

# **Intro to OSG Caching work**

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



# Big Picture

- OSG has integrated multiple technologies to produce the “illusion” of a read only global filesystem that has high performance and is completely federated.
  - CVMFS for namespace management/authentication and caching of the namespace.
    - CVMFS uses Apache SQUID under the hood.
    - Accesses via CVMFS use GeoIP under the hood to choose the closest Xrootd cache for data access.
  - Xrootd for large volume data access via either HTTPS or XROOT protocols
    - Supports partial file access and caching
    - Supports engineered redundancy at various levels
      - Multiple JBODs can be combined into one cache
      - Multiple Origins can serve overlapping data
      - Origins and Caches can be configured into trees to separate regions
    - Supports caching of only parts of namespace
      - Caches can be deployed to support data from only certain communities.
    - Supports plug-ins for caching algorithms other than LRU

**Versatile building blocks that allow flexible deployments**

# Current Use in OSG

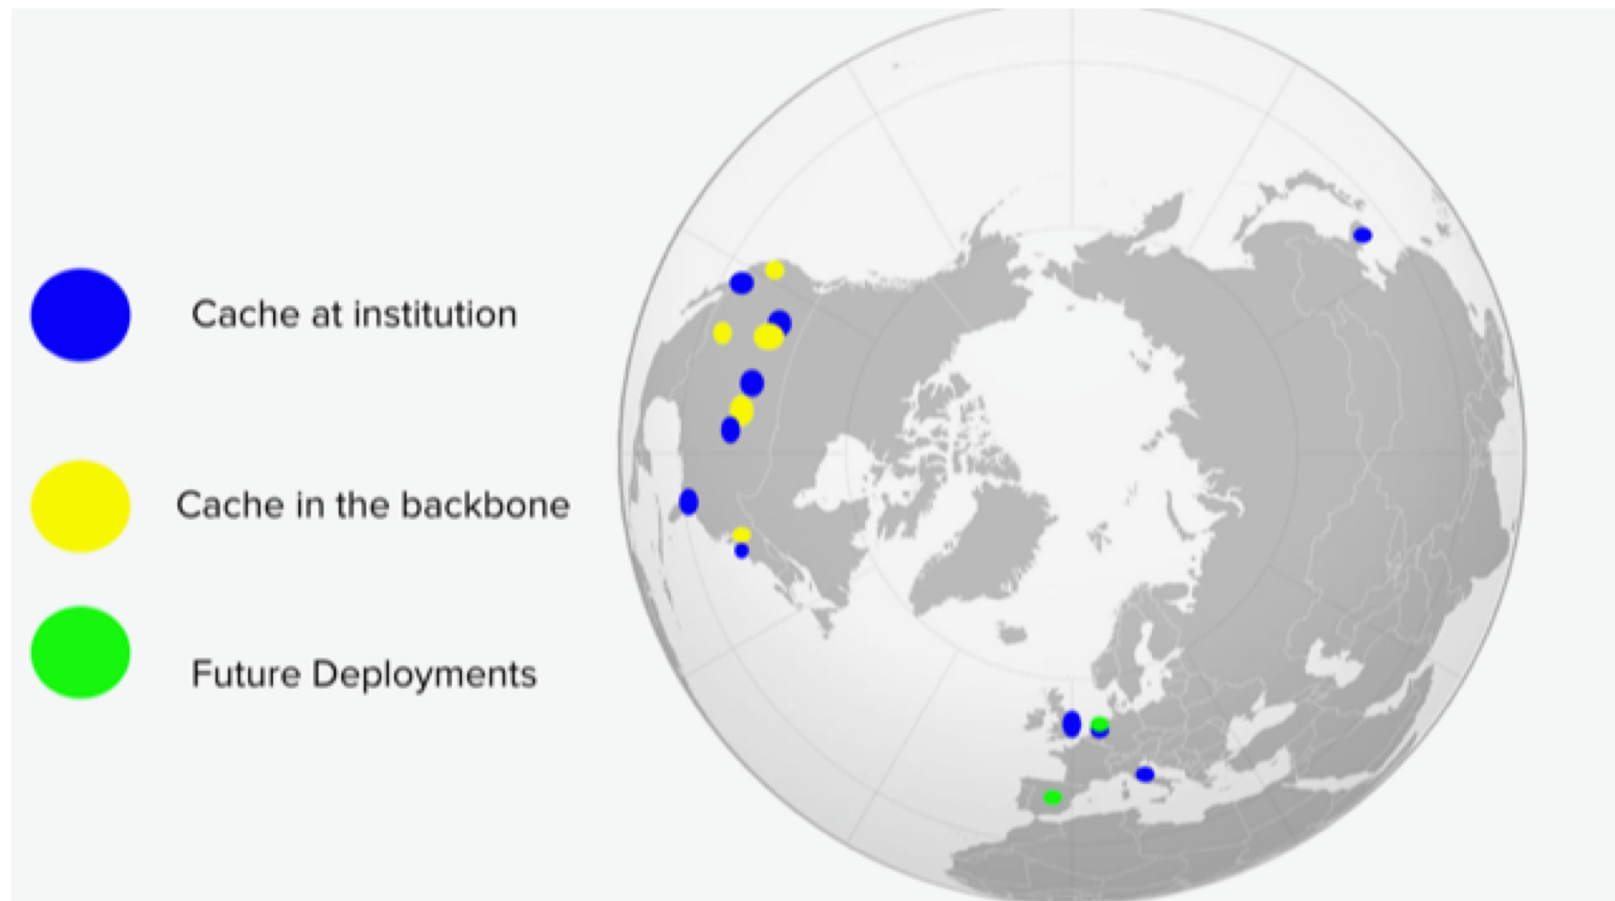
Collaboration	Working Set	Data Read	Reread Multiplier
DUNE	14.2GB	1.8PB	30k
LIGO (private)	18.2TB	596TB	30
LIGO (public)	7.2TB	96TB	14
MINERVA	435GB	305TB	700
DES	272GB	127TB	500
NOVA	4.8TB	8.5TB	2

Depending on the expected size of working sets, a given science use case may need larger caches than we have deployed.

**LIGO has driven international deployment, and DUNE is the major beneficiary of it.**

Typical cache sizes are large enough to fit the sum of all working sets at this point in time.

**More than a dozen caches  
deployed across 3 continents**



**Including several at POPs of Internet2 and GPN.**

# Largest Deployment

- The same federation technology is presently deployed by three different organizations: OSG, ATLAS, and CMS.
- ATLAS and CMS have typical working sets of:
  - Daily few tens of TB
  - Monthly few hundred TB
- CMS operates a 1PB cache distributed across Caltech and UCSD.
  - Two Compute centers access the same cache. File Servers in both are operated as one shared cache.

- We have a good starting point to build for new science use cases.
- The key technology we are lacking is fine grained authentication.
- The proposal we are talking about is meant to integrate that into what we have, to arrive at a more complete feature set, while at the same time gaining experience with a different science driver.