High Performance and Scientific Computing for FAIR

Thorsten Kollegger

Facility for Antiproton and Ion Research
& GSI Helmholtzzentrum

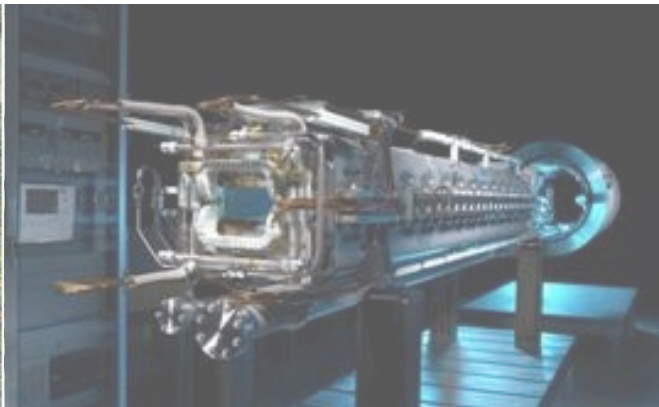

FAIR Scientific Pillars

APPA

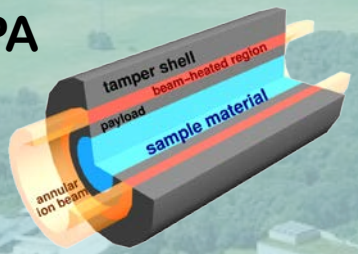

Atomic, Plasma Physics and Applications

CBM

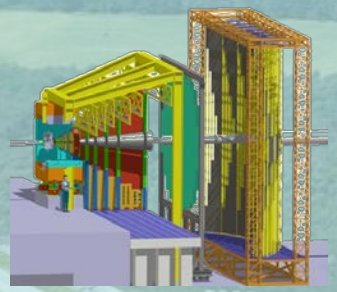

Compressed Baryonic Matter

PANDA

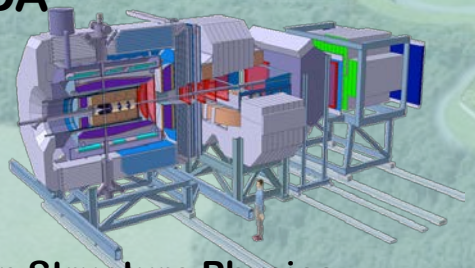

Hadron Structure Physics

NUSTAR

Nuclear Structure, Astrophysics and Reactions

Computing at FAIR: Big Data

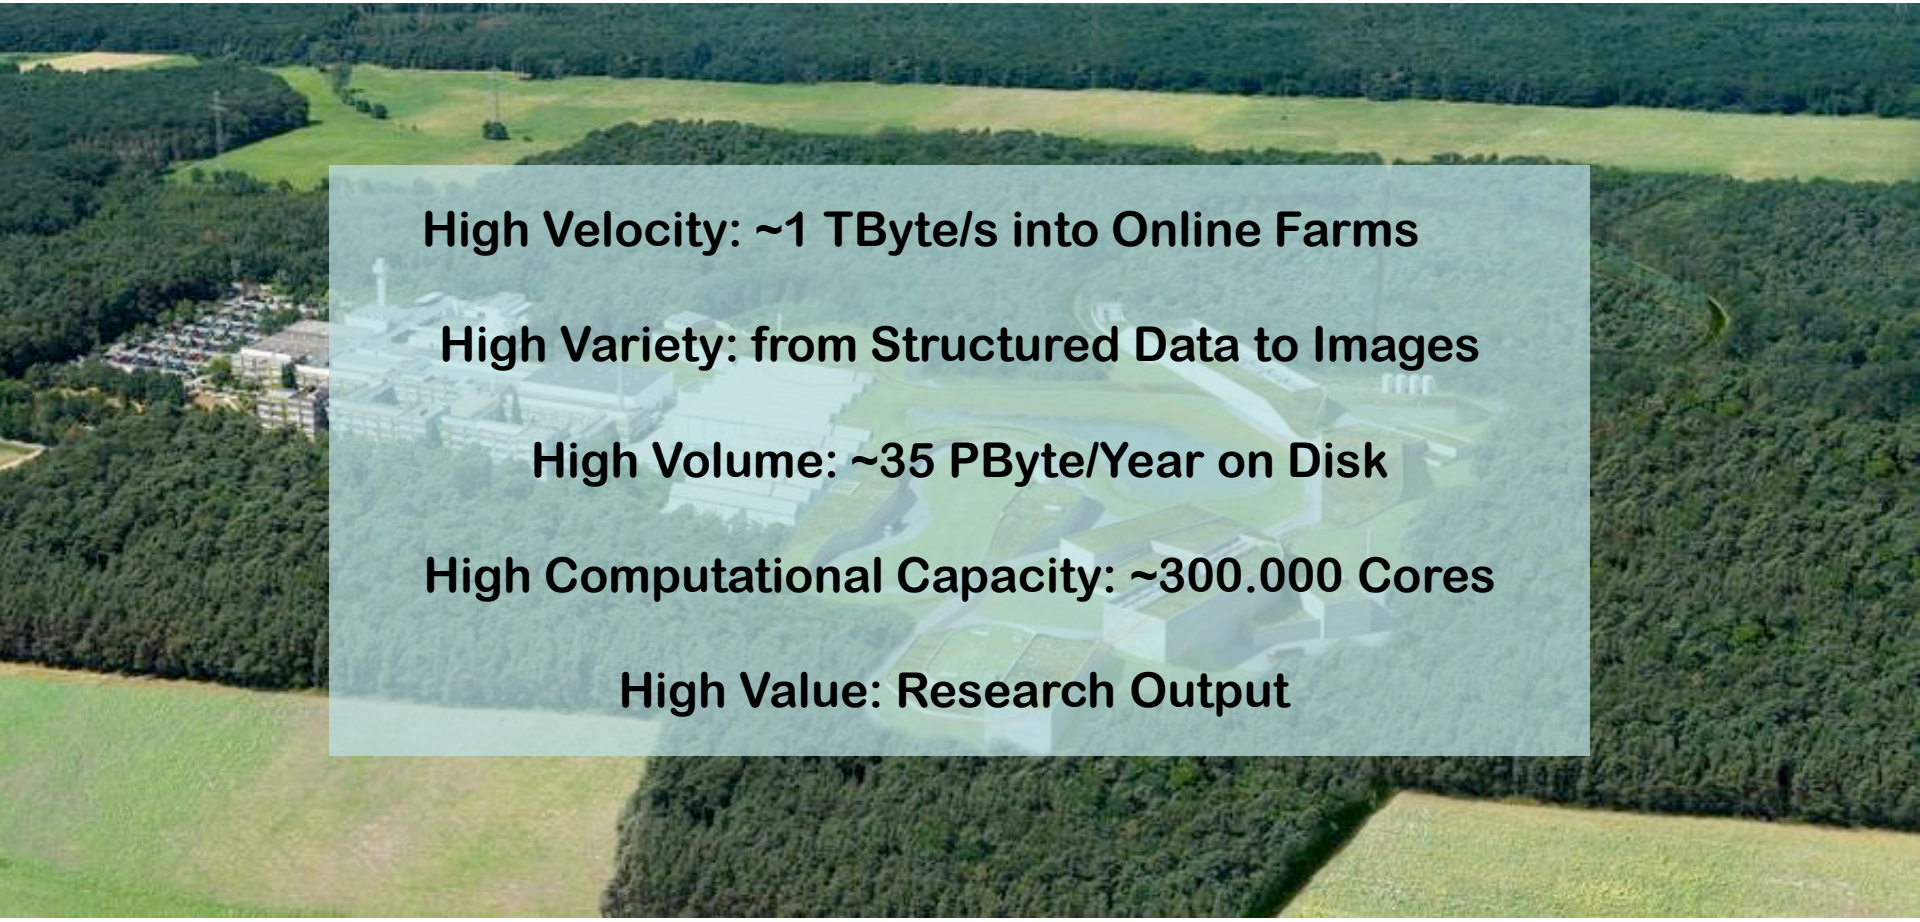

High Velocity: ~1 TByte/s into Online Farms

