.993TB
/user/ligo	5.612TB	83.564TB
/pnfs/fnal.gov/usr/nova	162.632GB	18.841TB

30k
1k
15
100

Depending on community, files were read 10-30,000 times during last 60 days.

Cache at I2 peering point with Cloud providers in Chicago

Summary

- We are on track.
- The IRIS-HEP part of OSG benefits from the various other efforts across the other 5 contributing NSF 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.