

Midwest group of universities to extend computing network and storage system to Utah for industry conference

Nov. 11, 2016

Contact: Dan Meisler, dmeisler@umich.edu, 734-764-7417

A group of midwestern universities will extend its innovative Big Data storage and networking system next week from Michigan and Indiana 1,500 miles west to the Supercomputing 16 (SC16) conference in Salt Lake City, Utah.

The NSF-funded project, called Open Storage Research InfraStructure or OSiRIS, is led by the University of Michigan and is meant to remove the bottleneck of data storage and transfer that currently slows down multi-institutional research collaborations that use massive data sets.

Principal Investigator Shawn McKee of the U-M Physics Department, said the demonstration of OSiRIS will show the speed and ease of access to large data sets that the system allows, and its scalability beyond the midwest.

“The goal is to get scientists connected to data as quickly and easily as possible, so they can focus on their research and not have to worry about data management,” McKee said.

The project is a partnership between U-M, Indiana University, Michigan State University, and Wayne State University, and is funded by a \$4.9 million National Science Foundation grant.

OSiRIS is built on software-defined storage, allowing off-the-shelf hard drives to be programmed with intelligent software — in this case, an open-source system called Ceph. U-M’s Advanced Research Computing - Technology Services (ARC-TS) unit is building the infrastructure for the system, and the project is supported by the Michigan Institute for Computational Discovery and Engineering (MICDE).

The SC16 demonstration will showcase the project's ability to run a distributed Ceph cluster by turning the U-M and MSU exhibition booth into another OSiRIS storage site. The SC16 site will participate equally in storing and replicating data with the existing sites at U-M, WSU, and MSU. The demo will include visualizations of Earth Science research data flows into OSiRIS coordinated by tools from the OSiRIS IU team.

“One of our key goals with OSiRIS is to enable compute clusters at one site to access remote datasets directly,” notes Patrick Gossman, Deputy CIO for Research, Wayne State University. “This accelerates the research process by eliminating the costly and time-consuming process of copying very large data sets.”

Andrew Keen, HPC Architect, Michigan State University adds, “Another exciting aspect of the OSiRIS project is the way it uses open source software defined storage and networking to provide object, file, or block-level access and the ability to control where and how the data is stored based on the network topology and researchers' needs, without the need for researchers to build or implement large scale data management themselves.”



Dr. Kenneth Merz, Co-PI and the director of the Institute for Cyber-Enabled Research (iCER) and Joseph Zichis Chair in Chemistry at Michigan State University, adds that "Collaboration amongst institutions of higher learning in Michigan has led to the OSiRiS project which will benefit not only each individual institution, but will facilitate collaboration in the data sciences within the state and without."

"The OSiRiS deployment promises to significantly improve existing application workflows by making available a large, managed storage resource that would otherwise be unobtainable to certain communities," said Ezra Kissel, Research Scientist at the Indiana University Center for Research in Extreme Scale Technologies. "As one such example, we are excited to highlight the use of OSiRiS in enhancing the distribution of, and access to, geospatial and remote sensing data sets for use by Earth Science applications."

For more information: <http://www.osris.org/>