High Variety: from Structured Data to Images

High Volume: ~35 PByte/Year on Disk

High Computational Capacity: ~300.000 Cores

High Value: Research Output

Computing at FAIR

„Data is inherently dumb.
Algorithms are where the real values lies“
Peter Sondergaard, Gartner, 2015

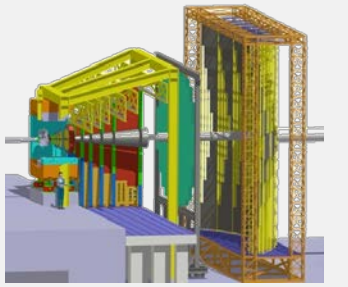

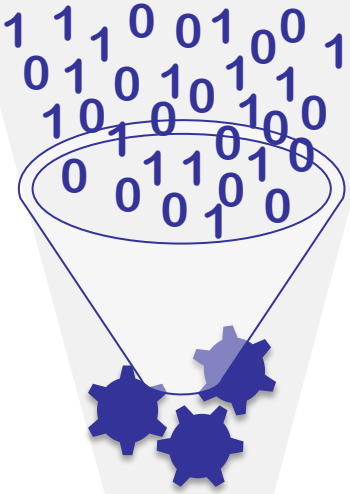

More complex algorithms necessary
to exploit the full scientific potential
of the instrumentation
(i.e. accelerators, detectors)

Fundamental change in detector design

From custom electronic to commodity
computing systems (“triggerless”)

