Global Research Platform 2019

Christine Kirkpatrick

This material is based upon work supported by the National Science Foundation under grants # 1747552, 1747493, 1747507, 1747490, 1747483 and the Schmidt Futures Foundation.

Outline

- OSN Project Recap
- Achievements
- Lessons Learned
- What's Next

Research Cyber-Infrastructure Today

Computation	Networking	Storage
Shared Resource (XSEDE, PRAC)	Over 200 universities with 40/100Gb Connectivity	Largely Balkanized
Standardized	Standardized	No Standards Requirement
NSF-Funded	NSF-Funded	No CI Funding

The Open Storage Network:
Distributed National-Scale Storage for Research

A Repository for Active Publicly-Funded Research Data

- ❑ Unified Administration
- ❑ Standardized and Scalable System
- ❑ Purpose-Built System
- ❑ Distributed Platform
- ❑ Fast, Reliable Data Transfer Performance
- ❑ Accessible and Sustainable

Prototype Technical Goals

Ξ Leverage Existing Cyber-Infrastructure Resources, Opening Access Ξ

High-Speed Networking | Funded Datasets

Ξ Efficient Systems Management Ξ

Centralized Management/Monitoring | Limited Use of On-Site Staff | Deployed in SUs

Ξ Provide Safe, Reliable, Consistent Storage Ξ

Policy-Based Redundancy | Data Locality | Geography-Aware Replication | Workload Specific

Ξ Encourage Familiarity with Petascale Systems Ξ

Community Benefits

The OSN will enable increased scientific discovery

- ❑ Assist in the Definition of Standards
- ❑ Leverage Community-Built Tools
- ❑ Enable Data Sharing
- ❑ Cultivate Data Discovery
- ❑ Improve Data Dissemination

Six Prototype Deployment Sites

Funded by NSF

Funded by Schmidt Foundation

- Johns Hopkins University
- Massachusetts Green HPC Center
- Northwestern University (Starlight)
- University of CA San Diego (SDSC)
- University of Illinois (NCSA)
- University NC Chapel Hill (RENCI)

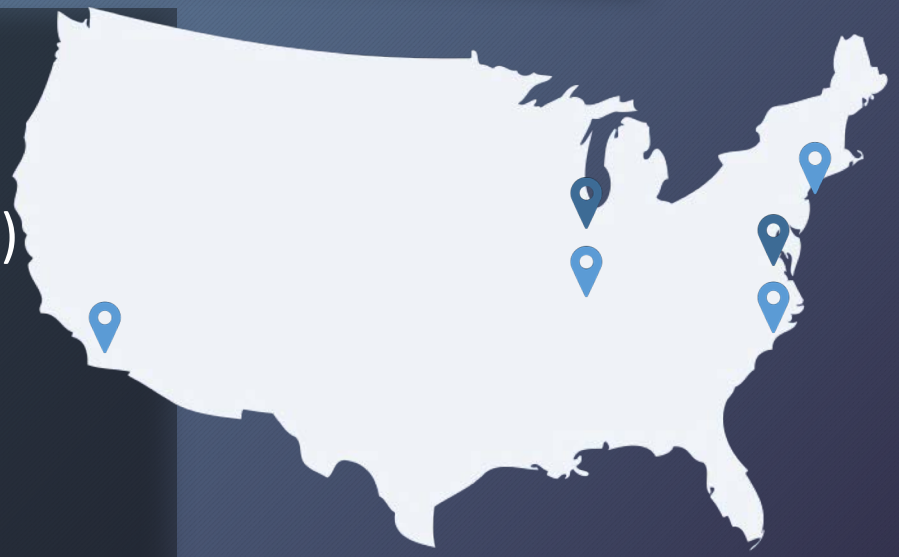

Prototype Details

Hardware is Deployed in “Scalable Units”

