Advancing Open Science through distributed High Throughput Computing

Frank Würthwein
OSG Executive Director
Professor of Physics
UCSD/SDSC

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.

The Scope of Open Science

- 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, … many more
  - Collaborations between multiple campuses
- The 4 "big science" projects:
  - US-ATLAS, US-CMS, LIGO, IceCube

Global Collaborations

Services OSG Operates (I)

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

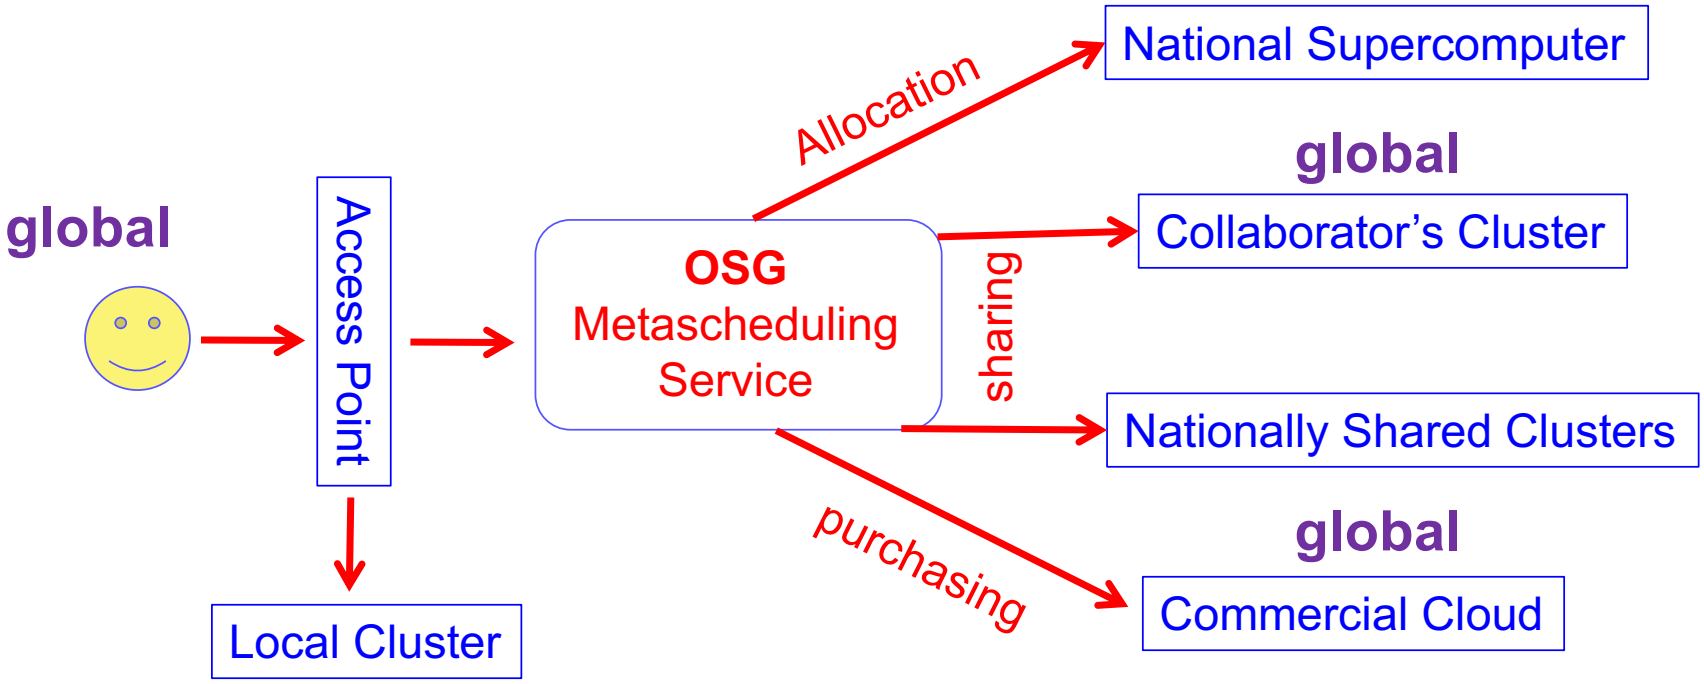

Federation = distributed control

- 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

OSG Data Federation

6 Data Origins
12 Data Caches

Caches deployed globally:
Amsterdam, Korea, Cardiff,
… more coming.

