



Contribution ID: 238

Type: **Presentation**

Bridging HPC and Cloud: Deploying an S3 Frontend for POSIX-compliant Storage

Friday 21 March 2025 11:30 (15 minutes)

As research communities increasingly rely on cloud-native tools and workflows, integrating High-Performance Computing (HPC) environments with cloud storage has become increasingly important. From the perspective of an IT support team, this presentation outlines our initial approach towards a lightweight, easily deployable S3-layer on existing POSIX-compliant storage, enabling scientists to leverage cloud-native tools while collaborating with backend storage users. A key driver for this work is also the need to share very big data with external collaborators who lack direct access to the internal storage systems. By providing an S3 interface to existing storage, we can facilitate secure and controlled data sharing across institutional boundaries, enabling researchers to work together more effectively.

We discuss the challenges and first solutions for mapping S3 objects to POSIX files, ensuring seamless data access and changes via both file system and S3. We look into existing open-source implementations like the Varsity S3 Gateway with the aim to provide read/write functionality on supported backend storage systems while addressing security considerations, user mapping, and compliance to the S3 standard. We will share our first experiences with running workflows of the reproducible analysis platform (REANA) on our initial setup as well as future directions for this project, which aims to bridge the gap between HPC and cloud storage, ultimately enhancing the productivity and collaboration of researchers.

Author: REDL, Robert (Ludwig-Maximilians-Universität in Munich, Faculty of Physics, Munich, Germany)

Co-authors: SACCHI, Elena (Leibniz-Institut für Astrophysik Potsdam (AIP), Potsdam, Germany); KHALATYAN, Arman (Leibniz-Institut für Astrophysik Potsdam (AIP), Potsdam, Germany); RAUSCHER, Felix (Ludwig-Maximilians-Universität in Munich, Faculty of Physics, Munich, Germany)

Presenter: REDL, Robert (Ludwig-Maximilians-Universität in Munich, Faculty of Physics, Munich, Germany)

Session Classification: HPC data access and integration

Track Classification: Main sessions: Scalable Storage Backends and Integration with Data Processing