

Advancing Open Science through distributed High Throughput Computing

**Frank Würthwein
OSG Executive Director
Professor of Physics
UCSD/SDSC**



Goals for this Presentation

- **Discuss if there are areas where we can benefit each other by working more closely together?**
- My approach in this talk:
 - Provide you my standard intro to OSG so you see how OSG differs from the rest of WLCG.
 - Go through some ideas for more co-ordination/collaboration, and see if anything sticks.

Advancing Open Science with OSG

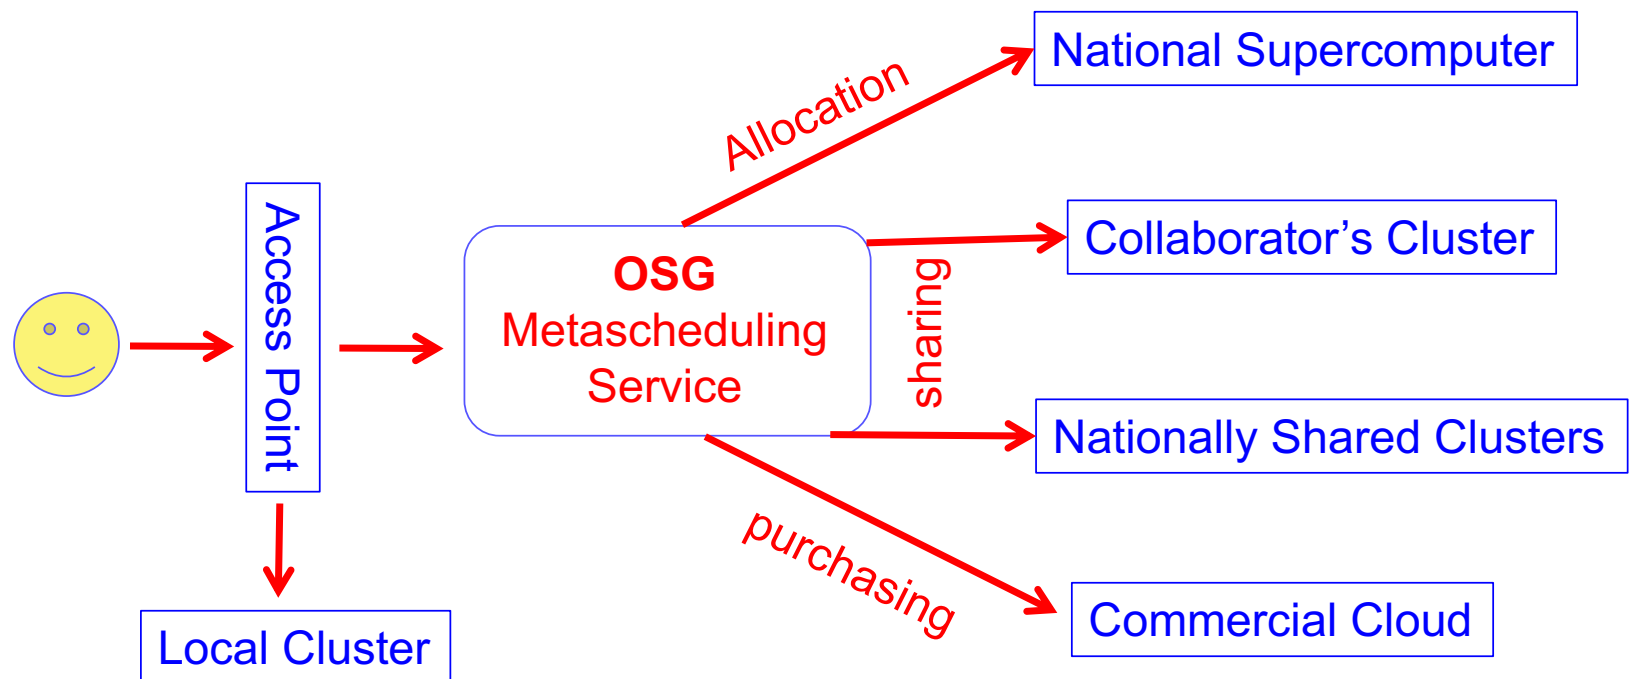
- All of open science irrespective of discipline
- Advance the maximum possible dynamic range of science, groups, and institutions
 - From **individual undergraduates** to international collaborations with thousands of members.
 - From **small colleges, museums, zoos**, to national scale centers of open science.
- Advancing this entire spectrum requires us to have a **diversified portfolio of services**

