


Franz Poeschel (CASUS/HZDR), Axel Huebl (LBNL), Lipeng Wan (GSU), Remi Lehe (LBNL), Norbert Podhorszki (ORNL), Junmin Gu (LBNL), Maxence Thévenet (DESY), Erik Schnetter (PITP), Michael Bussmann (CASUS/HZDR)

www.openPMD.org
github.com/openPMD

Heterogeneity through Standardized Data



Optional: Avoid File I/O via Streaming

Scientific compute workflows are complex:

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

HELPMI – a project by the HMC

- explore a F.A.I.R. standard for laser-plasma experiments
- project from April 2023 to April 2025

openPMD's role

- background: LPA simulations
- HELPMI aims to close the gap between openPMD and NeXus
- aim for interoperability between openPMD + HELPMI + NeXus

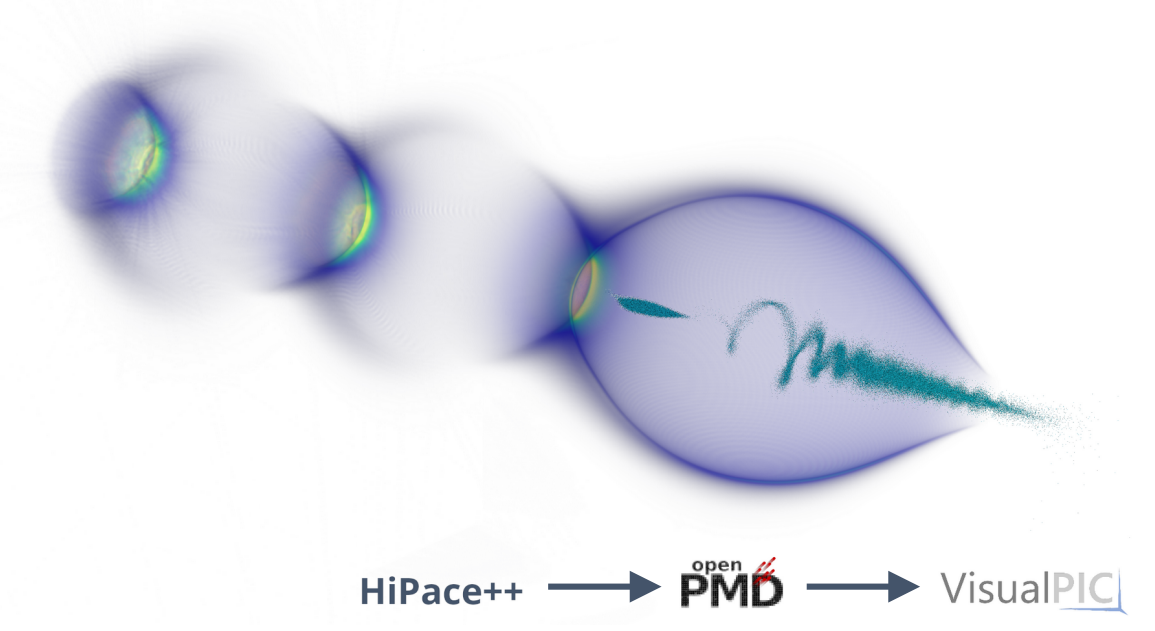


openPMD Ecosystem


File markup and definition:
openPMD standard (1.0.0, 1.0.1, 1.1.0)

Ref. Implementation and Bindings:
openPMD-api^{LBNL, CASUS, HZDR}:

- express **data description** in a C++/Python/(dev) Julia 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



- open development & tests: github.com/openPMD/openPMD-api
- available in common package managers and deployed on HPC systems



Data Processing and Visualization:

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

Open Simulations with openPMD:

