Open Storage Network: national data storage cyberinfrastructure for the 21st century

Santiago Nuñez-Corrales[4], Melissa Cragin[4], Alainna White[1], Michael Norman[2], Christine Kickpatrick[3], Kenton McHenry[5], Stanley Ahalt[6], Lea Shanley[7], John Goodhue[8], Derek Simmel[9], Alex Szalay[1]
[1]IDIES JHU, [2]SDSC UCSD, [3]NDS and WBDH, [4]MBDH NCSA UIUC, [5]NCSA UIUC, [6]RENCI, [7]UNC Chapel Hill, [8]GCHPCC MIT, [9]PSC UP-CM. Corresponding author: szalay@jhu.edu

Abstract

The increasing amount of scientific data emerging from research projects on all scales is spurring research universities to invest in multi--petabyte (PB) storage systems. At the same time, more than 200 US academic institutions have access to high-speed network connectivity for research purposes through NSF CC*NIE awards, and advanced computing resources through XSEDE. The Open Storage Network (OSN) will demonstrate the potential of a distributed storage infrastructure capable of leveraging high speed links to provide a transparent multi-petabyte data storage and access layer. The OSN proposes a scalable substrate composed of storage appliances that are robust and secure, intended to be simple to manage while supporting various data access patterns. One aim is to develop processes and solutions to encourage widespread adoption. The OSN prototype will leverage existing use cases across several research communities for early adoption, allowing the OSN team to evaluate and test this new type of cyberinfrastructure. At the service level, the OSN will enable flexible and scalable science driven collaborations, necessary for the creation of robust data science software that combines unique data resources regularly and easily.

Project timeline

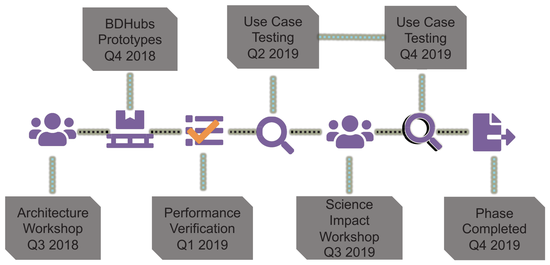

Next steps

- Deployment of scalable elements across deployment sites.
- Implementation of the software and service architectures for the OSN.
- Mapping of scientific communities at large to understand national data storage requirements.

Funded by NSF (IIS 1747552, IIS 1747493, IIS 1747507, IIS 1747490, IIS 1747483) *and the Schmidt Futures Foundation.*

Prototype deployment sites

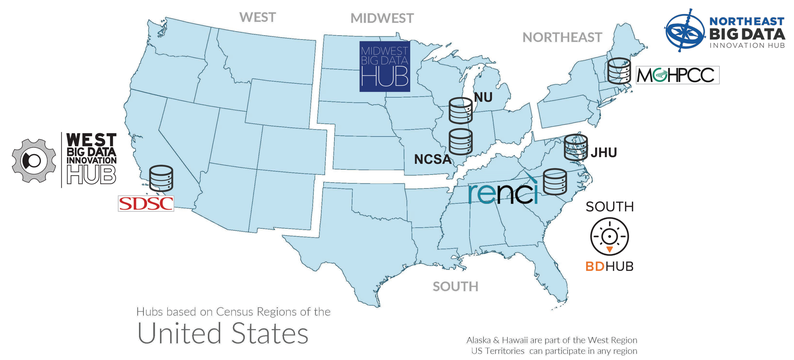

HW/SW architecture

OSN node specs

8 Nodes
Five 4U Data Nodes | Three 1U Monitoring Nodes

1.44 PB Storage
Raw | 7200 RPM 12Gb SAS | 36 Disks per Data Node

High Speed Uplinks
40Gb Public Uplink | 40Gb Backend Uplink

Ceph-Spec'ed Performance
One HT Core per OSD | .89 RAM to TB Storage Ratio

Service architecture

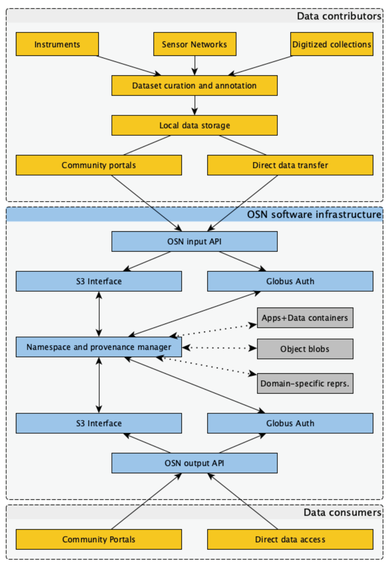

Scientific use cases

Project	Research area	Average size of data entities	Total data volume
Connectomics	Neuroscience	10 GB	2 PB
Critical Zone Observatories	Earth Sciences	10 MB	50 TB
TerraFusion	Earth Sciences	10 GB	1 PB
Global ocean modeling	Climatology and Oceanography	5 GB	4 PB
HathiTrust Research Center collection	Digital Humanities	200 MB	500 TB
Machine Learning	Neuroscience, Computer Science,	10 GB	1 PB
Sloan Digital Sky Survey	Astronomy	15 MB	70 TB
Large Synoptic Survey Telescope	Astronomy	2 TB	100 PB
Combined Array for Research in Millimeter Astronomy	Astronomy	50 MB	50 TB
Watershed Models at the Process Scale	Earth Sciences	1 GB	2 TB
Collaborative Gene Matching	Bioinformatics	1 GB	1 PB

Infrastructure design goals

Technical goals

Ξ Leverage Existing Cyber-Infrastructure Resources 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

+

Science goals

Prompt definition of data standards | Leverage community-built tools

Enable data sharing | Cultivate data discovery

Improve data dissemination | Foster multi-disciplinary interaction

References:

Thompson, K. (2012). Campus Cyberinfrastructure–Network Infrastructure and Engineering (CC-NIE). National Science Foundation, December 2012.

Szalay A.S., Du, David, Foster, I., McHenry, K., Neuhauser, C. (2016), "CC*Data: Prototype of a National- Scale Distributed Storage System for Open Research Data," quotes and portions from unfunded proposal to NSF, August 2016.

Links:

NSF CC*NIE: https://bit.ly/2qhQmwe
XSEDE: https://www.xsede.org
Connectomics: https://lichtmanlab.fas.harvard.edu
CZO: http://criticalzone.org/national/
Terra Fusion: https://go.nasa.gov/2ql6NIm
Global ocean modeling: https://bit.ly/2Oe3jRi
HathiTrust Research Center: https://wiki.htrc.illinois.edu
Machine learning: http://chemimage.illinois.edu
SDSS: https://www.sdss.org
LSST: https://www.lsstcorporation.org
CARMA: http://carma-server.ncsa.uiuc.edu:8181
Watershed models: https://www.hydroshare.org

https://www.openstoragenetwork.org/