OSG serves 4 distinct groups

- The **individual researchers** and small groups on OSG-Connect
- The **campus Research Support Organizations**
 - Teach IT organizations & support services so they can integrate with OSG
 - Train the Trainers (to support their researchers)
- **Multi-institutional Science Teams**
 - XENON, GlueX, SPT, Simons, and many many more
 - Collaborations between multiple campuses
- The 4 **“big science”** projects:
 - US-ATLAS, US-CMS, LIGO, IceCube

~1/3 of the funded effort in OSG

- OSG-Connect, a submission host for individual researchers.
 - You get an account, and we teach you how to use OSG.
- A Compute Federation



OSG is complicated

- We have:
 - Sites that deploy the OSG-CE or for whom we operate an OSG-CE as a service.
 - VOs that we operate glideinWMS for.
 - VOs that we operate only the gfactory but not the gWMS Frontend for.
 - VOs that we operate HTCondor pool and gWMS for.
 - VOs that we operate HTCondor pool for, but not gWMS.
 - ... and the OSG VO for which we are the VO and do absolutely everything, including training our scientists.
 - Aside: the OSG VO does not have user certificates.

How many wall hours are provided by OSG is a very poorly defined question.

- OSG works on three simple principles:
 - Resource Owners determine policy of use
 - This means that all policy of use is set locally by the clusters that join the federation.
 - Resource Consumers specify the types of resources they are willing to use.
 - How much RAM? How many cores per node? ...
 - OSG submits HTCondor batch system as payload into all local batch systems that match requirements.
 - Jobs are submitted locally, queue centrally, and execute anywhere that matches requirements after resource becomes available.

OSG operates overlay system(s) as services for all of science



Open Science Grid

OSG Data Federation (XRoot)



