



Contribution ID: 23

Type: **not specified**

The OSiRIS Project: Meeting the Multi-Institutional Data Collaboration Challenge

Monday, April 18, 2016 4:35 PM (25 minutes)

The OSiRIS (Open Storage Research Infrastructure) project started in September 2015, funded under the NSF CC*DNI DIBBs program. This program seeks solutions to the challenges many scientific disciplines are facing with the rapidly increasing size, variety and complexity of data they must work with. As the data grows, scientists are challenged to manage, share and analyze that data and become diverted from a focus on their scientific research to data-access and data-management concerns. Even more problematic is determining how to support many scientists sharing and accessing this ever increasing amount of data across multiple institutions.

We will describe how the OSiRIS project is tackling this challenge using a combination of Ceph, software-defined storage, various open-source management, security and monitoring components and software-defined networking to enable an infrastructure that supports multi-institutional access for scientists working to collaboratively extract scientific results from large, distributed or diverse data. The presentation will cover the current status of OSiRIS as we bring up it's initial implementation, describe the technical details and choices we have made, and summarize our longer term goals and plans for this 5-year project.

Length of presentation (minutes, max. 20)

20

Primary author: MC KEE, Shawn (University of Michigan (US))

Co-author: MEEKHOF, Benjeman Jay (University of Michigan (US))

Presenter: MC KEE, Shawn (University of Michigan (US))

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems