F.A.I.R. Scientific I/O at the Exascale

Franz Poeschel (CASUS) | Axel Huebl (LBNL)

CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

www.openPMD.org
github.com/openPMD

Heterogeneity through Standardized Data

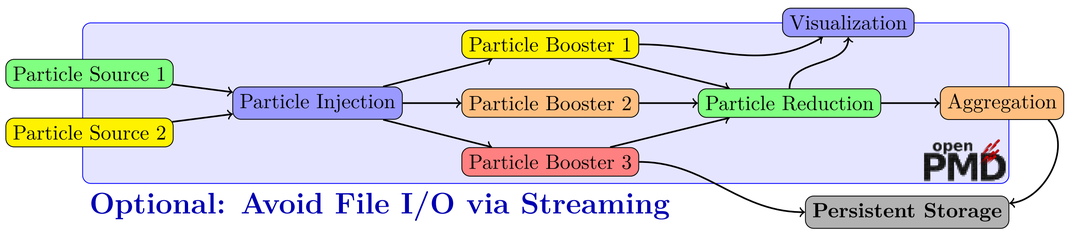

Particle Accelerators are complex:

- need to span different **time** and **length scales**
- particle accelerator modeling requires **multiples codes**, collaborating in a **data processing pipeline**
- **bridge heterogeneous models** by standardization of data

F.A.I.R I/O with openPMD

Findable: Standardized metadata to identify the data producer

```
string    /author              attr   = "franz"
string    /software            attr   = "PIConGPU"
string    /softwareVersion     attr   = "0.5.0-dev"
```

Accessible: Open standard, implementable in various formats

ADIOS · The HDF Group · *currently implemented, but not limited to

Interoperable: Data exchange spans applications, platforms and teams

Reusable: Rich and standardized description for physical quantities

Axel Huebl et al. "openPMD: A meta data standard for particle and mesh based data". 2015. doi: 10.5281/zenodo.591699. url: https://openPMD.org

File markup and definition:

openPMD standard (1.0.0, 1.0.1, 1.1.0)

Implementation and Language Binding:

openPMD-api[LBNL, CASUS, HZDR]:

- express **data description** in a C++/Python API **backend-agnostically**, **configuring the I/O backend at runtime**
- still use full functionality of underlying I/O libraries **(compression, aggregation, staging, strides, ...)** and their native tooling

file validators[HZDR, LBNL]

openPMD Ecosystem

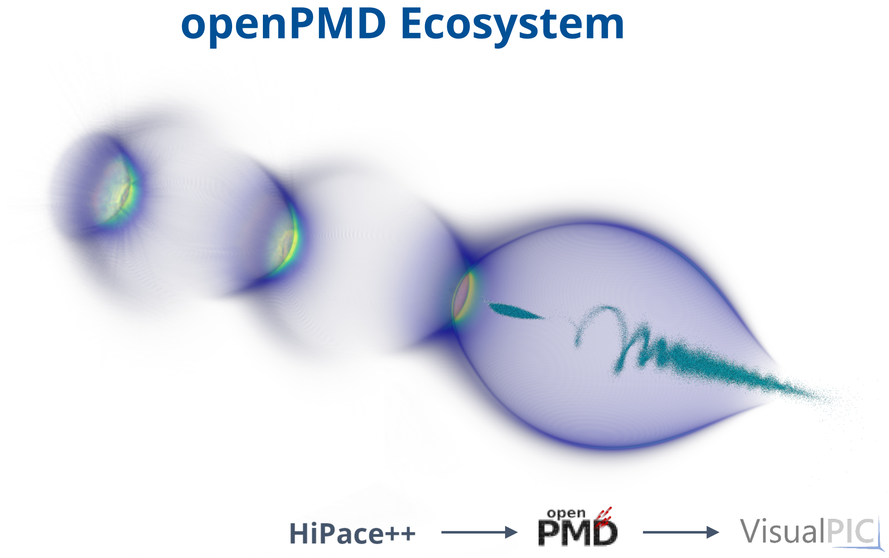

Data Processing and Visualization:

openPMD-viewer, VisIt, pyDive, postpic, yt project, ParaView, VisualPIC

Native tooling: HDF Compass, bpls

Open Simulations with openPMD:

Examples: PIConGPU[HZDR], HiPACE++[DESY, LBNL], SimEx Platform[EUCALL, European XFEL], BMAD[Cornell], Wake-T[DESY], FBPIC[LBNL, CFEL Hamburg University], WarpX[LBNL, DESY], ...

