

# OSG

**Frank Würthwein**  
**OSG Executive Director**  
**UCSD/SDSC**

**March 4<sup>th</sup> 2021**

# OSG “Statement of Purpose”

OSG is a consortium dedicated to the advancement of all of open science via the practice of distributed High Throughput Computing (dHTC), and the advancement of its state of the art.

# To learn more about OSG

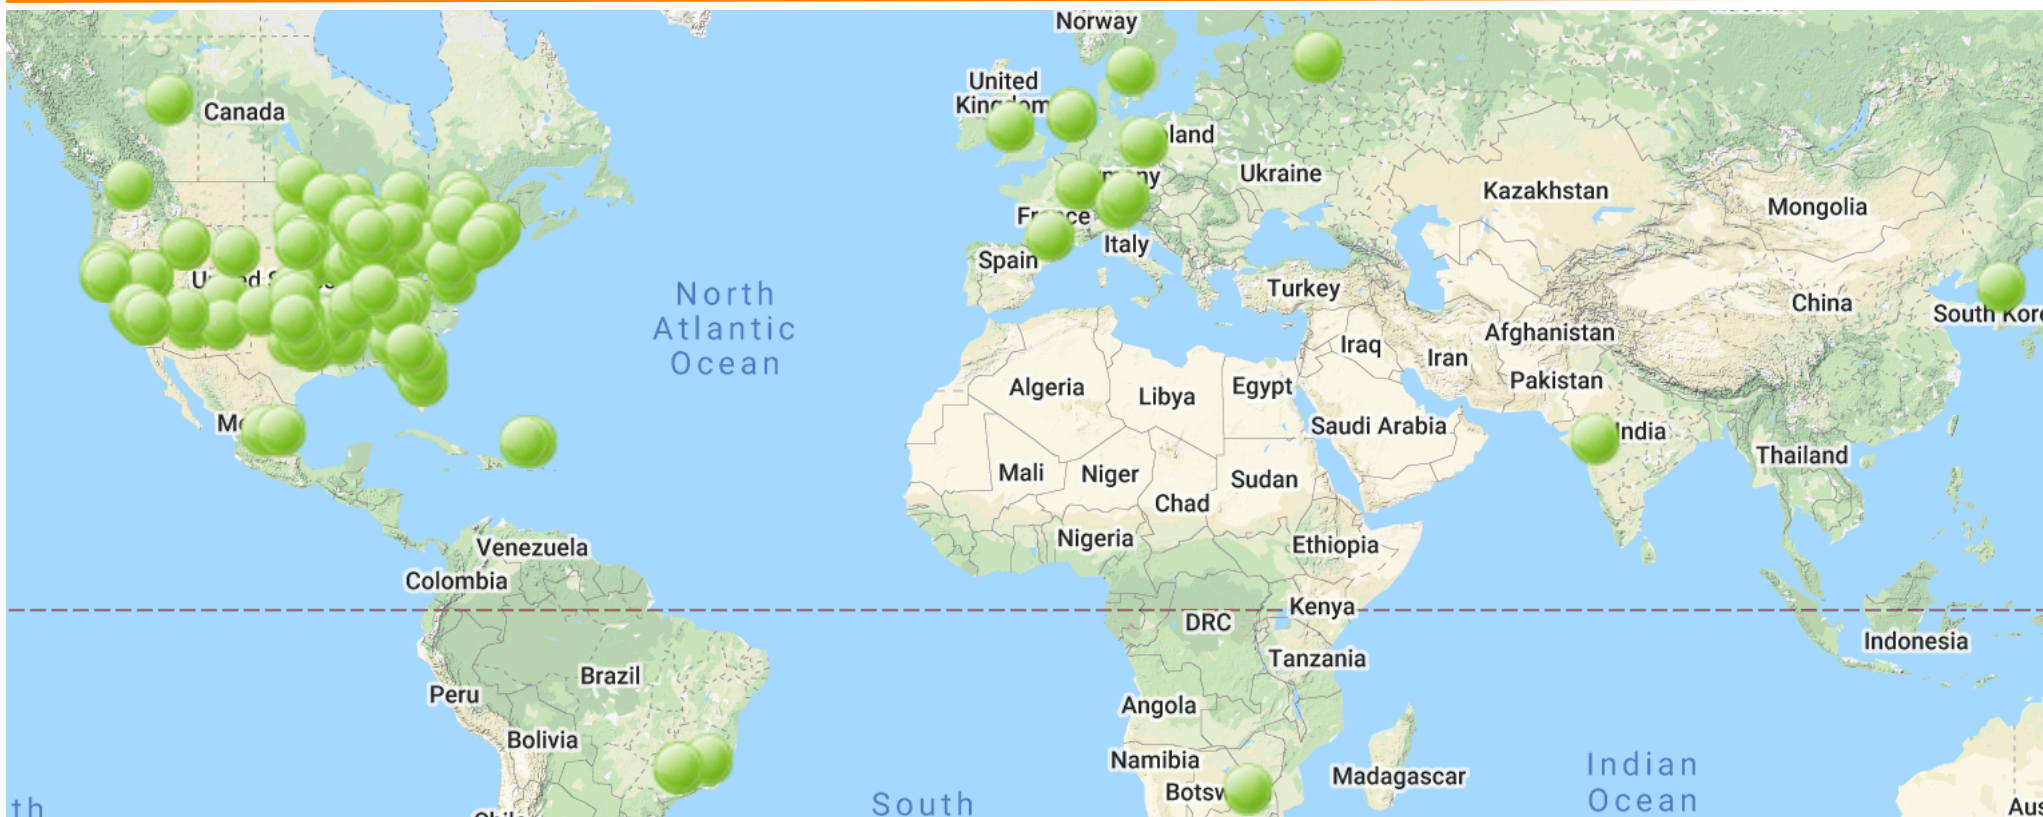
- We are having our annual meeting this week: OSG AHM 2021
  - Big Picture Overview:
    - State of OSG
    - OSG as an Agent for research computing democratization
    - OSG Services for Campuses
    - Technical Overview of Services
  - Detailed deep dive into Evolving Technology
  - Science Sessions
    - Individuals on Monday
    - Collaborations on Thursday Morning
    - LHC on Friday