- Software defined -> Agility
- Huge data rates into (quasi-)real time analysis
- FAIR, LHC upgrades, Nuclear Physics …

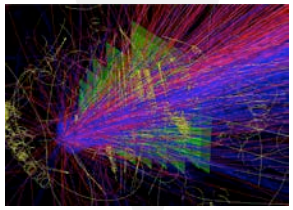

Computing at FAIR

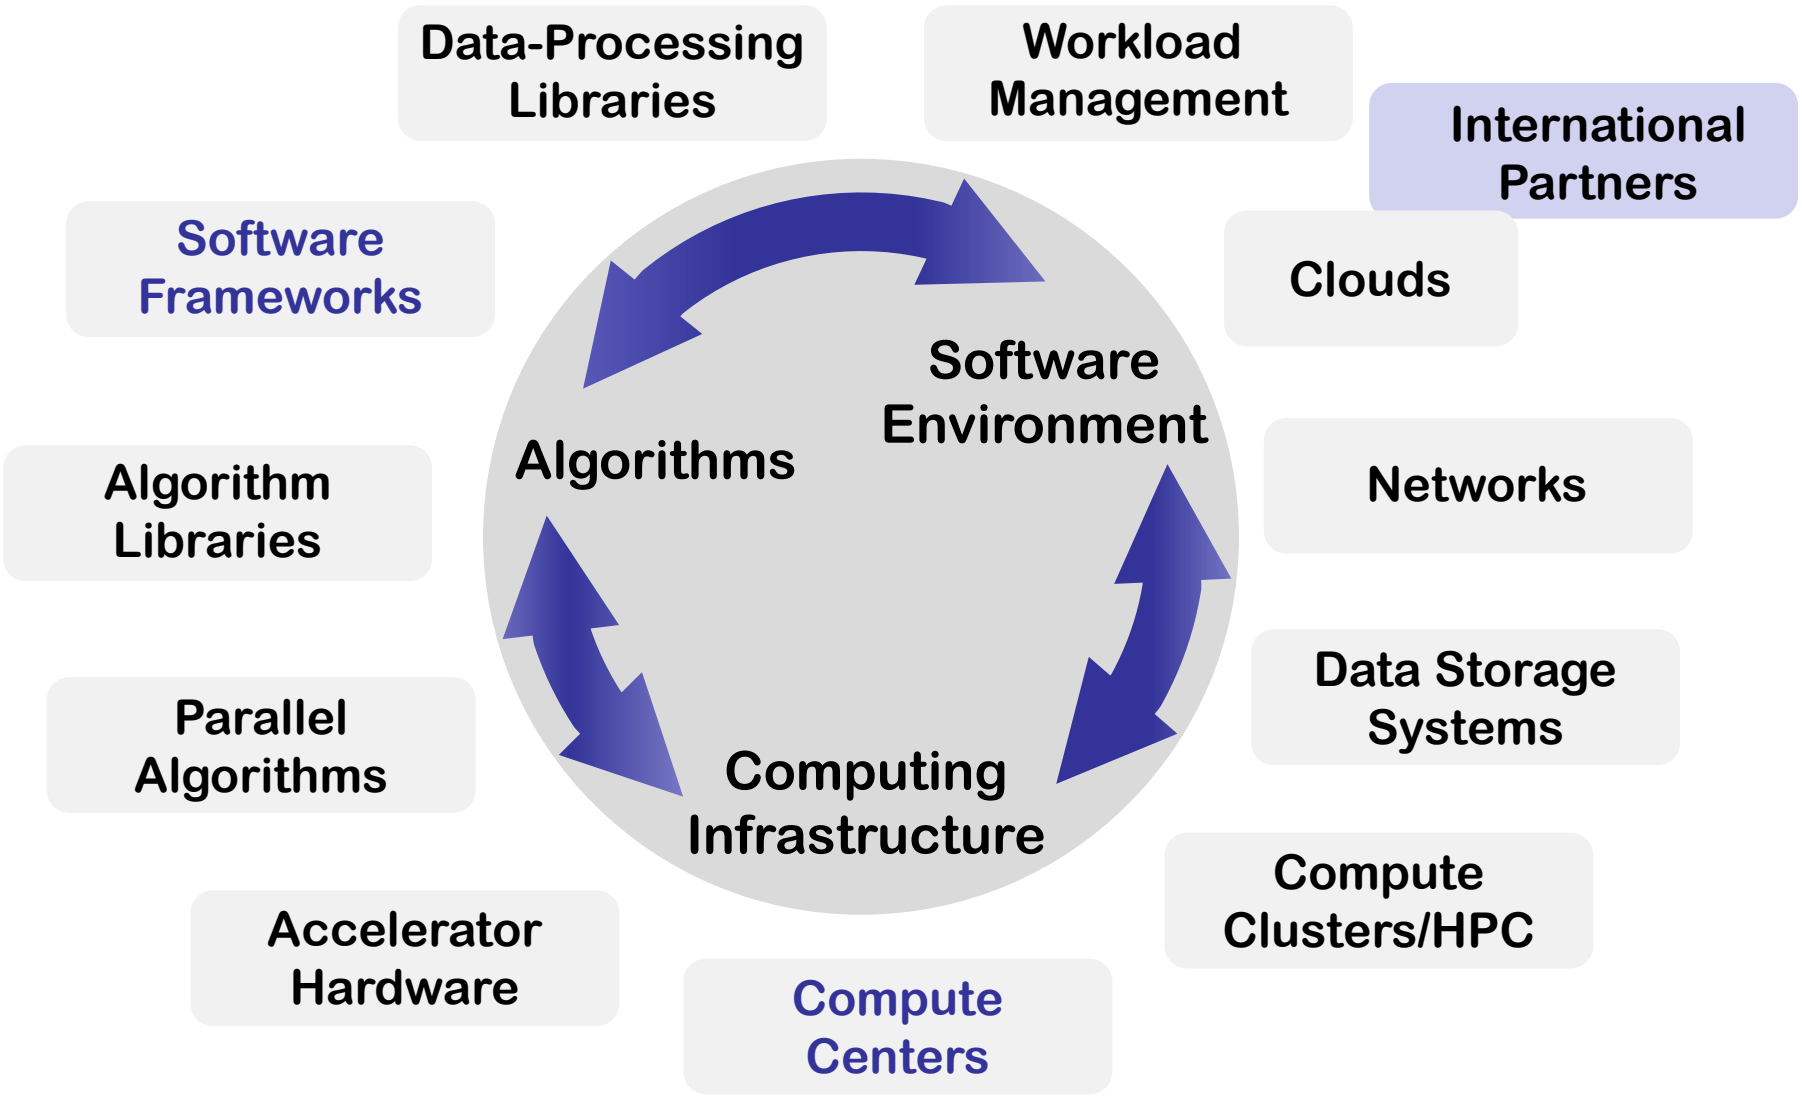

GSI/FAIR Green Cube

12 MW common data center for FAIR and GSI

- In operation since Feb ‘16

PUE < 1.07

- Very good Power Usage Efficiency (PUE)
- Reduced CO_2 Emission
- Low Operating Costs

More detailed information:
Jan Trautmann, HEPIX Spring 2016
https://indico.cern.ch/event/466991/contributions/1143585/

GSI/FAIR Green Cube

Constructed:	**Dec ‘14 – Nov ’15**
Building:	**6 Floors, 4.645 sqm** **768 19” racks** **(256 racks in 1st stage)**
Cooling & Power:	**12 MW (4 MW in 1st stage)** **PUE < 1.07 (Design)** **PUE ~ 1.04 (Commissioning)** **Water cooled** **- Passive rear-door heat-exchanger** **- Evaporative cooling towers** **N+1 Redundancy**
Cost:	**16 M€ (1st stage: 11.5 M€)**

