

# “RAPID EARTH SCIENCE DATA DISTRIBUTION OVER A MULTI-INSTITUTIONAL OPEN STORAGE RESEARCH INFRASTRUCTURE”

Jeremy Musser, Miao Zhang, Jayashree  
Candadai, Ezra Kissel, Martin Swany  
**Indiana University**  
{jemusser, mzhang, jayaajay ezkissel,  
swany}@iu.edu

Patrick Gossman, Matthew Lessins,  
Carlo Musante, Michael Thompson  
**Wayne State University**  
{pgossman, mjl, carlo, michael} @wayne.edu

Shawn McKee, Benjeman Meekhof, Lillian  
Huang, Charles Antonelli  
**University of Michigan**  
{smckee, bmeekhof, lihuang, cja} @umich.edu

Kenneth Merz, Andrew Keen, Charlie Miller  
**Michigan State University**  
{merzjrke, keenandr, cdmiller} @msu.edu

## Abstract

Research scientists face a number of obstacles in obtaining and working with large data sets over existing network and storage infrastructure. This data-access problem is often amplified when coordination among distributed, multi-institutional resources is in the critical path for conducting collaborative research. Our NRE entry highlights an application integrated with the Multi-Institutional Open Storage Research Infrastructure (MI-OSiRIS) [1] as a demonstration of how optimized data interfaces within a software-defined, scalable storage buildout can address many of the data-intensive and collaboration challenges faced by researchers and their respective communities. The Earth Observation Depot Network (EODN) [2] is our representative application that disseminates multi-resolution remote sensing data across distributed Ceph storage resources within the MI-OSiRIS deployment, which span well-connected sites in Michigan, Indiana, and local SCinet and Utah Cloudlab infrastructure in Salt Lake City. We intend for our demonstration to emphasize a network-intensive workflow that intelligently stages terabyte-scale data sets for processing and visualization.

## 1. Overview

The OSiRIS project provides a distributed, multi-institutional storage infrastructure that lets researchers write, manage, and share data from their own computing facility locations. Its goal is transparent, high-performance access to the same storage infrastructure from well-connected locations on any participating campus. It includes network discovery, monitoring and management tools as well as the creative use of Ceph features. The project provides data sharing, archiving, security and life-cycle management, implemented and maintained with a single distributed service.

The EODN aims to enable open access, reduced latency, and fast downloads to valuable and compelling Earth science data from satellites for meteorological and atmospheric researchers, for example remote sensing data sourced from NASA's Earth Science Program (EOS-DIS). The EODN effort is motivated by the fact that rapid access to satellite imagery is crucial in developing near-real time decision-support for emergency management, disaster response, and operational forecasting. Less latency-sensitive consumers of high-resolution Landsat imagery, in particular, include applications within agriculture, environmental sciences, land use/cover studies, urban planning and development, and education and outreach.

With the OSiRIS deployment, we are able to augment EODN with substantial, fast storage capacity connected by high-speed networks thereby enhancing the reach and capabilities of the remote sensing data workflow. EODN file metadata distributed among Ceph object stores is managed within the Unified Network Information Service (UNIS) [3] along with network topology and measurement metadata, allowing for integrated data movement services and workflow visualization tools to operate within a unified framework. At SC16, we intend to demonstrate the current state of multi-site data distribution based on pre-defined workflow policy and network measurement feedback collected between OSiRIS installations, including the ability to make use of ad-hoc instances as instantiated on cloud-based and SCinet infrastructure.

## II. Innovation

Instituting effective multi-institutional, at-scale research collaboration has been a challenge for a number of reasons:

- Fast, scalable, and accessible (i.e. well-connected) shared storage is not readily available
- Buildouts are often static and prone to periods of outages and/or service degradation
- A lack of staging mechanisms that moves data to where it is needed (data locality)
- Poor discovery and indexing mechanisms of existing and newly generated data
- A lack of effective sharing and access controls in a multi-institutional setting (federation)

The OSiRIS effort is addressing these points by (i) leveraging advanced networks buildouts within participating campuses and integrating proven distributed storage technologies (e.g. Ceph), (ii) making network topology awareness plus cross-site monitoring and measurement that directs dynamism through SDN capabilities a central feature within the OSiRIS architecture, (iii) supporting data movement and content management services that implement a distribution policy (replication factor, specific data locality, etc.) across available storage resources, and (iv) integrating InCommon federated authentication and authorization mechanisms at the data access layer. Of key value at SC16 will be the evaluation of OSiRIS, and specifically Ceph and associated configuration, in providing a performant shared storage capability over WAN latencies and unbalanced I/O between the deployed installations.