~1 PB of Object Storage per Scalable Unit

One Scalable Unit =
Capable of 40Gbps Sequential IO

Geography-Aware Replication =
Data Protection, Data Locality

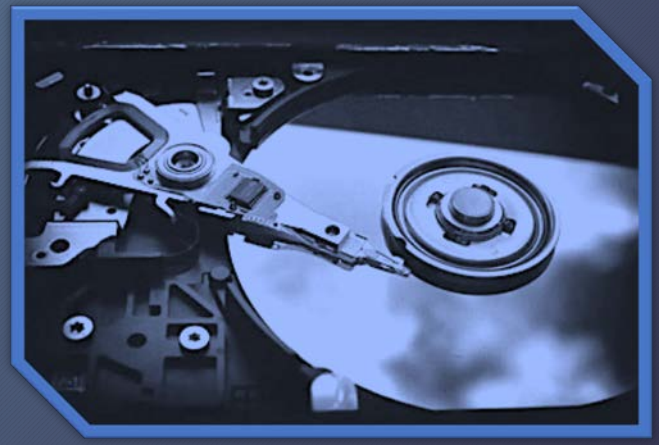

OSN POD Storage View

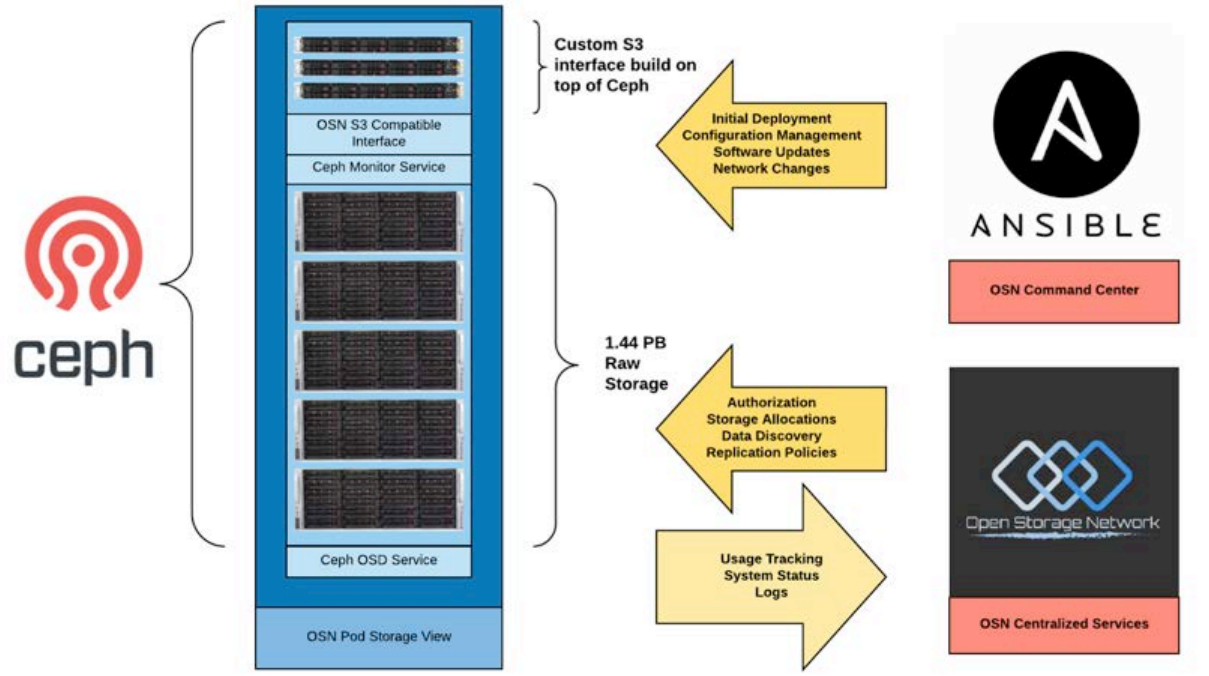

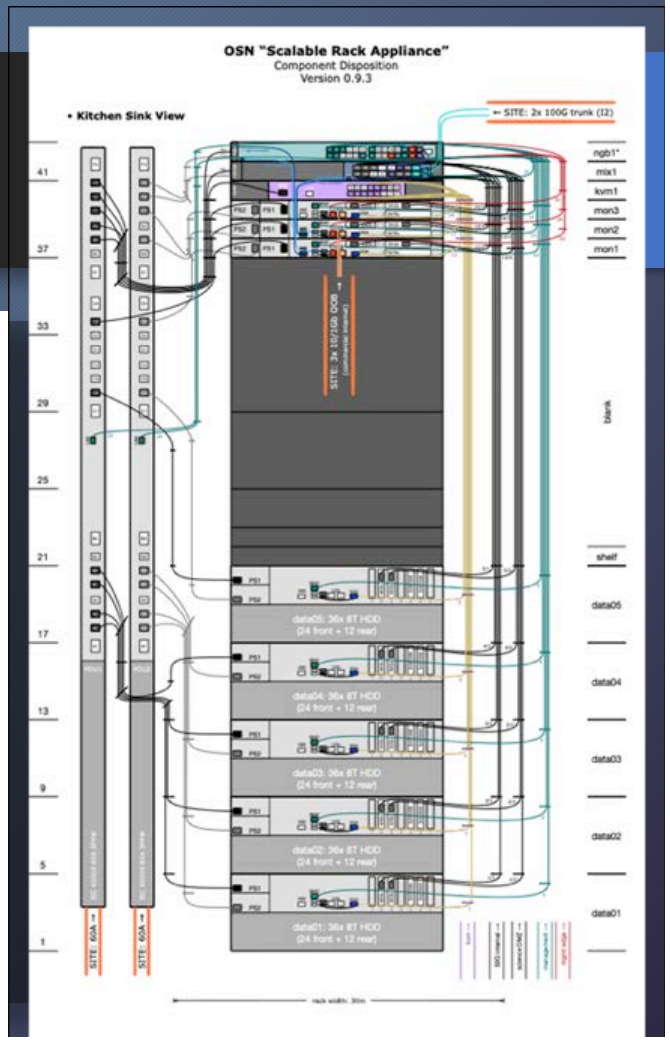

Provisioning Automation

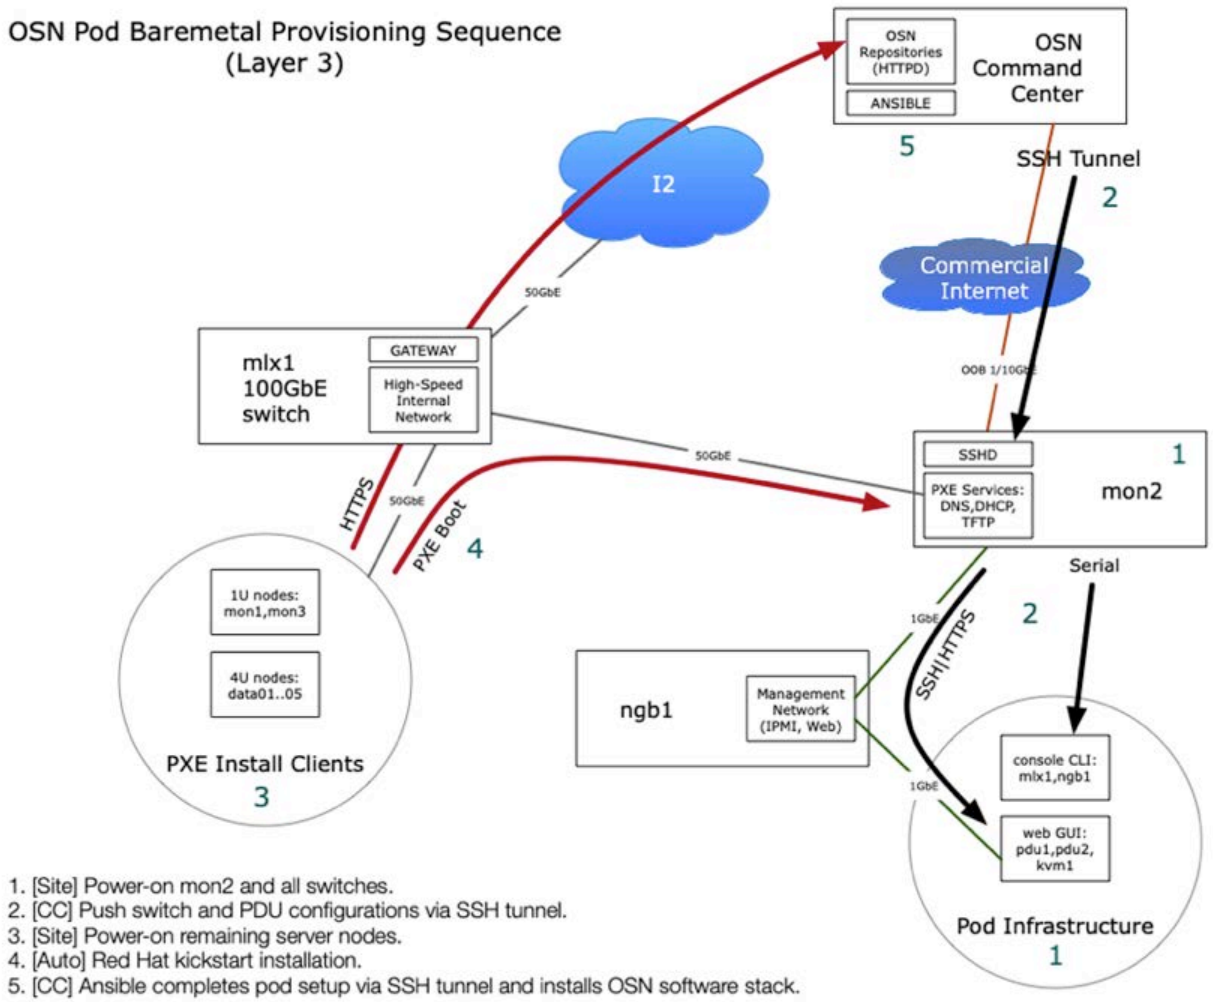

1. [Site] Power-on mon2 and all switches.
2. [CC] Push switch and PDU configurations via SSH tunnel.
3. [Site] Power-on remaining server nodes.
4. [Auto] Red Hat kickstart installation.
5. [CC] Ansible completes pod setup via SSH tunnel and installs OSN software stack.

Use Cases (Midwest Big Data Hub)

Project	Average size of data entities	Total data volume	Storage problem being solved	Use case
Critical Zone Observatories	10 MB	50 TB	Provide storage space and access to CZO datasets and community-generated data	Community long-tail data
TerraFusion	10 GB	1 PB	Transport datasets across the US at high speed, obtain data slices with high probability of reutilization	Experiment-to-site, Slice-and-compute
HathiTrust Research Center collection	200 MB	500 TB	Provide storage space and access to the HTRC dataset and further community-generated derivatives	Common resource access
Machine Learning	10 GB	1 PB	Make available a well-curated dataset for testing machine learning algorithms	Dataset-as-benchmark
Large Synoptic Survey Telescope	2 TB	100 PB	Transport datasets across the US at high speed, obtain data slices with high probability of reutilization, facilitate inter-site data processing	Experiment-to-site, Slice-and-compute, Workflow staging space
Combined Array for Research in Millimeter Astronomy	50 MB	50 TB	Transport datasets across the US at high speed, obtain data slices with high probability of reutilization	Experiment-to-site, Slice-and-compute