**It is impossible to cover this in 15 minutes !!!**



Open Science Grid

# OSG Compute Federation



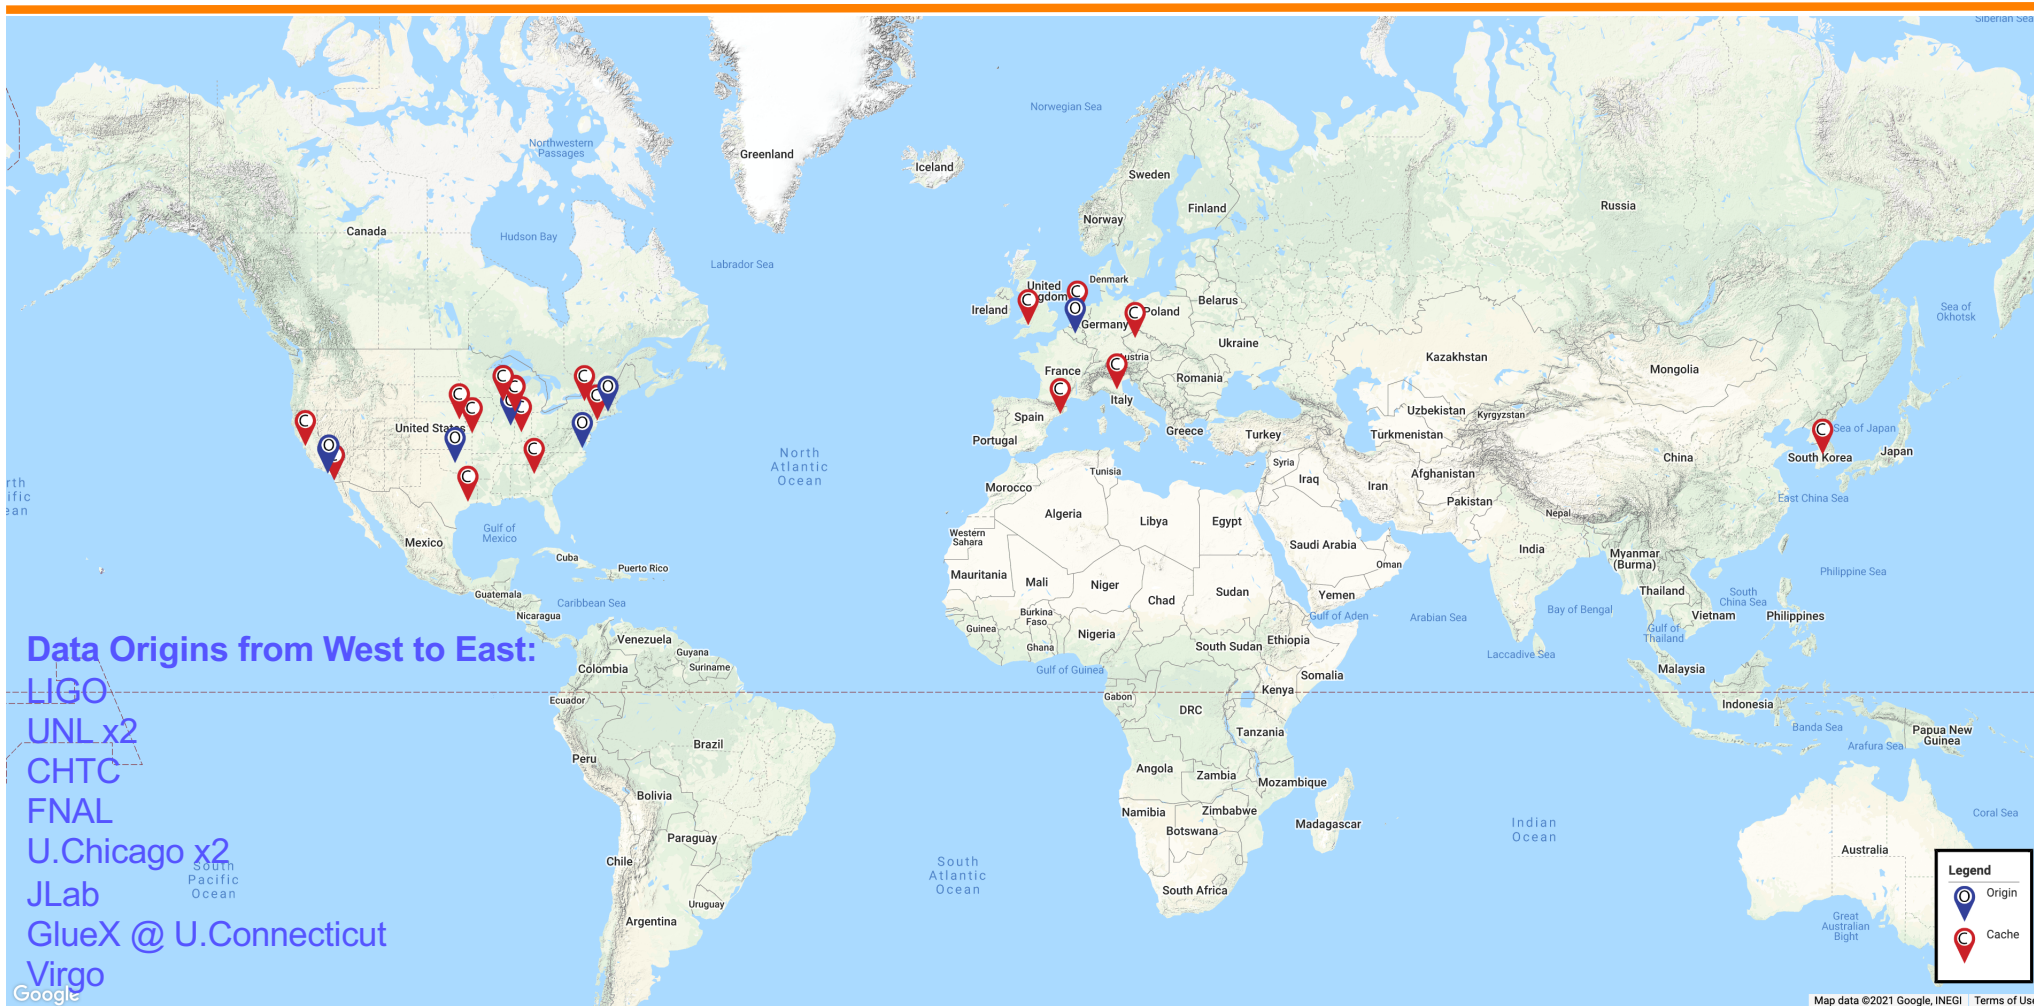
**136 clusters listed on this map**

... but the real story is so much more complicated ....