Green Cube Cooling Concept

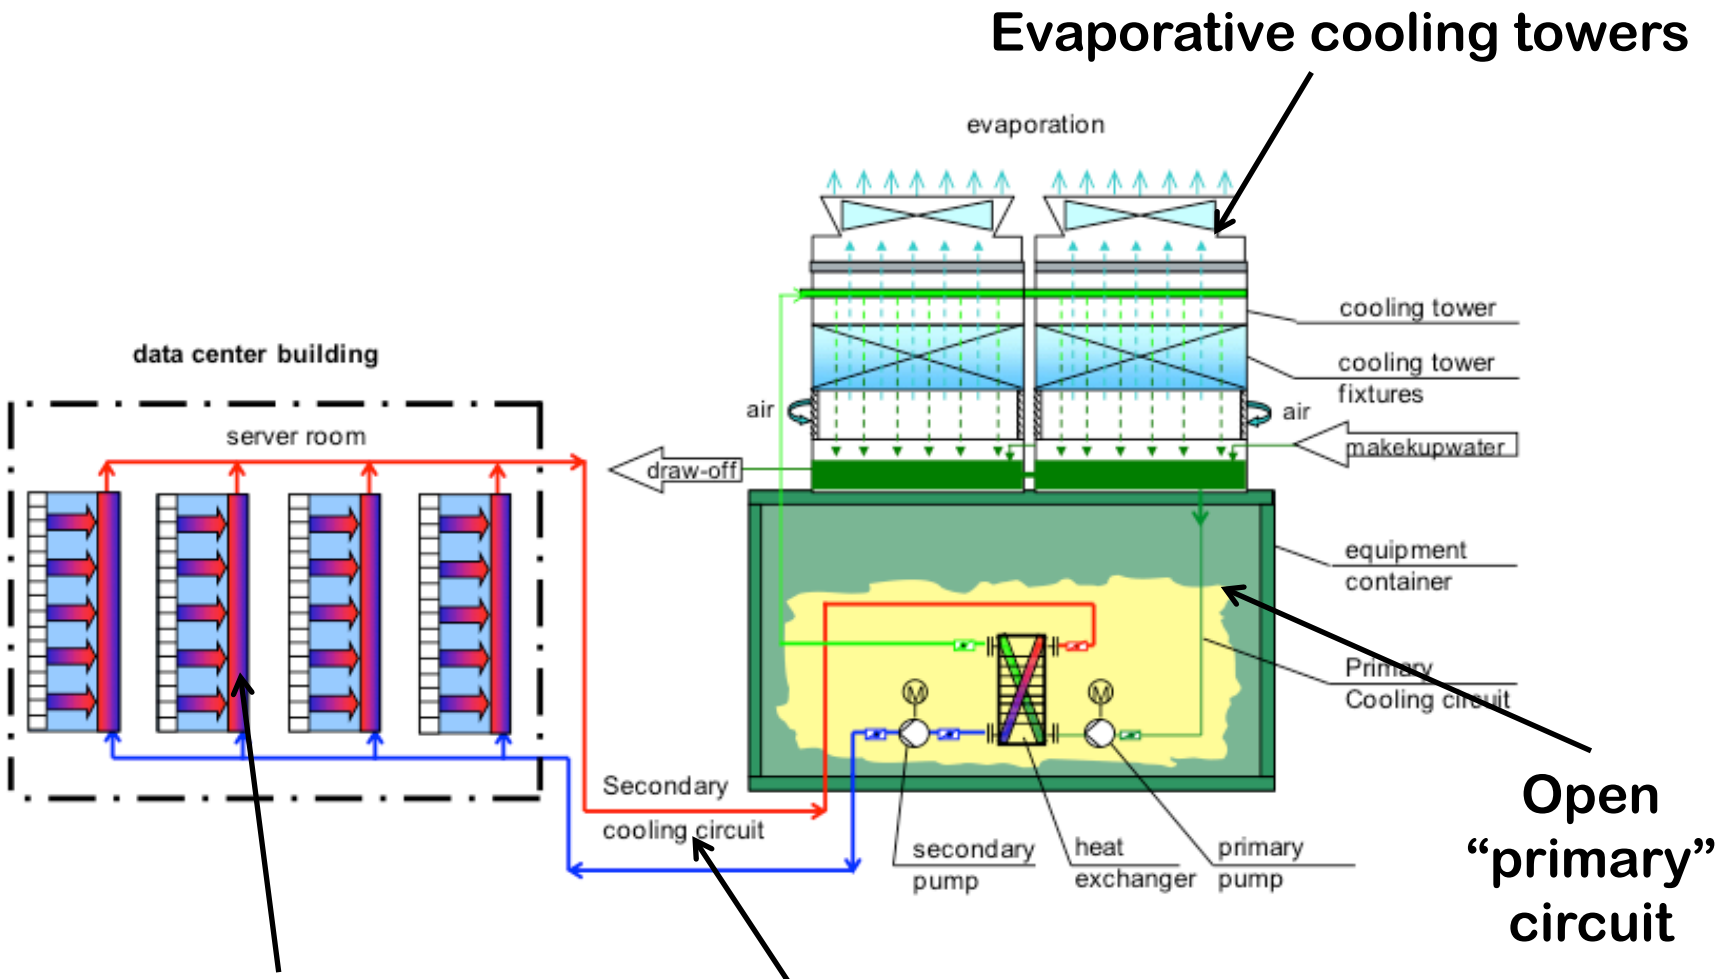

Green Cube Buildings

Cooling Towers

GT1 (Cooling)

GC 1 (Computing)

