OSG Cache on Internet Backbone developments

Jan 16 2019

Edgar Fajardo on behalf of OSG and PRP

Introduction to Stashcache

- Caching infrastructure based on SLAC Xrootd server & Xrootd protocol.

- Cache servers are placed at several strategic cache locations across the OSG.

- Jobs utilize GeoIP to determine the nearest cache
- Job talks to the cache using HTTP(S) via CVMFS

Image taken from Brian's slides

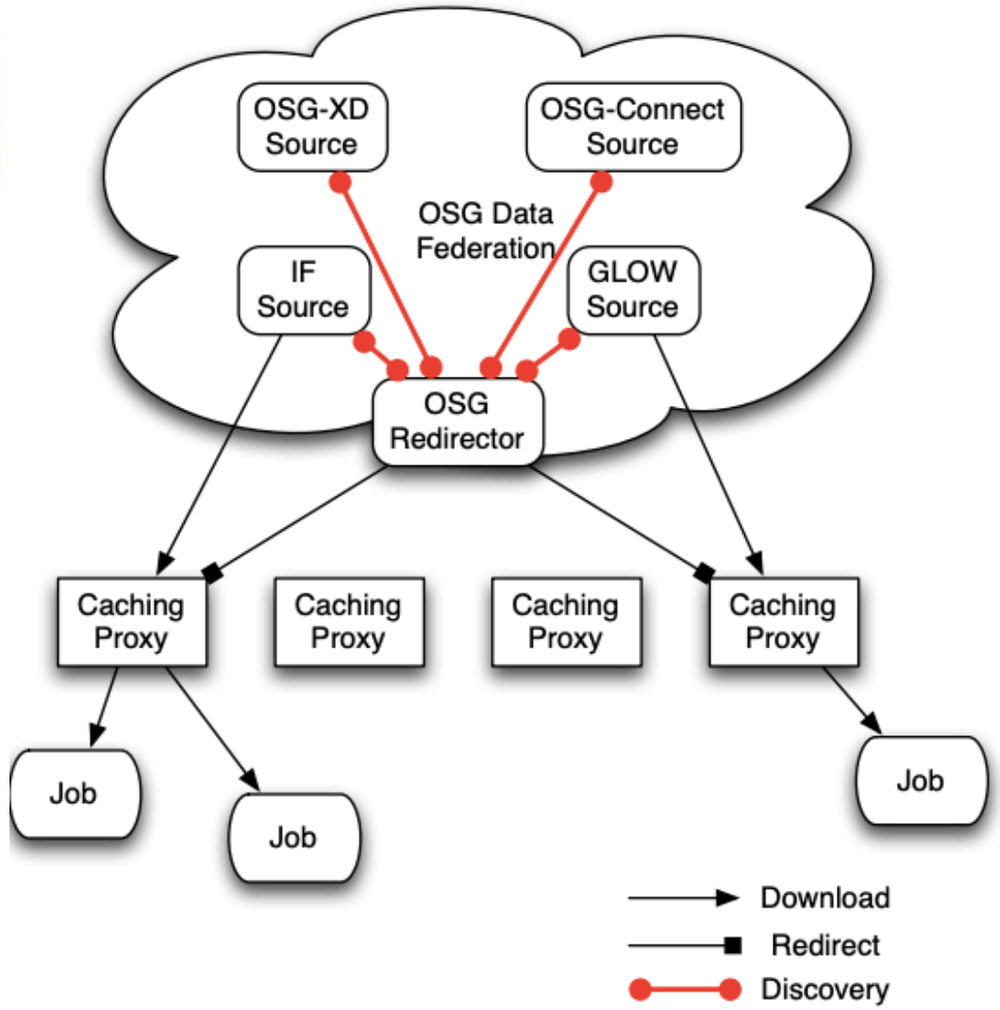

Implications

- An organization can join the federation with their own “data origin” and their own partition of the global namespace.
- A cache owner can decide on caching policies for different parts of namespace.
- This allows the owner to selectively serve only a subset of the community that use the federation.

We can build community specific “storage overlays” with this technology.