Use Case Types

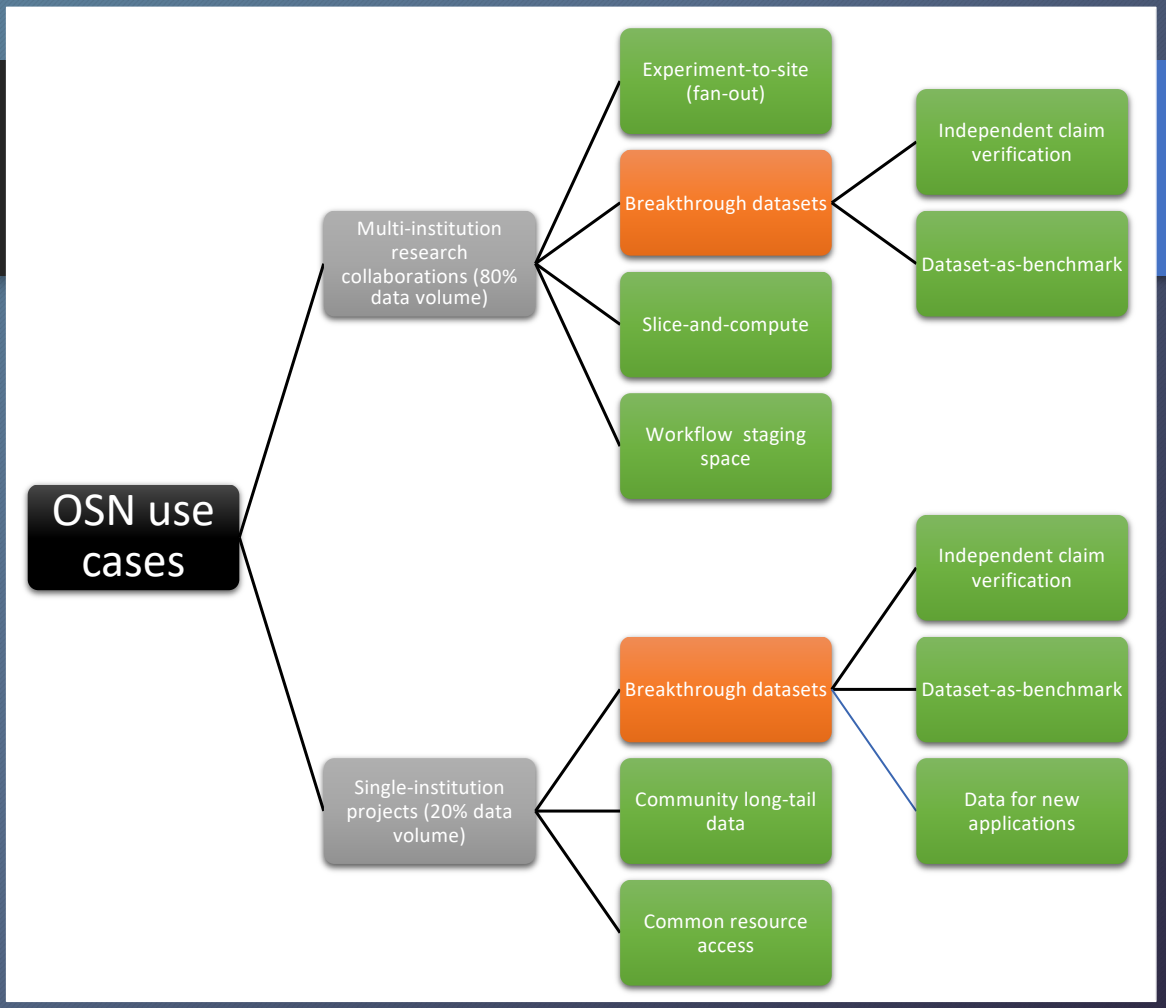

Experiment-to-site

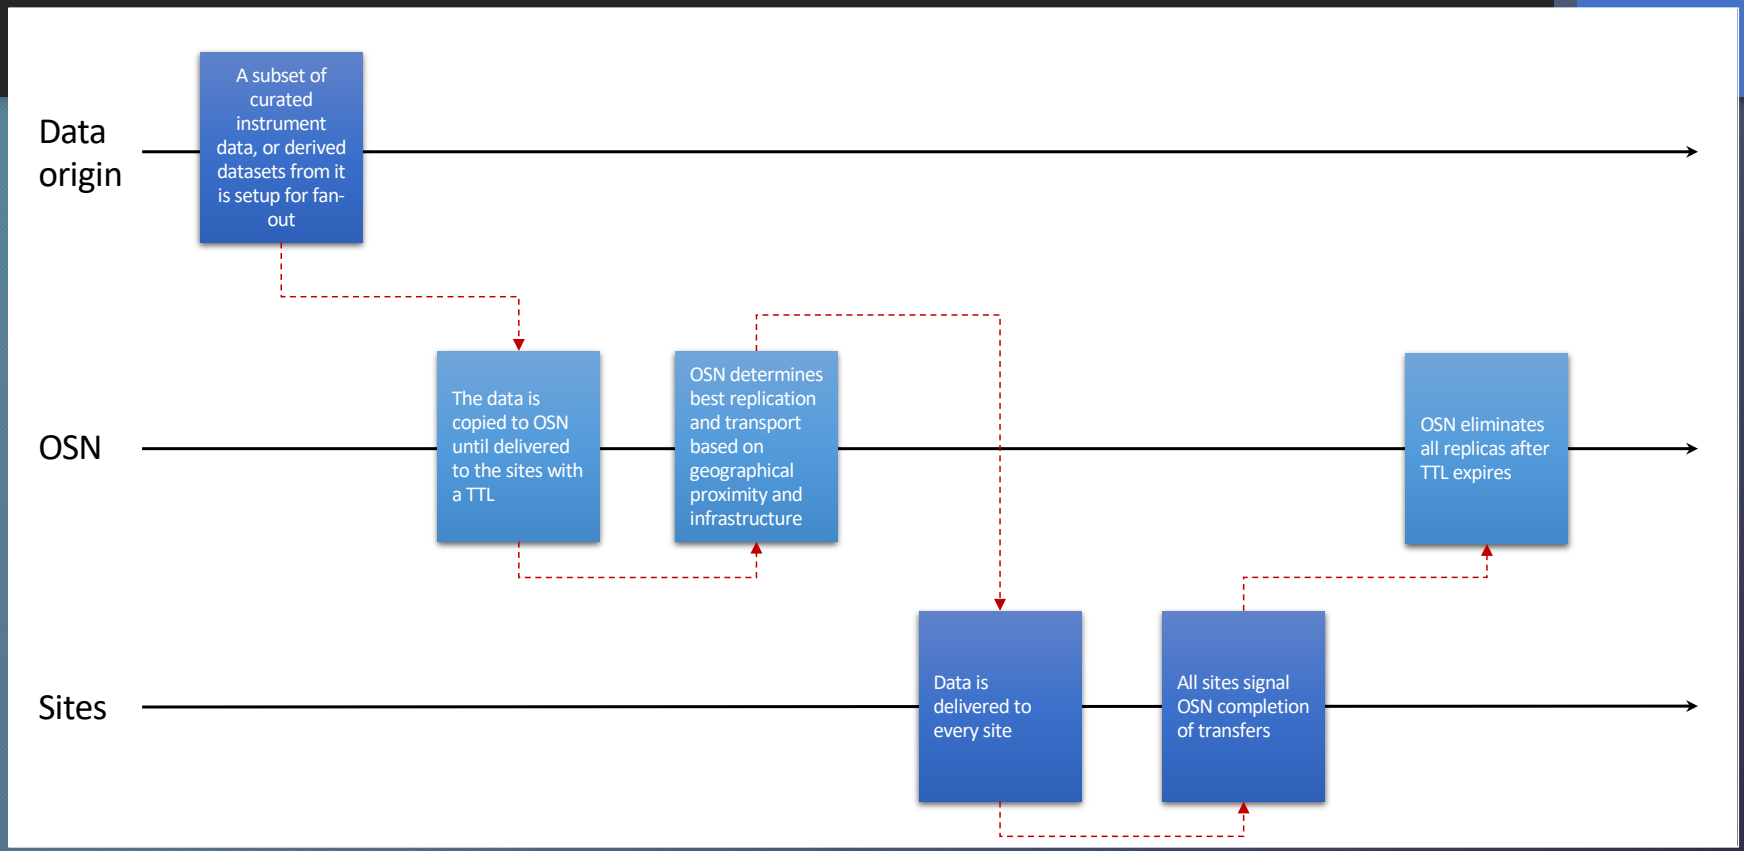

Slice-and-compute

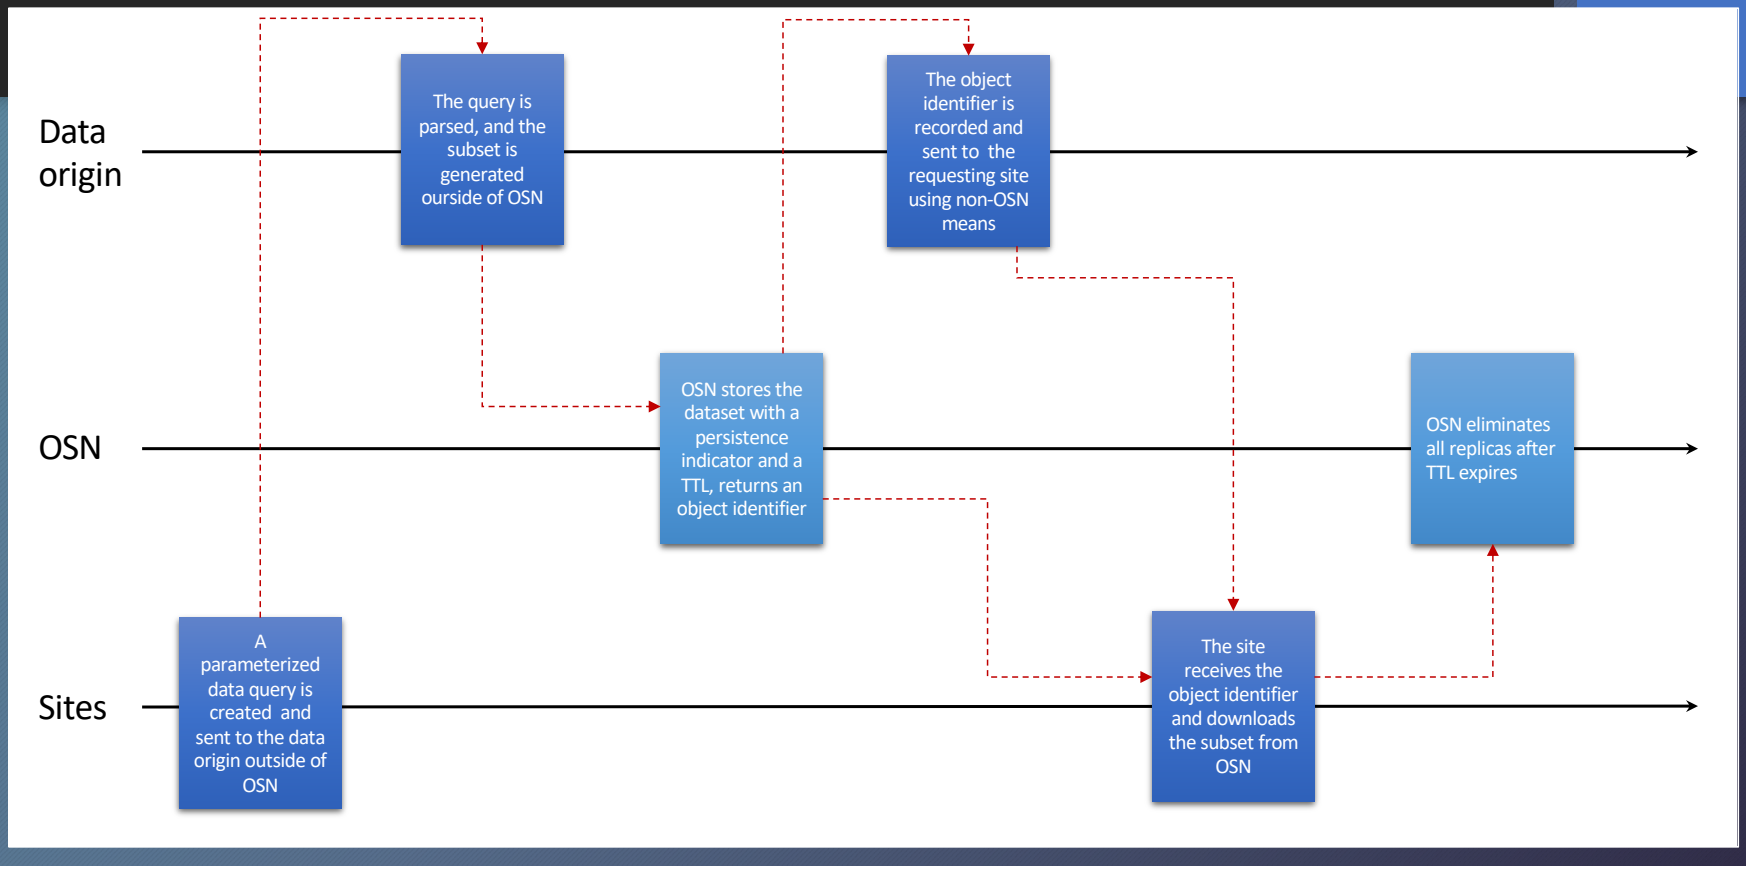

Workflow staging space

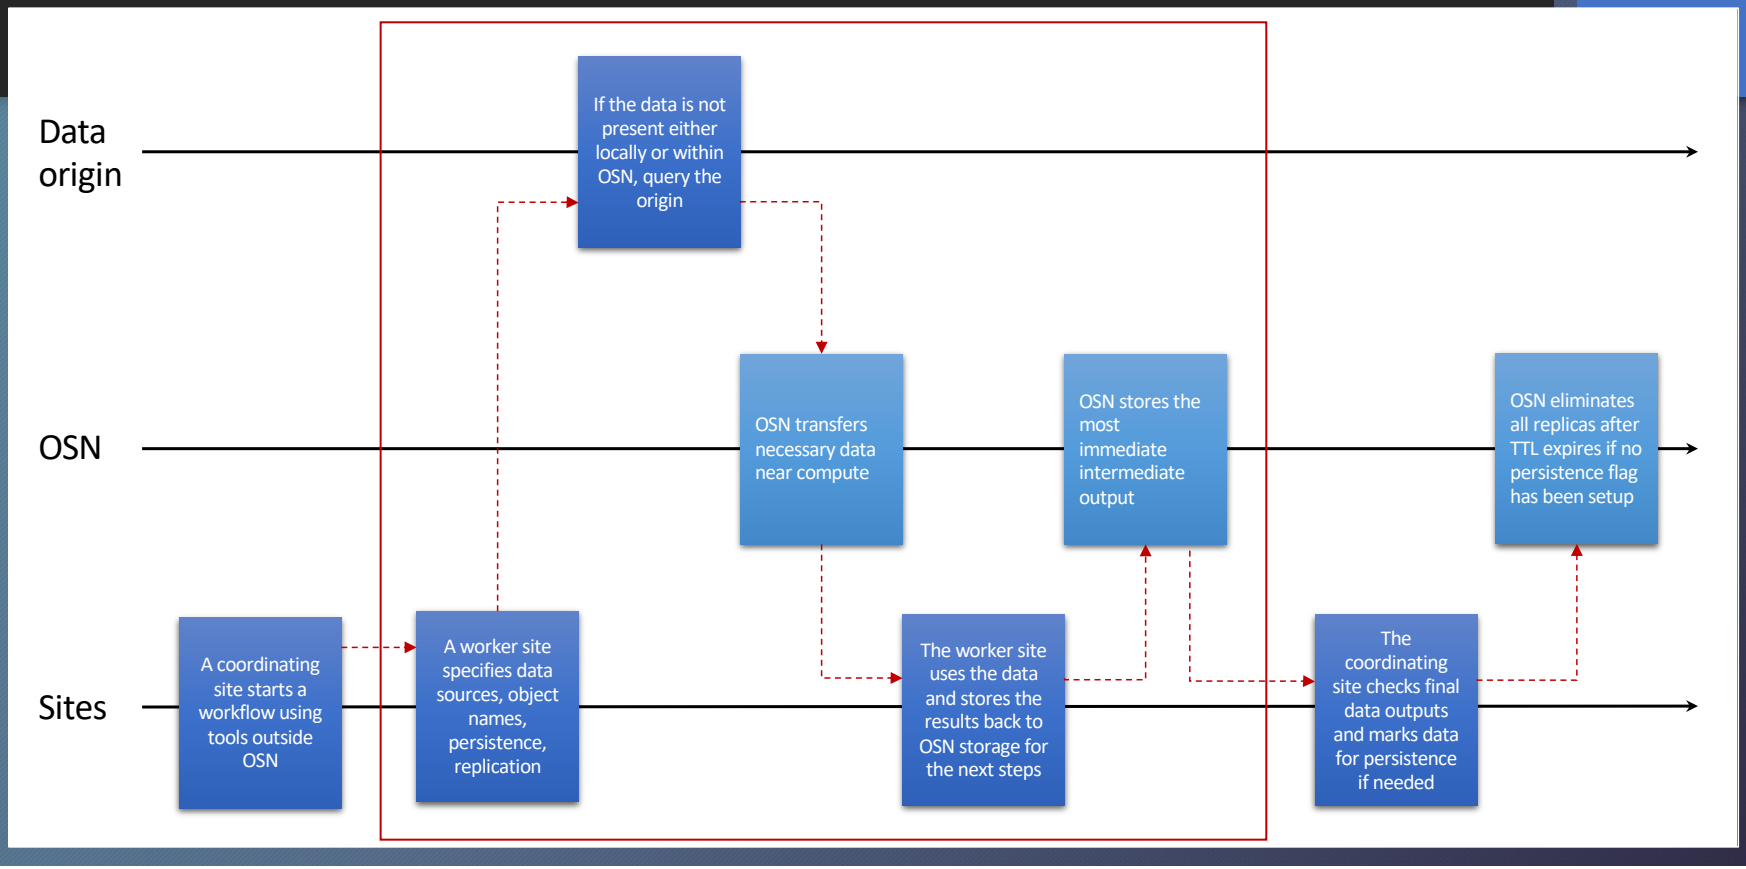

Common resource access

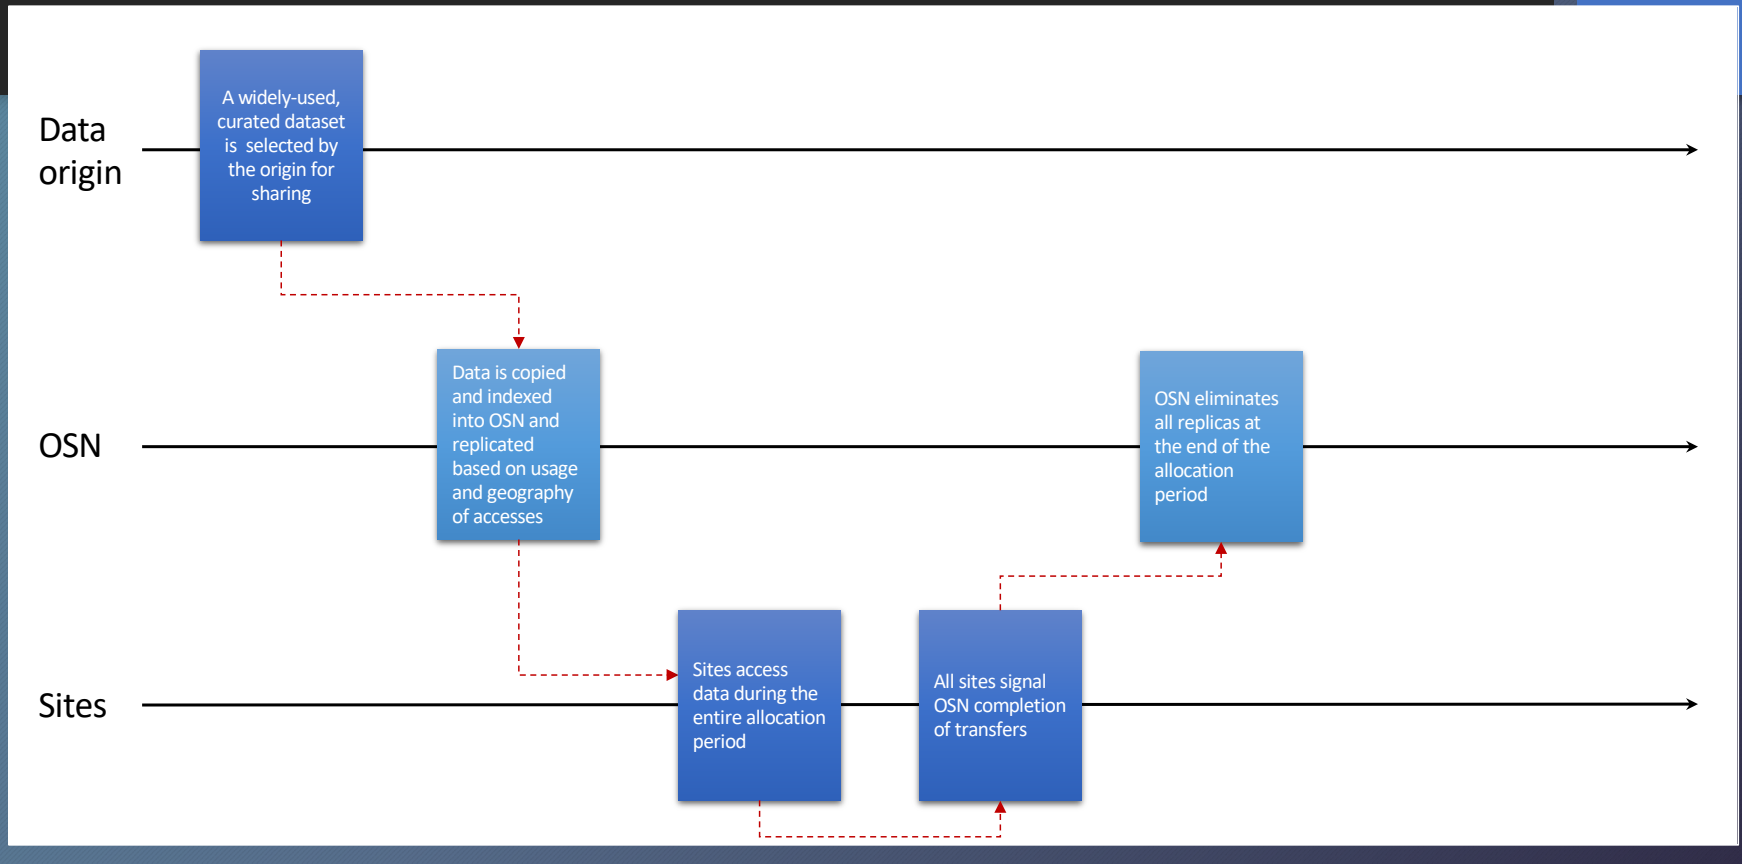

Community long-tail data

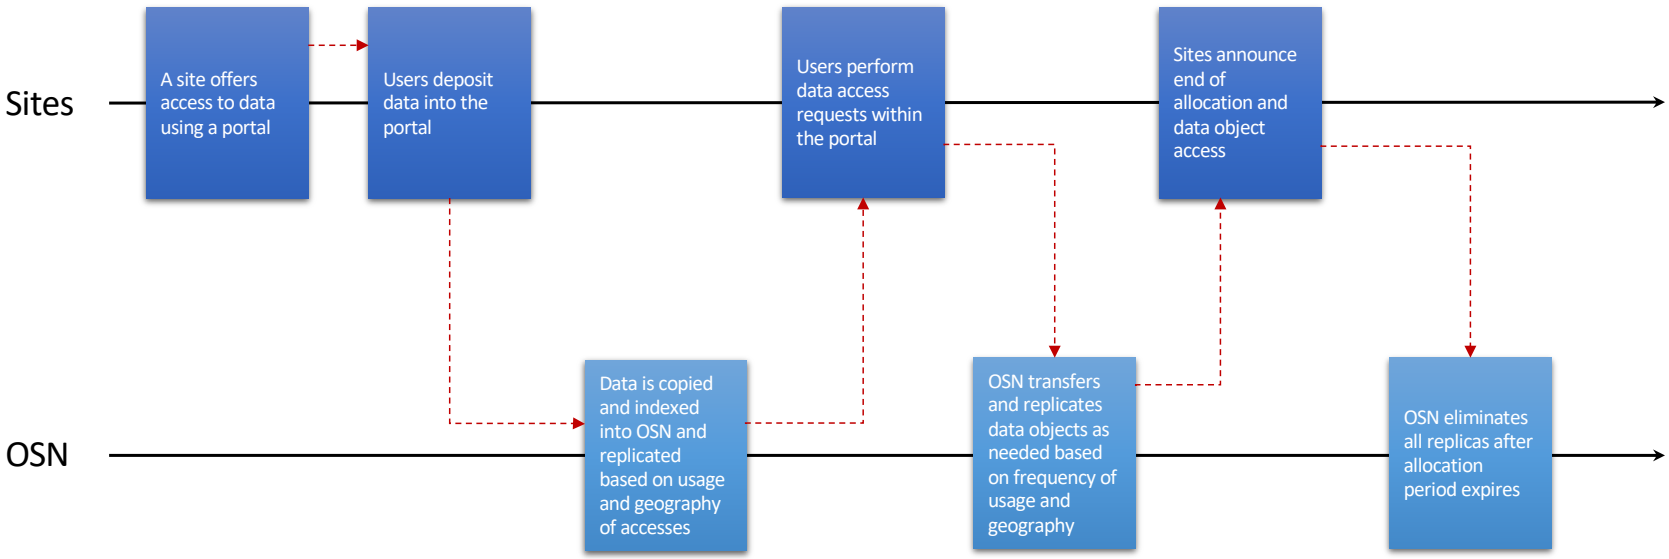

Dataset as benchmark