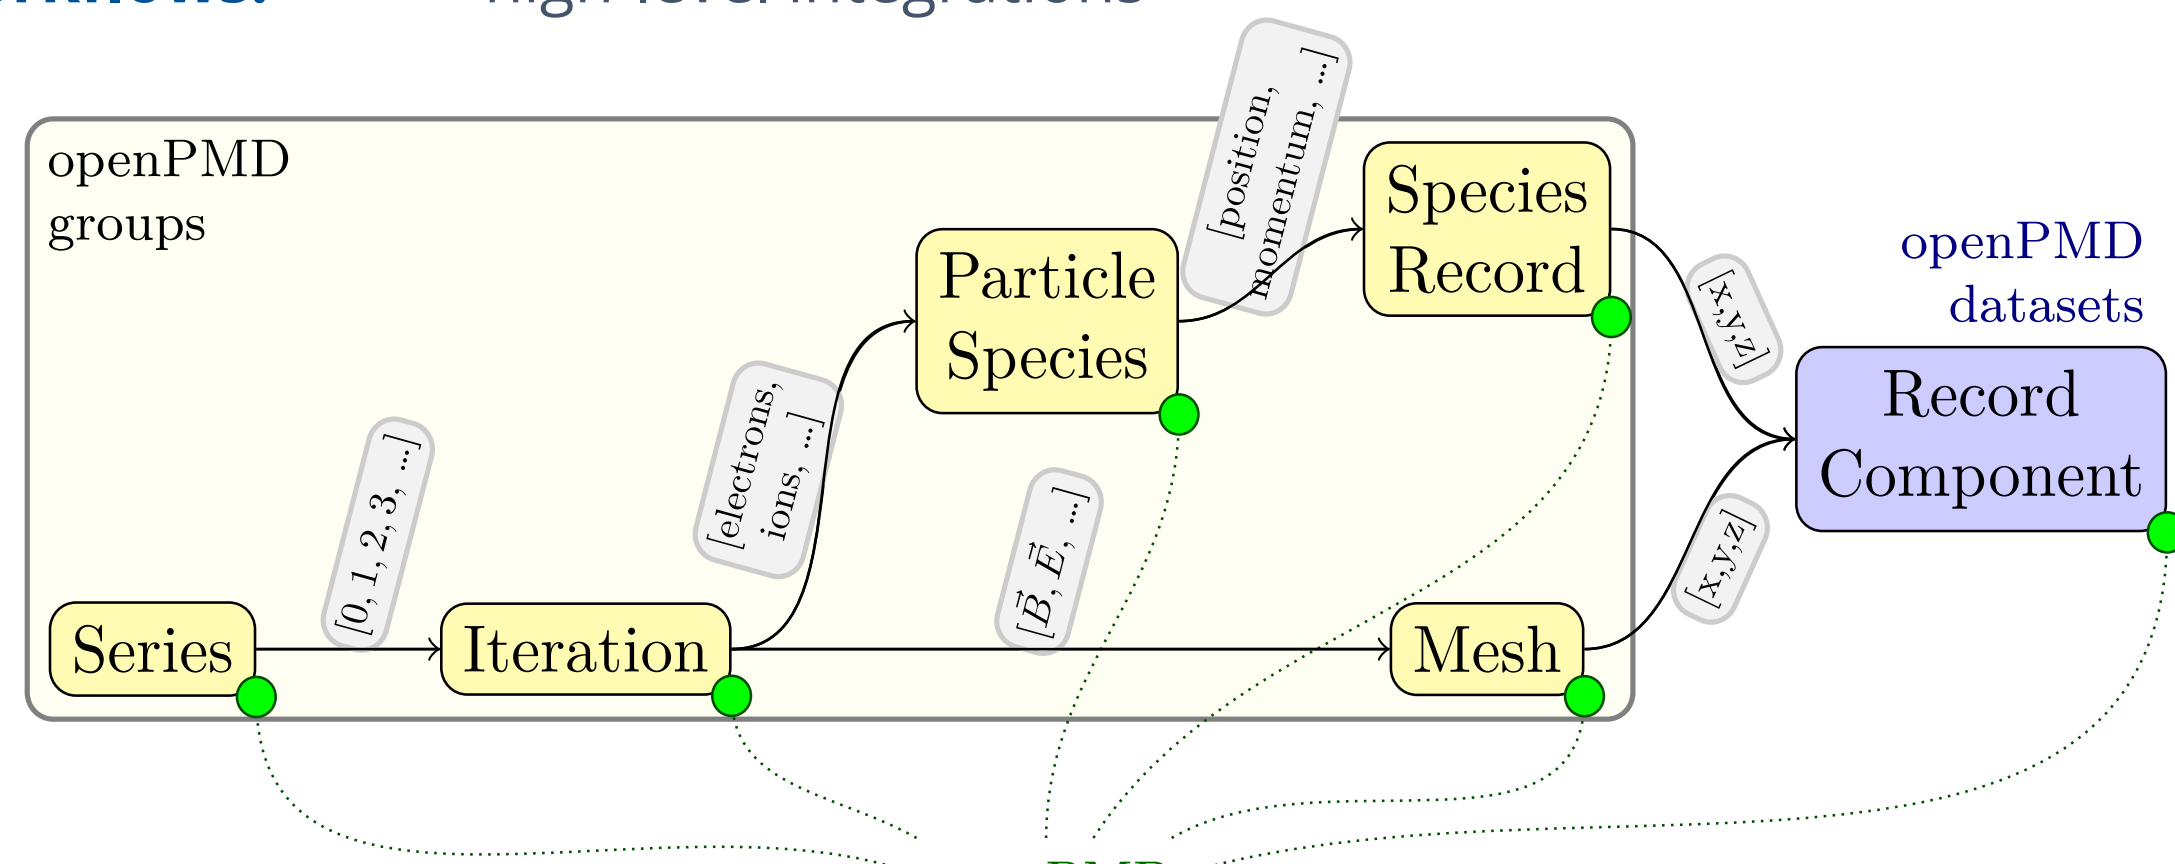
Examples: PIConGPU^{HZDR}, HiPACE++^{DESY, LBNL}, MALA^{CASUS}, BMAD^{Cornell}, Wake-T^{DESY}, SimEx Platform^{EUCALL, European XFEL}, FBPIC^{LBNL, CFEL Hamburg University}, WarpX^{LBNL, DESY}, ...