6 Data Origins 12 Data Caches

FNAL: HEP experiments

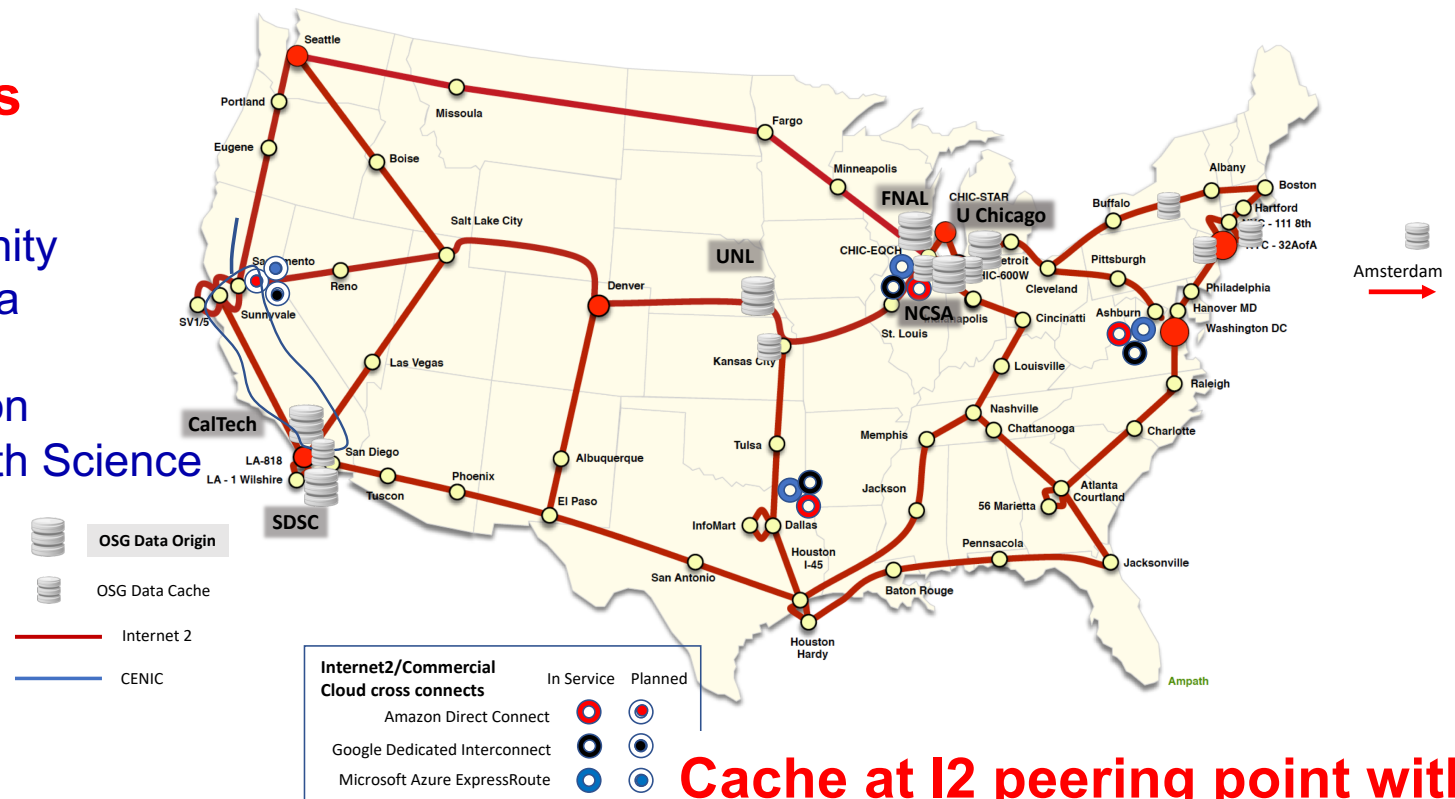
U.Chicago: OSG community

Caltech: Public LIGO Data

UNL: Private LIGO Data

SDSC: Simons Foundation

NCSA: DES & NASA Earth Science



Reads from Data Federation 9/1/2018-2019

Dune ~ 2.6PB

LIGO public ~ 1.5PB

LIGO private ~ 0.5PB

DES ~ 1.1PB

Minerva ~ 1.0PB

Cache at I2 peering point with Cloud providers in Chicago

Depending on community,
files were read 10-30,000 times
during typical 60 day period.

Data Federation Goals

- People come with their data on their storage systems.
- OSG offers to operate a Data Origin Service to export your data into the OSG Data Federation.
 - We give you a globally unique prefix for your filesystem namespace, and then export your namespace behind it.
 - We allow you to decide who can access what.
- OSG then strives to guarantee **"uniform" performance across the nation by operating caches** to:
 - Hide Access Latencies
 - Reduce unnecessary network traffic from data reuse
 - **Protect the data origins from overloads**

OSG operates overlay system(s) as services to all of science

Ideas for Topics to work more closely on together

Better Support US based Experiments globally

- LIGO/Virgo, IceCube, Dune, GlueX, clas12, ... all have collaborators outside the US that are providing resources to them.
 - Deal with the AuthZ issues better.
 - CILogon Basic issue mentioned before today/yesterday
 - CVMFS support is splintered between CERN, EGI, and OSG. We need to track who has what and from where.
 - Who gets their software from where?
 - Who gets their repo(s) from where?
 - Concerned about long term maintenance of two CVMFS configs, EGI and OSG. Can we converge onto one?
 - Submission host(s) outside the US.
 - E.g. discussion between Ligo/Virgo and NIKHEF
 - XCache shared across experiments.

What else?

- We have presently one OSG site that is entirely K8S ... we had to fix CVMFS to make that work. We contributed this back to CERN CVMFS team.
 - Unclear who provides long term support for CVMFS on K8S
- We have half a dozen XCache deployments in US, EU, and Asia on K8S.
 - Some in the network backbone at POPs, some at clusters, i.e. “edge services” at endpoints
- We want to deploy all of our services optionally via K8S.
 - Our strategy for “small site” problem is for us to host the services that sites would traditionally have to host.

Our **Aspirational** Goal:

A Research IT Organization should not have to learn anything "non-standard" in order to have their researchers benefit from OSG, or have their resources be available via OSG.

Well, we got some ways to go before we reach that goal ...

- We rely heavily on CVMFS, and would like to eliminate the need for worker node installs at sites.
 - OSG VO operates a module environment to produce uniform runtime environment for all of science.
- Why can't CVMFS be provided as a singularity container ?
 - Eliminate OSG specific worker node installations !!!



Summary

- We'd like to identify practical issues where we can work together more closely to benefit us all.
- The ideas presented are just fkw's first thoughts.
- Am open for additional or different ideas.

Open for Discussion