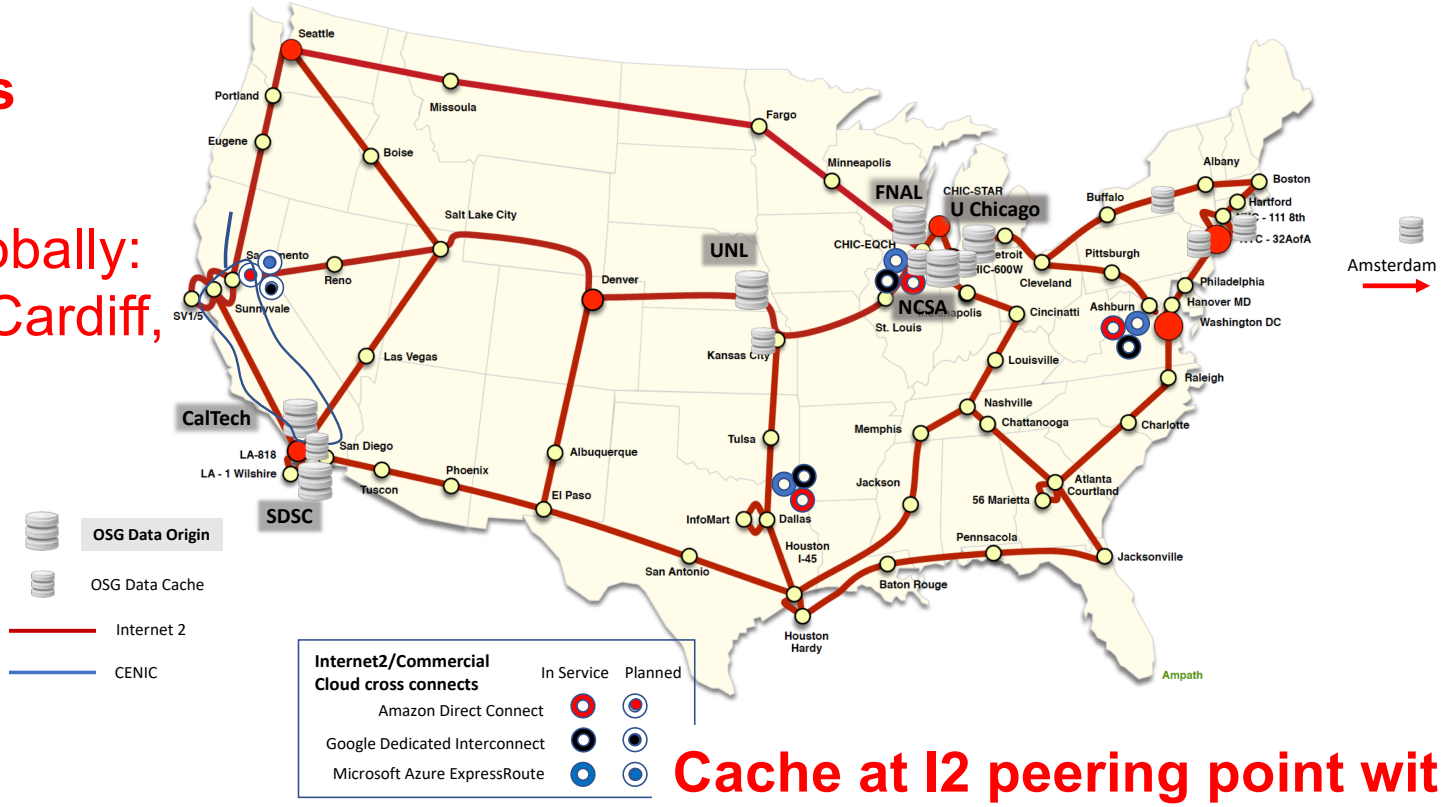

Cache at I2 peering point with Cloud providers in Chicago

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

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

Deployment and Operations

The PRP Vision

- Cheap storage is deployed all over the network
  - At end-points inside Science DMZs
  - At various peering points in the network
- Services can be dynamically deployed on top of the storage, e.g. via K8s.
  - A specific cache can be grown by adding storage.
    - Storage distributed “regionally” can be combined into logical cache.
    - What’s practical as region is determined by latency tolerance of the applications.
  - Additional logical caches can be added to the tree.
- The people who own the storage hardware need not know anything beyond container orchestration system, e.g. K8s.
- A given server may generally run one container to measure network performance, and a second container to provide the cache service.
  - The two containers may be managed by different organizations.

Long Term Vision

- Capacity Providers
  - Commercial cloud “competing” with on-premise
  - Different regions in the world will invest differently, and yet, capacity needs to be integrated globally.
- Service Providers
  - Software based services
  - Human based services (“consulting, training, …”)
  - “Content” providers
- Scientists organized at all scales
  - Individuals to 1000’s of collaborators

Managing Data across Origins

- The architecture supports multiple origins serving overlapping parts of the total namespace.
  - Origins in different parts of the world may provide replication for performance or redundancy or …
- There is an implicit notion that data across origins are possibly managed by a “data replication system” (e.g. globus online, Rucio, …) which uses bulk data movement engines like the ones described in the other talks in this session.

Summary

- Containerization, including container orchestration allows for:
  - New division of labor to support science globally
  - Lowers barriers to adoption of new software & services.
- OSG Data Federation is just the beginning of a variety of services that utilizes capacity globally
  - In cloud
  - On-premise
  - In the network