OSG Computing Resources

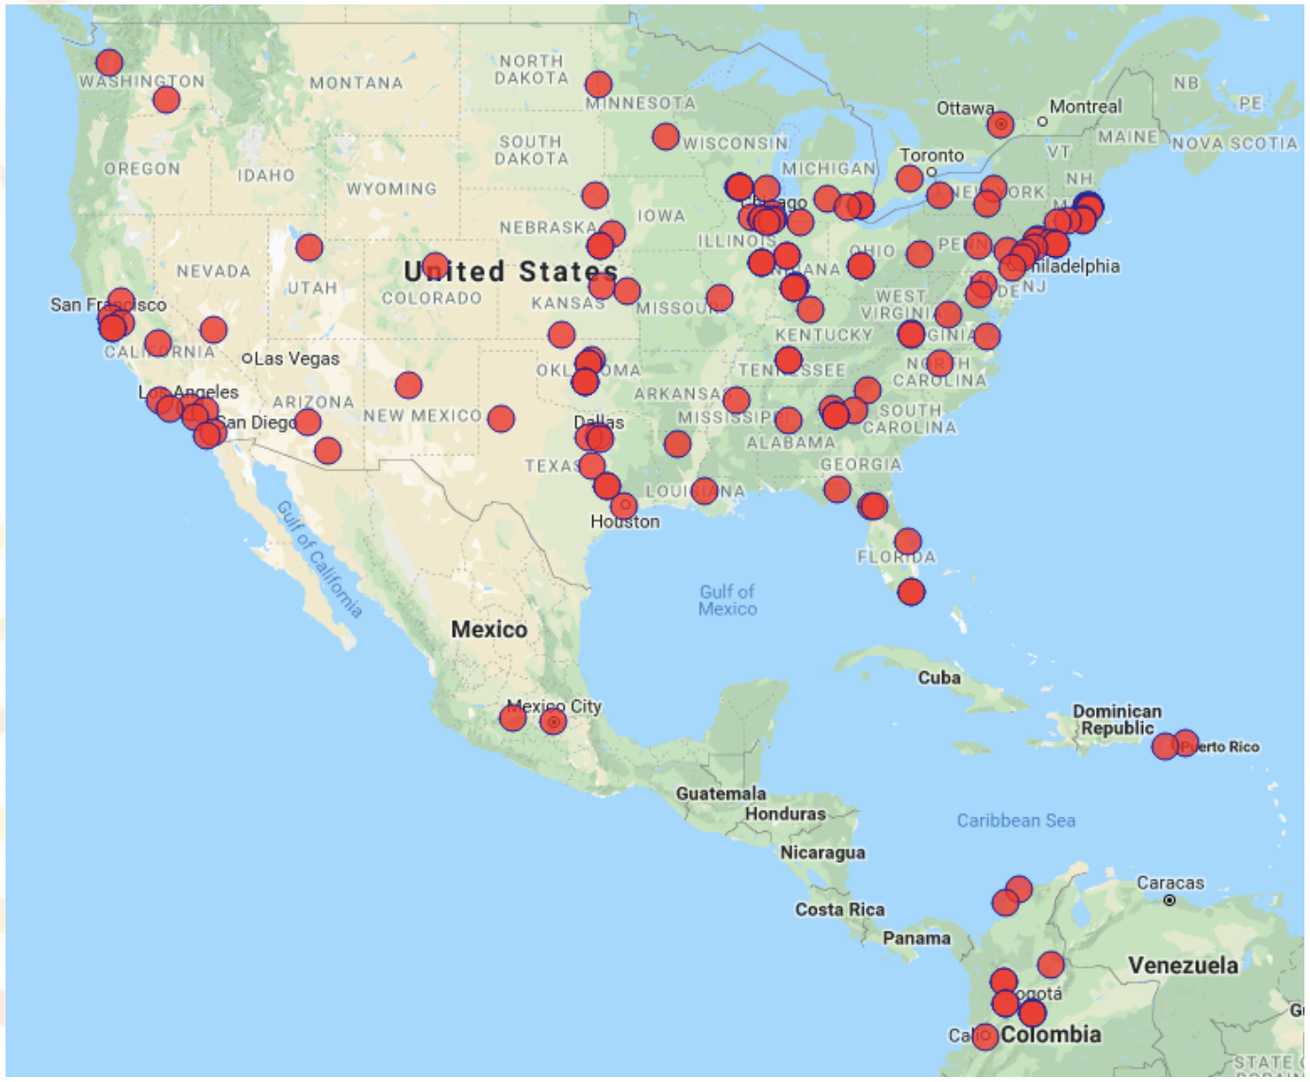

OSG Data Origins

FNAL: Fermilab based HEP experiments

U.Chicago: general OSG community

Caltech: Public LIGO Data Releases

SDSC: Simons Foundation (planned)