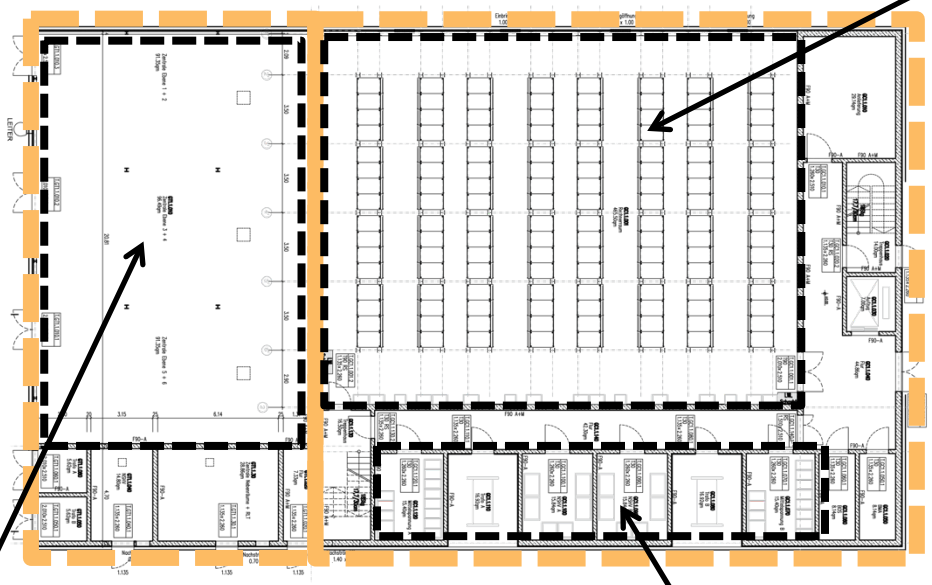

Computing Room

Cooling Circuits Pumps & Heat Exchangers

Electrical Power

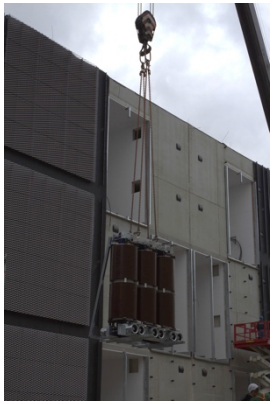

Green Cube CG1

Green Cube CG1