## III. HPC and Science Relevance

In the age of “Big Data”, the efficient transfer and sharing of research data is nearly universally understood to be of critical importance. Reducing the time and complexity needed to acquire and distribute files enables scientists to collaborate more effectively and rapidly test additional hypotheses within their field. Our design and evaluation of OSiRIS, and the science workflows it supports, allows us to further best-practices in deploying new technologies and services within increasingly complex and inter-related network and storage buildouts.

## IV. SCinet and R&E Requirements

Depending on final network configuration and bandwidth requirements, we may request demo reservation times involving the IU and UMICH booths. Given 100G connection(s) a custom configuration on the Corsa SDN equipment is a possibility. Large flow requirements and connectivity from Cloudfab remain TBD.

## V. Network Topology

A preliminary, high-level diagram of our network topology is shown in Figure 1. The “far-end” of our demo will primarily involve OSiRIS deployments in Michigan (MiLR) via 710 NLS. Details TBD. We are also in contact with Cloudfab to bring connectivity via one or more VLANs to SCinet.

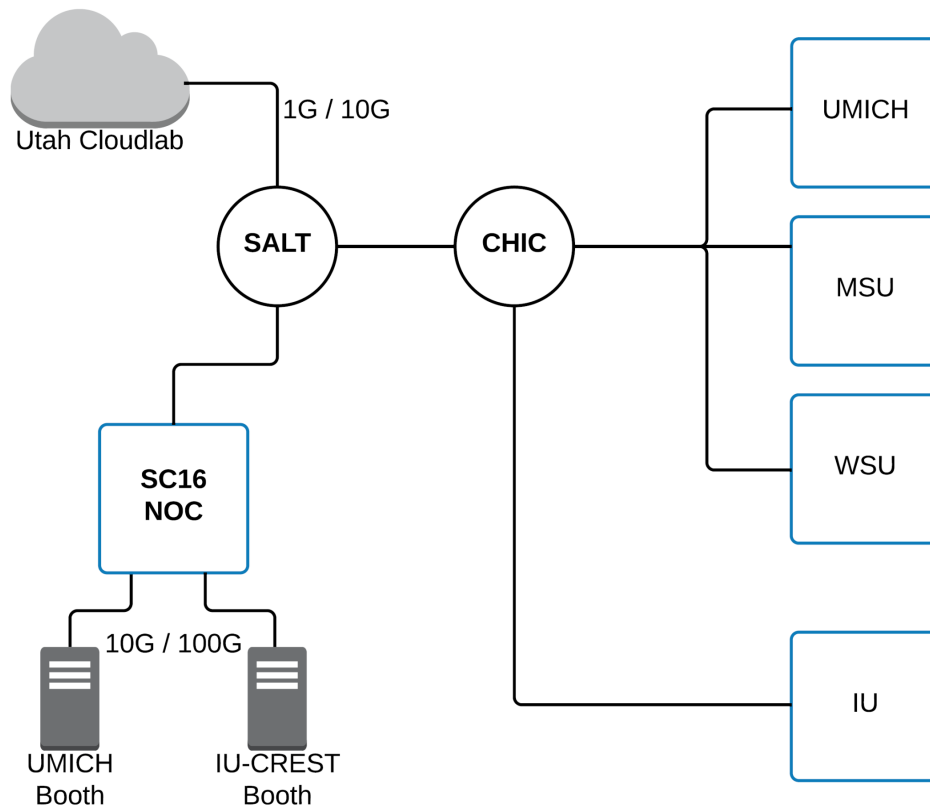


Figure 1

### References (Optional)

[1] MI-OSiRIS <http://www.osris.org/>

[2] Earth Observation Depot Network (EODN) <https://data-logistics.org/?q=EODN>

[3] A. El-Hassany, E. Kissel, D. Gunter and M. Swany, "Design and Implementation of a Unified Network Information Service," in 11th IEEE International Conference on Services Computing, 2014.