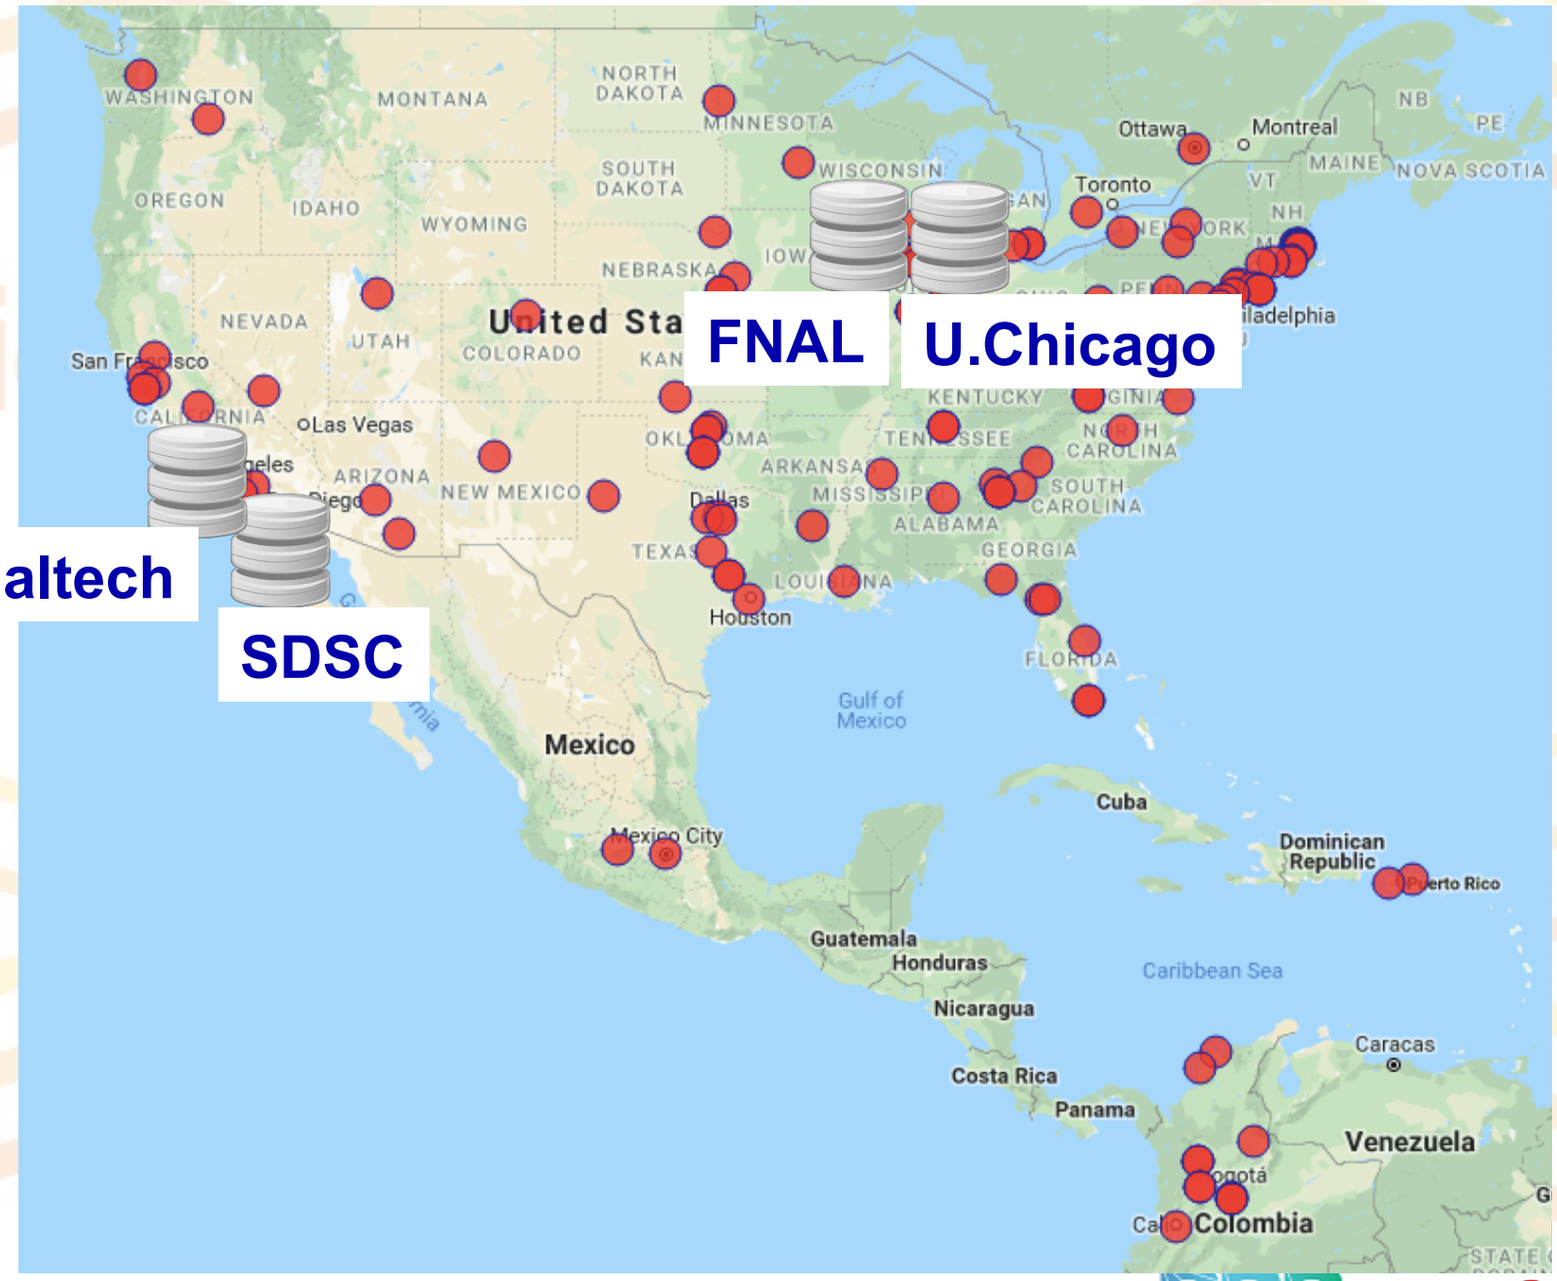

Network Backbone Caches

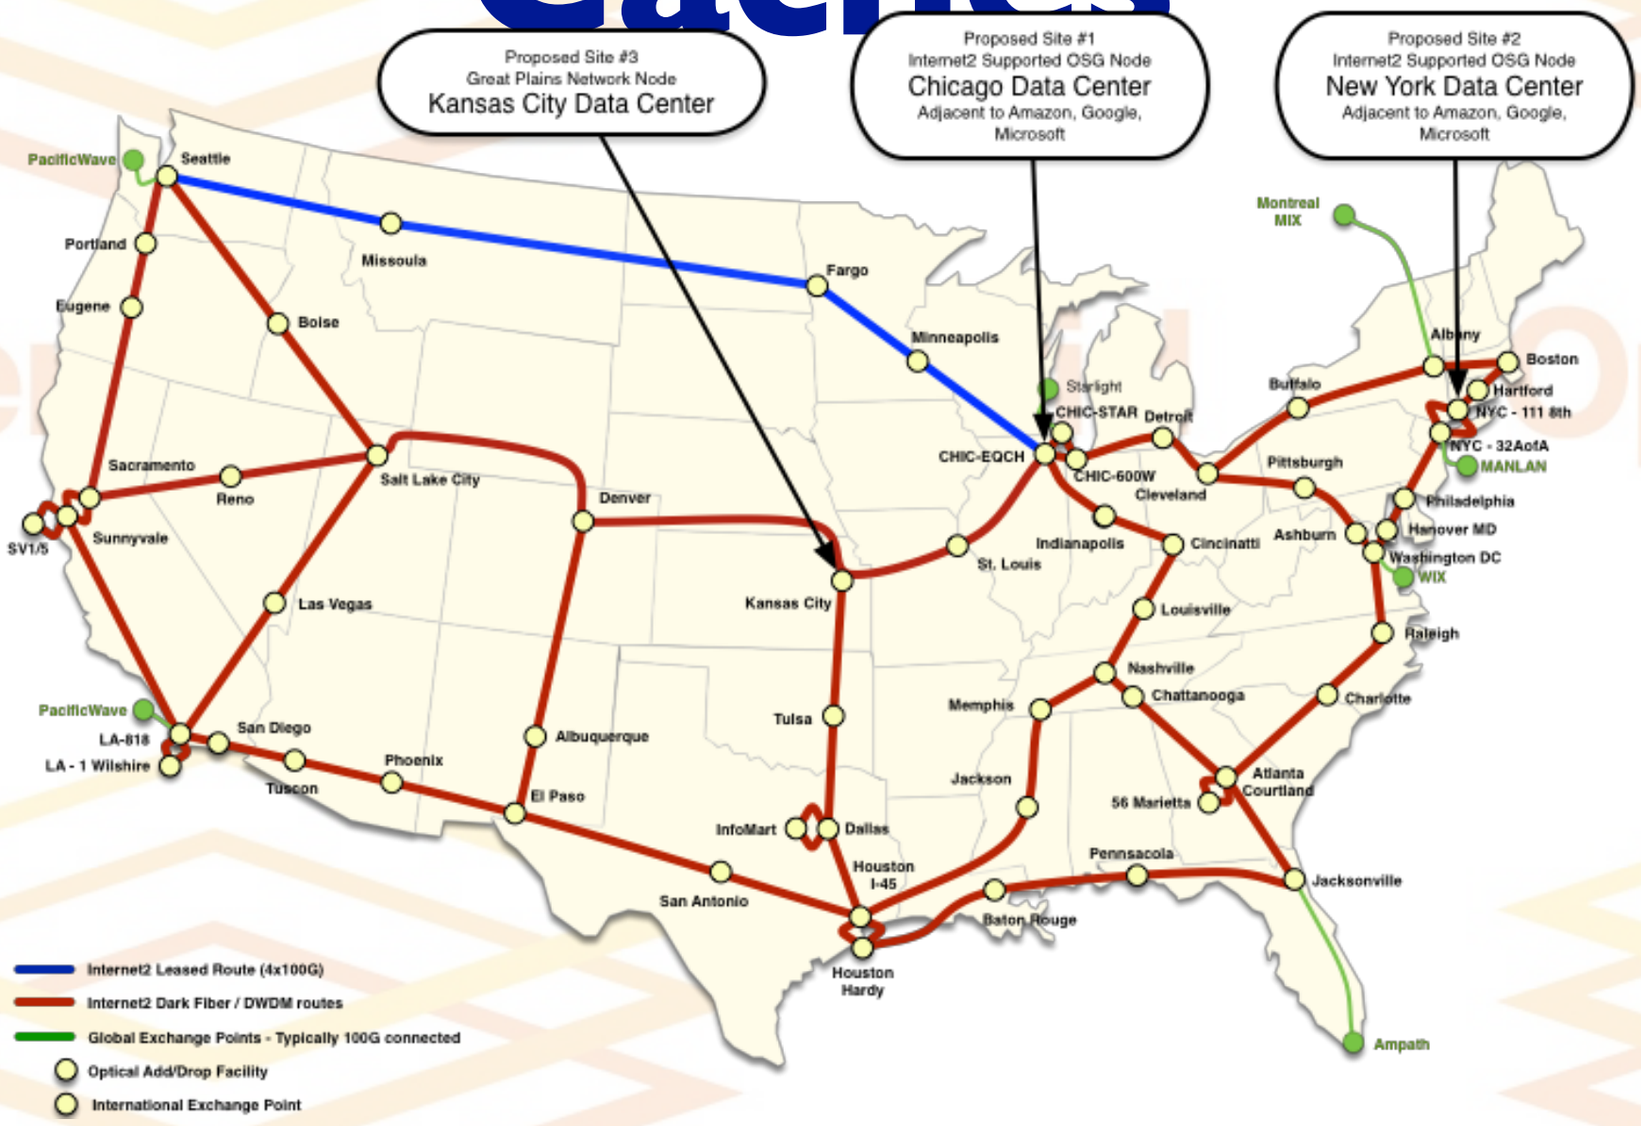

A collaboration between OSG and Internet 2

Endpoint Caches

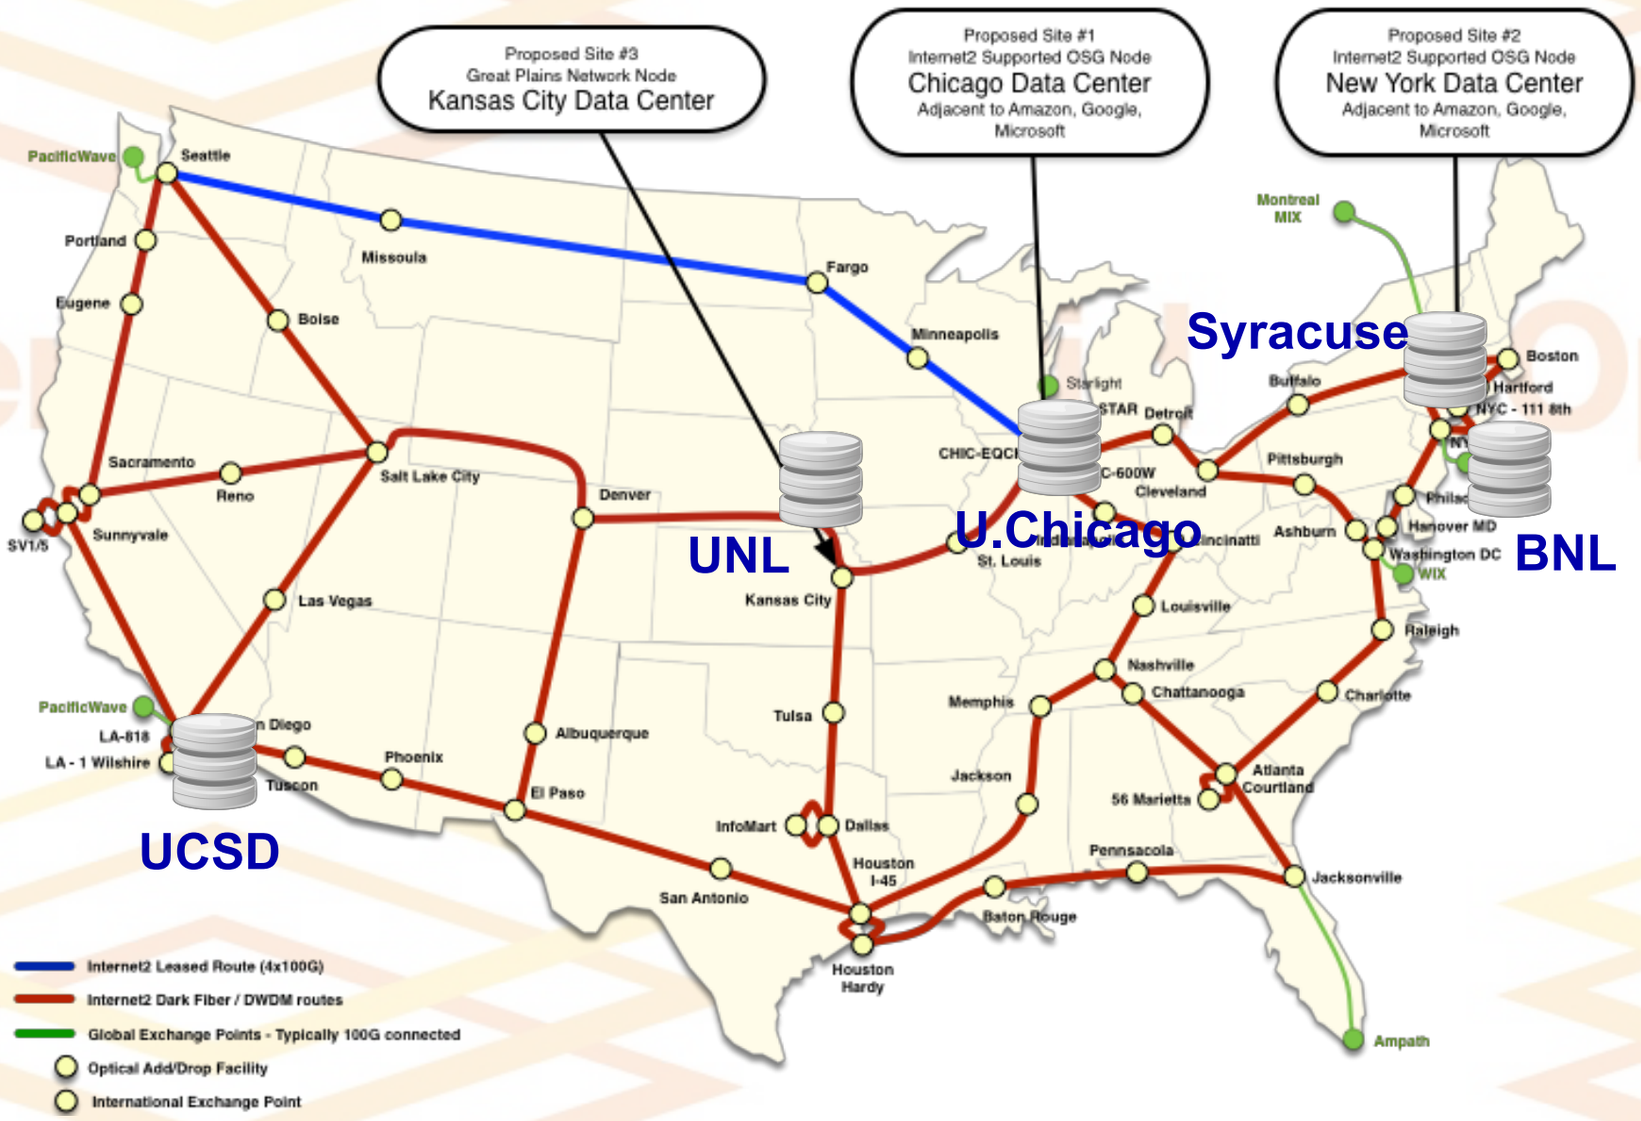

GeoIP

- The two clients for StashCache, CVMFS and StashCP, use GeoIP to determine a nearby cache to utilize.
- Utilizes CVMFS's existing GeoIP services

Deployment and Operational model

- We are leveraging an already existing PRP Kubernetes federation infrastructure maintained by SDSC.

- A shift in traditional grid deployment. The hardware and software responsibilities are split. All software (including the cache) run on docker containers (k8 pods) and it is maintained centrally. The local admins take care of hardware issues (Ex: disk).

- On every node there is a perfsonar pod. This helps deliver a quality of service since now network responsible(PRP) have full access.

- Our current Stashcache Containers can be found at https://github.com/opensciencegrid/prp-stashcache