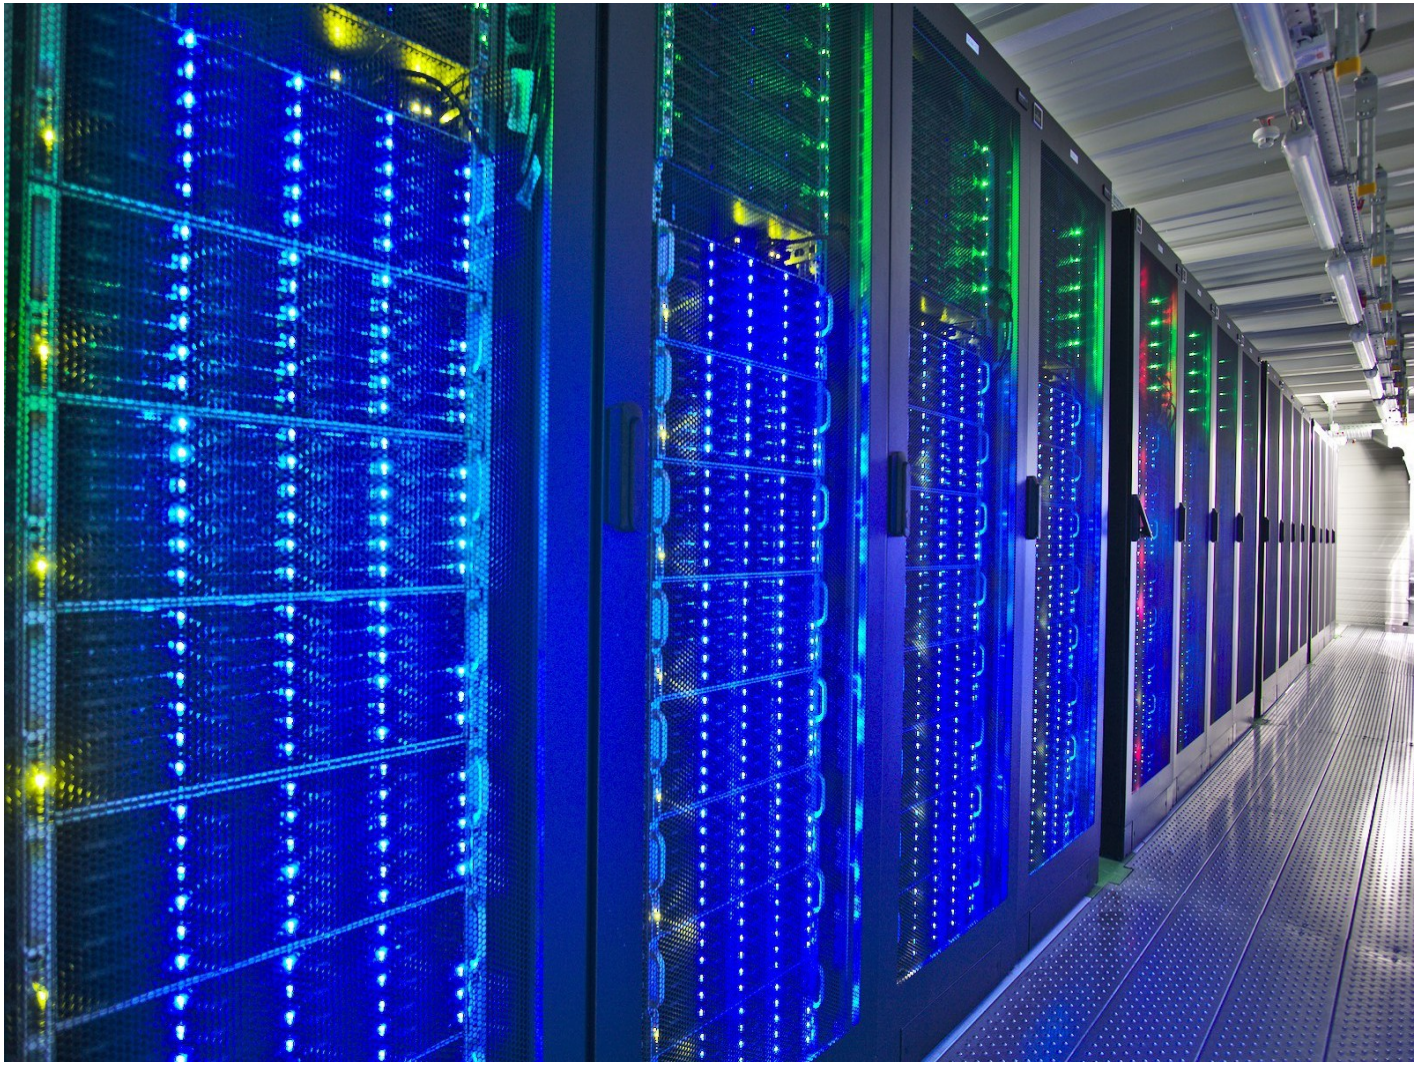

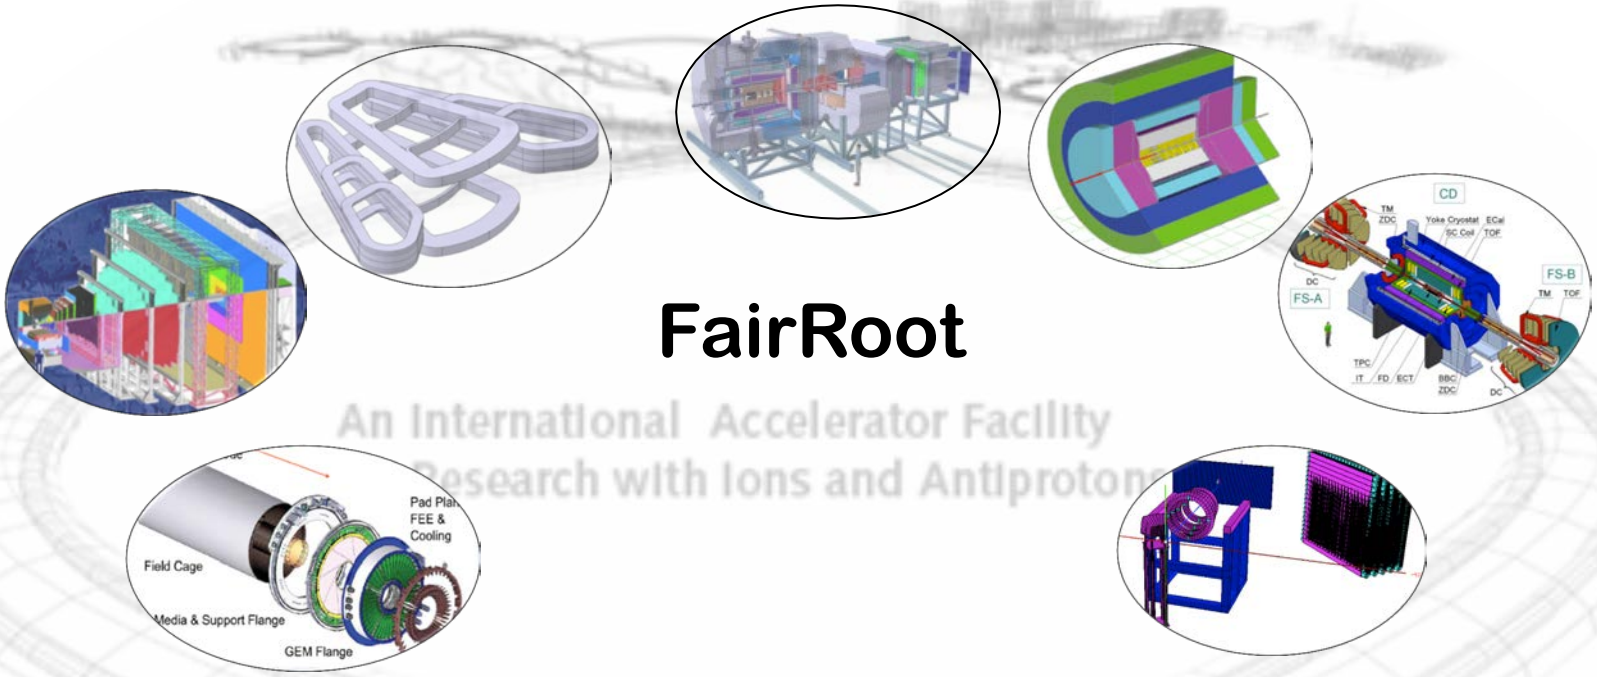

Common simulation, reconstruction and analysis software framework for the FAIR experiments (and beyond)