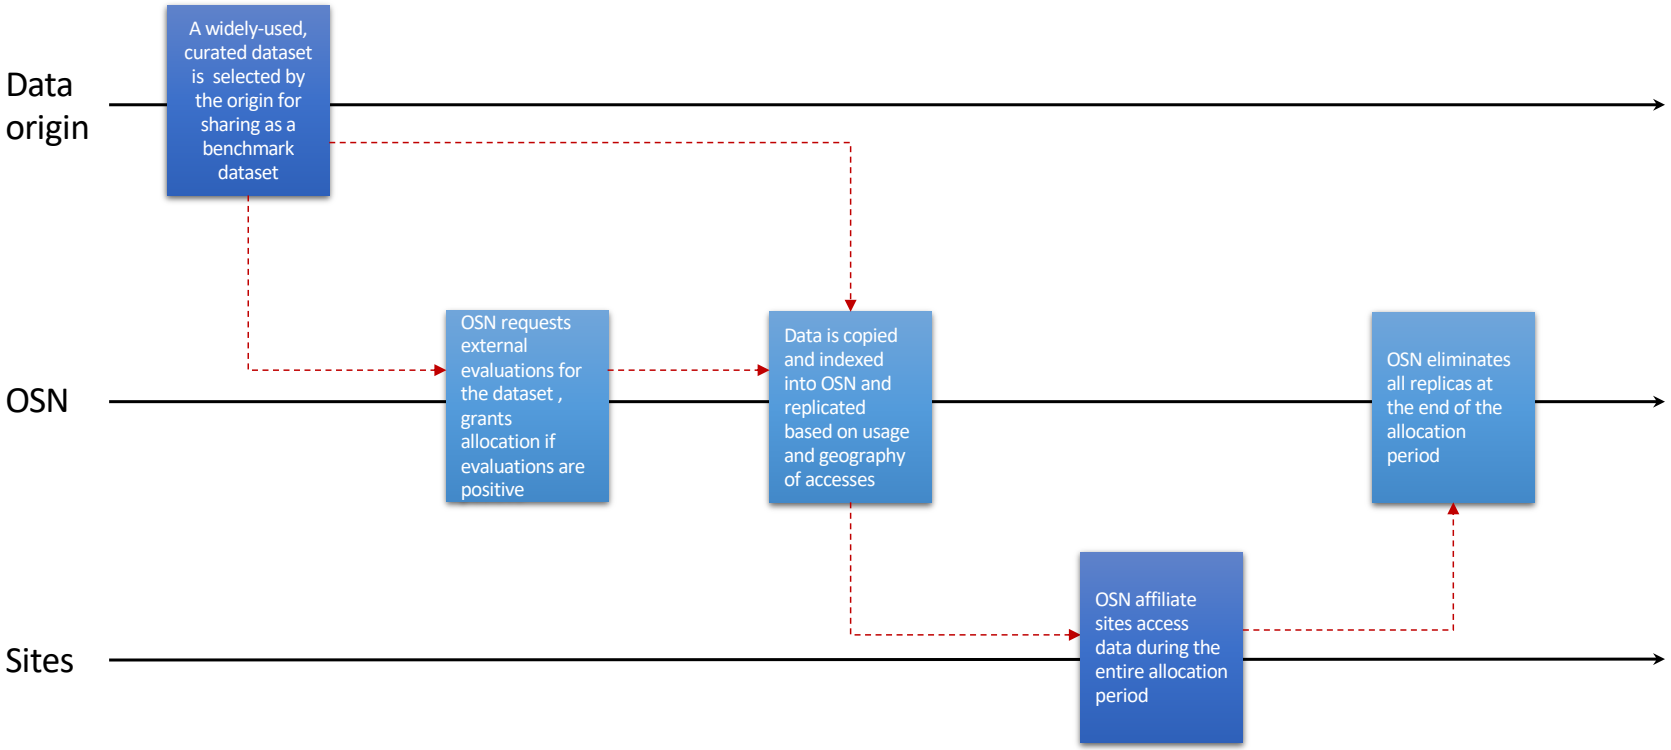

Achievements

- Architecture workshop – Baltimore, September 2018
- Appliances installed
- Built for scale out
  - Installation checklist, kickstart on USB
  - Ansible automation
  - Established command center at SDSC and JHU
- Built for speed – benchmarking between Schmidt nodes
- Governance in progress