OSG enables LIGO to run on VIRGO resources

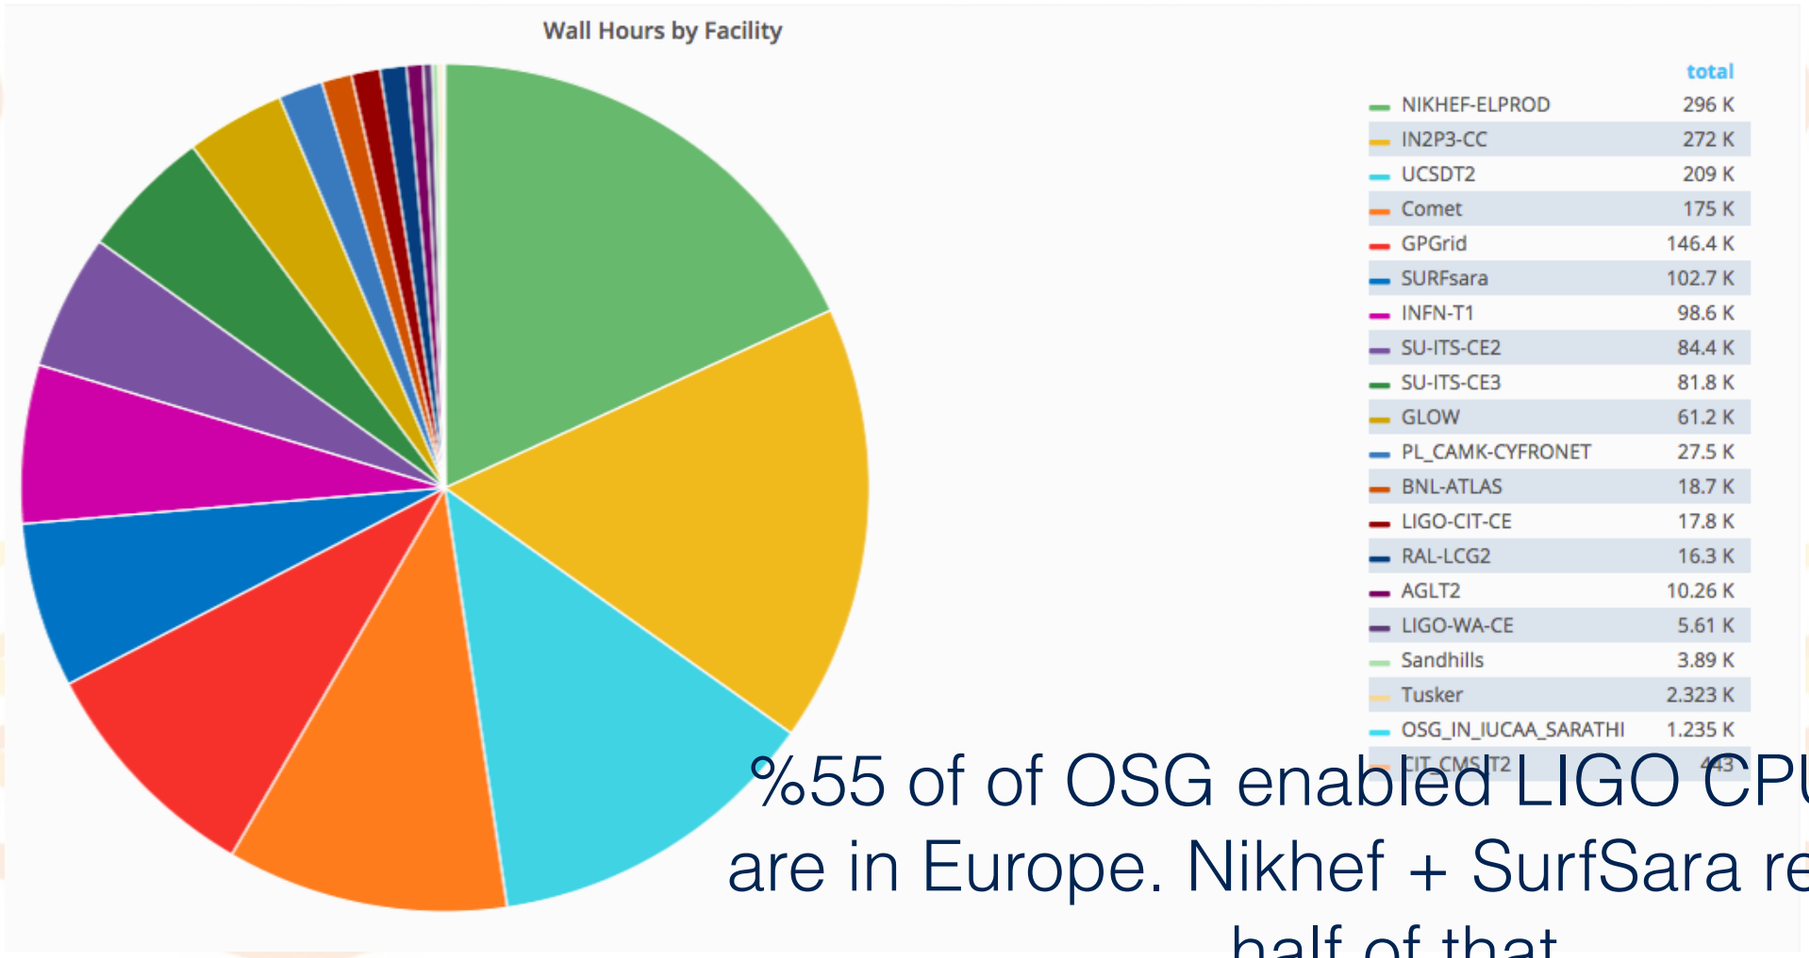

%55 of of OSG enabled LIGO CPU hours are in Europe. Nikhef + SurfSara represent half of that

Caches in Europe

- Given the physical location of the Virgo computing resources and the established partnership between PRP and University of Amsterdam it made sense to place a cache there.
- The cache is managed via k8s with same configuration as the I2 backbone ones.
- Right now it is on acceptance testing phase. We ran into some certificate problems that are now fixed.
- One more cache to be deployed in the first half of 2019 on University of Cardiff and one in Geant in London.

Summary and Takeaway

- Our non ambitious goal is to become the Netflix of science.
- A turn towards a Dev Ops mode of operations in which experts on a given service run those services across all the grid and DEBUG and DEPLOY new versions without the site system administrator intervention.
- A pilot on how potentially all grid services could be deployed on the Grid.
- We are looking to kubernitize and deploy more services this way. Some low hanging fruit are hosted HTCondor-CE and squid servers.