Full list: github.com/openPMD/openPMD-projects

Hierarchical organization of openPMD data

- **meta-standard:** truly self-describe data
- **open-access:** unified description
- **workflows:** high-level integrations

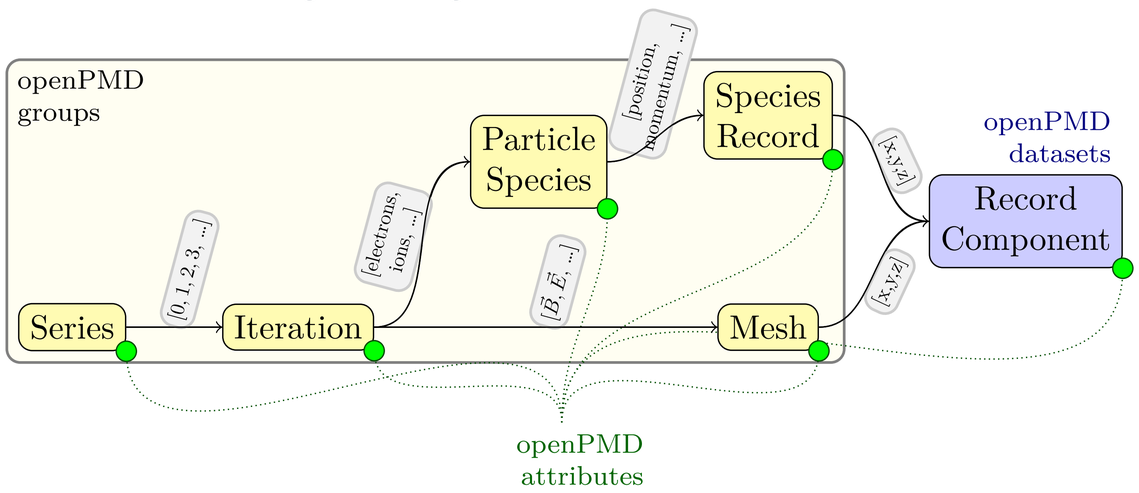

- Particle and mesh based data
- data format agnostic
- frictionless data exchange

Scientific workflows need to bridge various applications and algorithms, ideally both **automatically-** und **human-readable**.

openPMD defines **scientific self-description**, usable in common storage/transport formats such as HDF5, ADIOS, JSON.

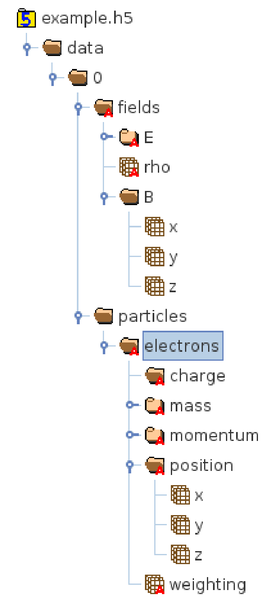

Compute Performance Outpaces Storage Performance

	Titan	Summit	Frontier
Peak Performance:	27 Pflop/s	200 Pflop/s	1.6 Eflop/s
FS Throughput:	1 TiByte/s	2.5 TiByte/s	5~10 TiByte/s
FS Capacity:	27 PiByte	250 PiByte	500~1000 PiByte

Break through Filesystem Bandwidth with Streaming: >2.5TiByte/s

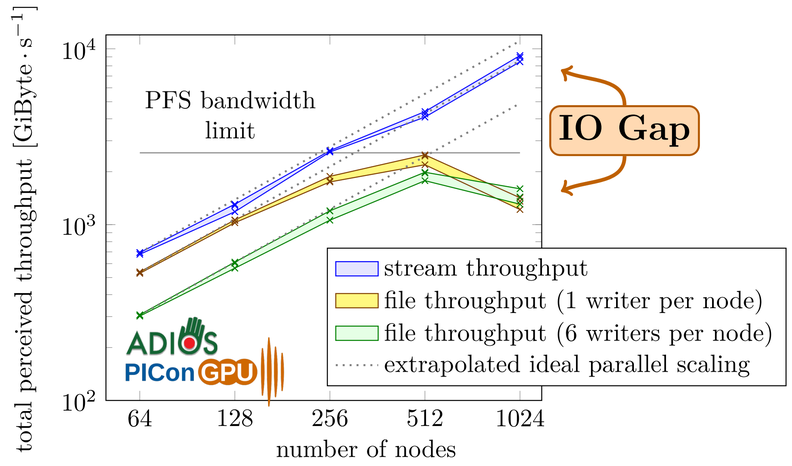

Franz Poeschel et al. "Transitioning from file-based HPC workflows to streaming data pipelines with openPMD and ADIOS2". 2022. doi: 10.1007/978-3-030-96498-6_6.

SPONSORED BY THE Federal Ministry of Education and Research

SACHSEN This project is co-financed by the Saxon State government out of the State budget approved by the Saxon State Parliament