



Contribution ID: 57

Type: **Invited talk**

openPMD – F.A.I.R. and open scientific I/O at the Exascale Era

Tuesday, 20 September 2022 11:30 (10 minutes)

This talk presents openPMD, an open and F.A.I.R. standard for particle-mesh data, and its impact in heterogeneous scientific workflows.

Particle accelerator codes need to span various time and length scales, leading to data processing pipelines consisting of multiple heterogeneous codes.

Standardization of physical data helps bridging the different models with a commonly-understood markup, creating interoperable and flexible workflows.

The openPMD standard is made accessible to scientific software via the openPMD-api, a library for the description of scientific data.

The backend implementations of the openPMD-api are based on established I/O frameworks such as HDF5 and ADIOS2, and also include a scalable streaming backend for HPC workflows, provided by ADIOS2.

The talk gives an insight into the existing ecosystem of openPMD and describes the basic concepts of the data markup.

It shortly illuminates recent trends in large-scale I/O and their impact on scientific compute workflows. While traditional attempts at counteracting such trends, e.g. through compression, remain available in the openPMD-api, we propose loose coupling and online analysis via streaming workflows as a sustainable solution that avoids parallel filesystem bottlenecks.

Primary author: POESCHEL, Franz (CASUS/HZDR)

Co-author: HUEBL, Axel (LBNL)

Presenter: POESCHEL, Franz (CASUS/HZDR)

Session Classification: Special Topic