- Working groups at work – monthly leadership calls, weekly implementer calls, software, policy WGs.
- Engagement with science use cases

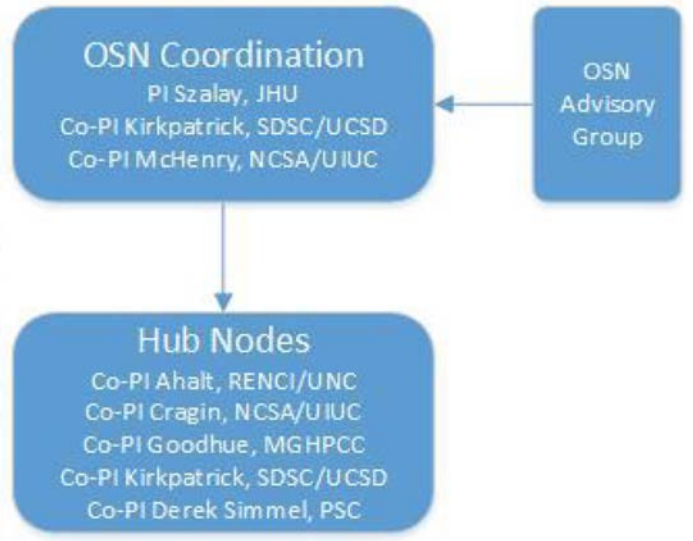

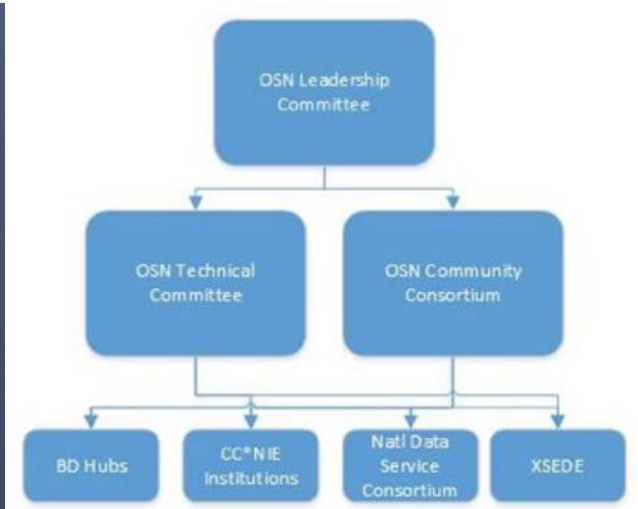

Lessons Learned

- Six nodes – homogenous hardware / 60 nodes – built to requirements
  - End of life hardware
  - Meeting price point, tradeoffs to meet capacity, account for varied sales tax, etc.
  - Preferred vendors, approved integrators
    - Constraints on purchases more diverse than anticipated → multi-institutional negotiation to overcome hurdles, framework established
- Distributed and shared project management
  - Cray hired our PM in Q1 of the project funding
  - Rotating WG chairs, rely on shared PM tools
- We're on the right track.
  - Researchers, projects, other agencies clamoring for OSN

What's Next

- Application layer – Clowder, Globus, Dataverse, iRODS, WholeTale
- Refine governance and shared administration
- Progress on policies, procedures around data lifecycle
- Partnership with XSEDE – Service provider application in progress
- Commercial cloud partnership (AWS)
- Outreach and user-focused workshop – October 10 at TACC
- See us at SC 19!

THANKS FOR YOUR TIME!

openstoragenetwork.org