Open Science Grid

# OSG Data Federation



**17 Caches ... 6 of which are in R&E network backbone**

**10 Data Origins ... one of which is for all of open science**

# The “real story”

- OSG supports a **modular software stack** and a “**Fabric of Services**” that allows organizations to create their own dHTC environment.
  - dHTC environments from multiple organizations typically have non-trivial overlaps.
    - They may share resources
    - They may share services
    - They share data in a global namespace for public and private data
- **OSG operates one feature complete instance of such an environment for the common good of all of open science**
  - We support the common good by democratizing access
  - We eat our own dog food by operating software & services we support
  - We teach others to follow our lead to support their researchers and collaborations.



# A Feature-Complete dHTC Environment

- **Open Science Pool**

- Submission infrastructure that functions as compute “Access Point”
- Workload management system
  - Complex workflows across heterogeneous resources possible.
    - Easy to run workflows comprised of 100,000 jobs or more with complex dependencies between sets of jobs (full support of arbitrary DAGs).
- Homogeneous runtime environment across heterogeneous resources
  - Includes a dozen or more types of GPUs from NVIDIA and AMD
  - Includes some FPGAs
  - 100’s of application “modules” and 100’s of curated containers

- **Open Science Data**

- Storage that functions as “Data Entry Point”
- Transparent “Data Access” via Global Content Delivery Network
- Supporting Public and Private Data

**Any Researcher in the US, as well as their international collaborators,  
can request access to this dHTC environment**

# Joining the Compute Federation

- **We host the CE** for your cluster
  - Require an ssh-key accessible account on your cluster that can access the batch system.
- **You can deploy a CE** via either the EGI or OSG software stack
  - Actually, you can even deploy HTCondor-CE used in OSG and then integrate into the EGI software stack.
    - Done in several places in Europe these days
  - You can deploy the OSG CE via K8S or other container or cloud deployment mechanism.
- **We provide a container that you can deploy** to join our HTCondor pool(s)
  - The site decides when to invoke containers and how.

**We strive for maximal flexibility !!!**



- Built on **XrootD** products at the core.
  - But **integrated with CVMFS** for scalability of Metadata.
- We **offer to operate all services** if you provide us with a K8S environment for service deployment
  - Data Origin is read-only mount to our service
  - Data-cache is (obviously) a read/write mount to our service.
- Scale in 2020: **100++ projects with 200++TB working set and 20++ PB reads.**
  - Virgo & LIGO data analyzed together in this data federation despite the fact that origins of the data are in different continents, and compute distributed across 3 continents (4 soon)

- **Authz:** Moving entirely to token authz.
  - Should be interoperable with EGI
- **Runtime environment:** We support CVMFS.
  - Should be interoperable with EGI
- **Data:** Unclear to me what EGI is doing here
  - We are transitioning from gridFTP to HTTP(S), and that is fully integrated with RUCIO.
  - Unclear how our XrootD based Data Federation in OSG would integrate with EGI DataHub.
- **Accounting:** We send accounting records for the LHC to EGI accounting portal. Could do so for others as well.
  - We also collect metrics on perfSonar, XRootD, ...

**Interoperability exists today & we want to maintain/expand it.**

# References for OSG Data Federation

- Data Access for LIGO on OSG
  - [arXive version of published paper](#)
- Creating a content delivery network for general science on the internet backbone using Xcaches
  - [arXive version of published paper](#)
- The first paper includes discussion of CVMFS integration with XRootD.
- The second paper focuses on the K\*S deployment and operations.
- Additional details are in the references of these two papers.

# Questions ?

# Acknowledgements

- This work was partially supported by the NSF grants OAC-2030508, OAC-1841530, OAC-1836650, and MPS-1148698

