

OSiRIS: A Distributed Ceph Deployment Using Software Defined Networking for Multi-Institutional Research

Monday, October 10, 2016 3:30 PM (15 minutes)

We will report on the first year of the OSiRIS project (NSF Award #1541335, UM, IU, MSU and WSU) which is targeting the creation of a distributed Ceph storage infrastructure coupled together with software-defined networking to provide high-performance access for well-connected locations on any participating campus. The project's goal is to provide a single scalable, distributed storage infrastructure that allows researchers at each campus to read, write, manage and share data directly from their own computing locations. The NSF CC*DNI DIBBs program which funded OSiRIS is seeking solutions to the challenges of multi-institutional collaborations involving large amounts of data and we are exploring the creative use of Ceph and networking to address those challenges.

While OSiRIS will eventually be serving a broad range of science domains, its first adopter will be ATLAS, via the ATLAS Great Lakes Tier-2 (AGLT2), jointly located at the University of Michigan and Michigan State University. Part of our presentation will cover how ATLAS is using the OSiRIS infrastructure and our experiences integrating our first user community. The presentation will also review the motivations for and goals of the project, cover the technical details of the OSiRIS infrastructure, the challenges in providing such an infrastructure, and the technical choices made to address those challenges. We will conclude with our plans for the remaining 4 years of the project and our vision for what we hope to deliver by the project's end.

Tertiary Keyword (Optional)

Network systems and solutions

Secondary Keyword (Optional)

Object stores

Primary Keyword (Mandatory)

Storage systems

Primary authors: SWANY, Martin (Helsinki Institute of Physics (FI)); MC KEE, Shawn (University of Michigan (US))

Co-authors: MEEKHOF, Benjeman Jay (University of Michigan (US)); Dr KISSEL, Ezra (Indiana Univesity)

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

Session Classification: Track 4: Data Handling

Track Classification: Track 4: Data Handling