Native tooling: HDF Compass, bpls
Integration into: Pandas, DASK, RAPIDS
file validators^{HZDR, LBNL}

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

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

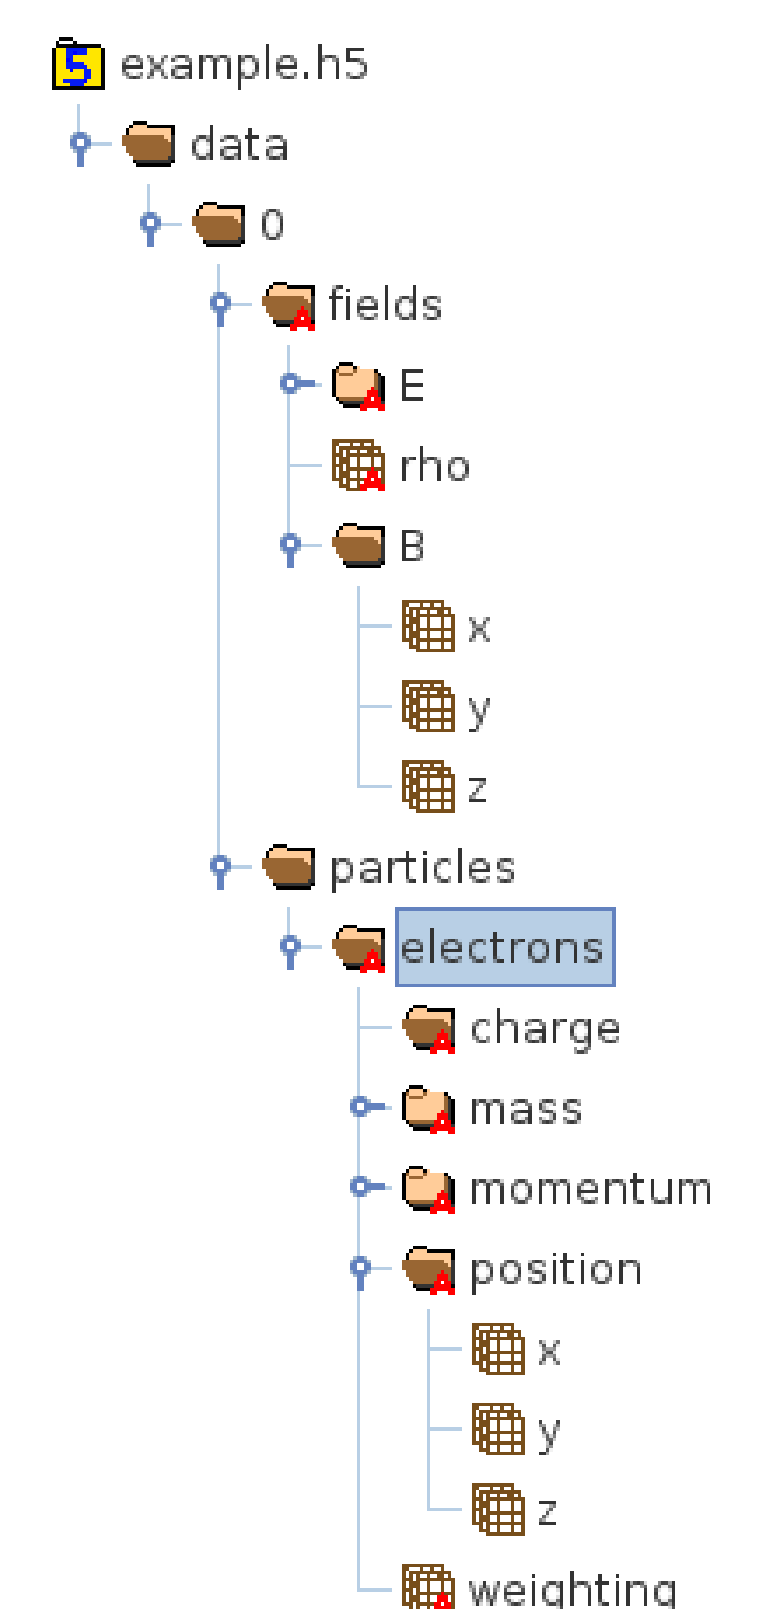
ADIOS The HDF Group *currently implemented, but not limited to

Interoperable: Data exchange spans applications, platforms and teams

Reusable: Rich and standardized description for physical quantities

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

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



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>

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

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.

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