https://github.com/FairRootGroup/FairRoot

FairRoot

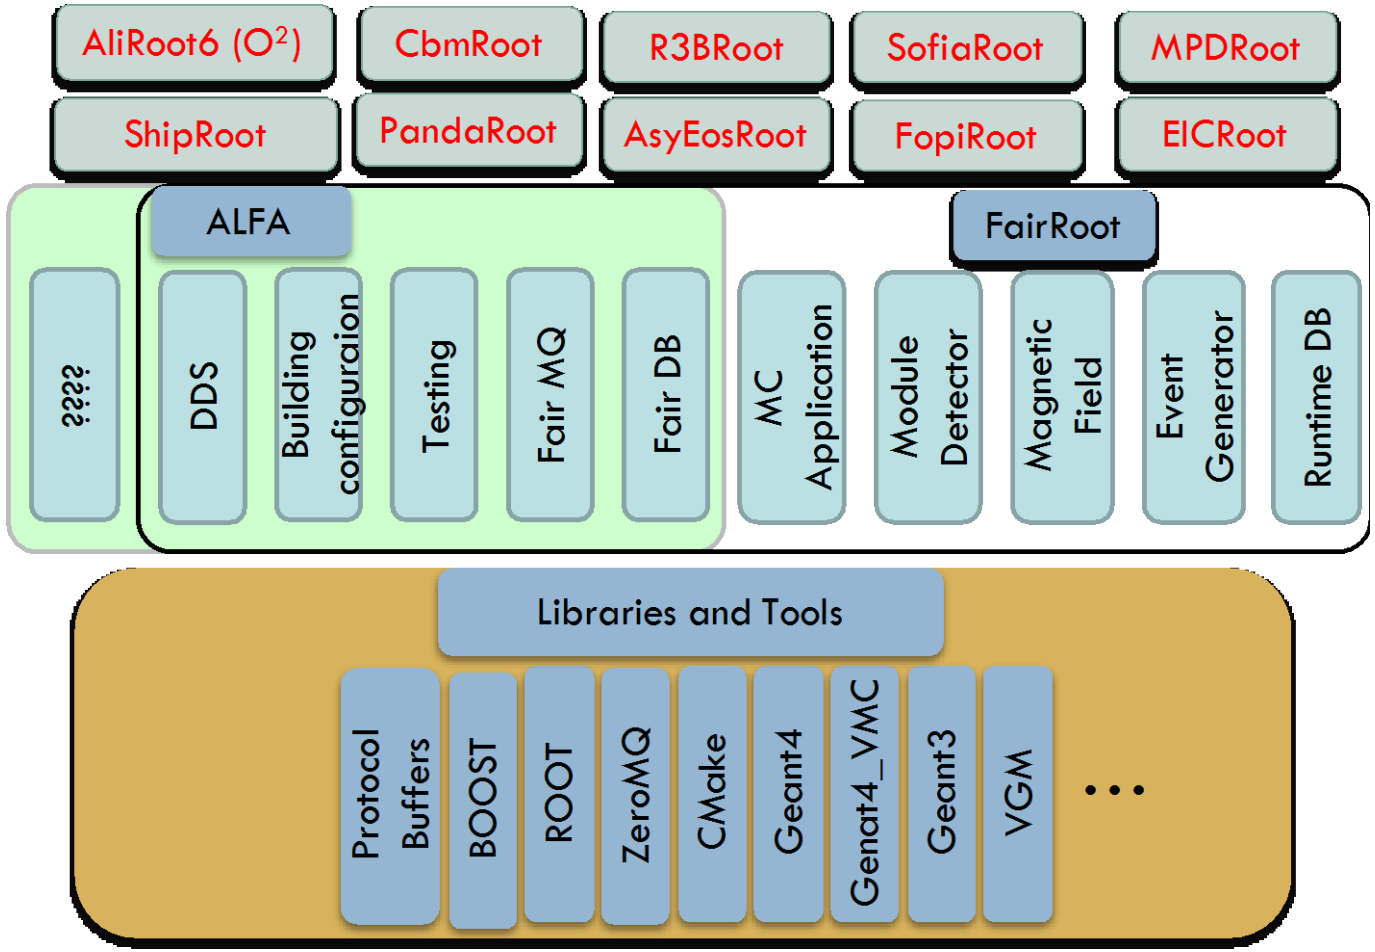

Modularized, LGPLv3 licensed common software stack for the FAIR experiments and others, development steered by GSI

FairMQ

The Data Processing Component of FairRoot

- Multi-process concept (specialized devices)
- Data-flow model: Message queues for data exchange, technology agnostic

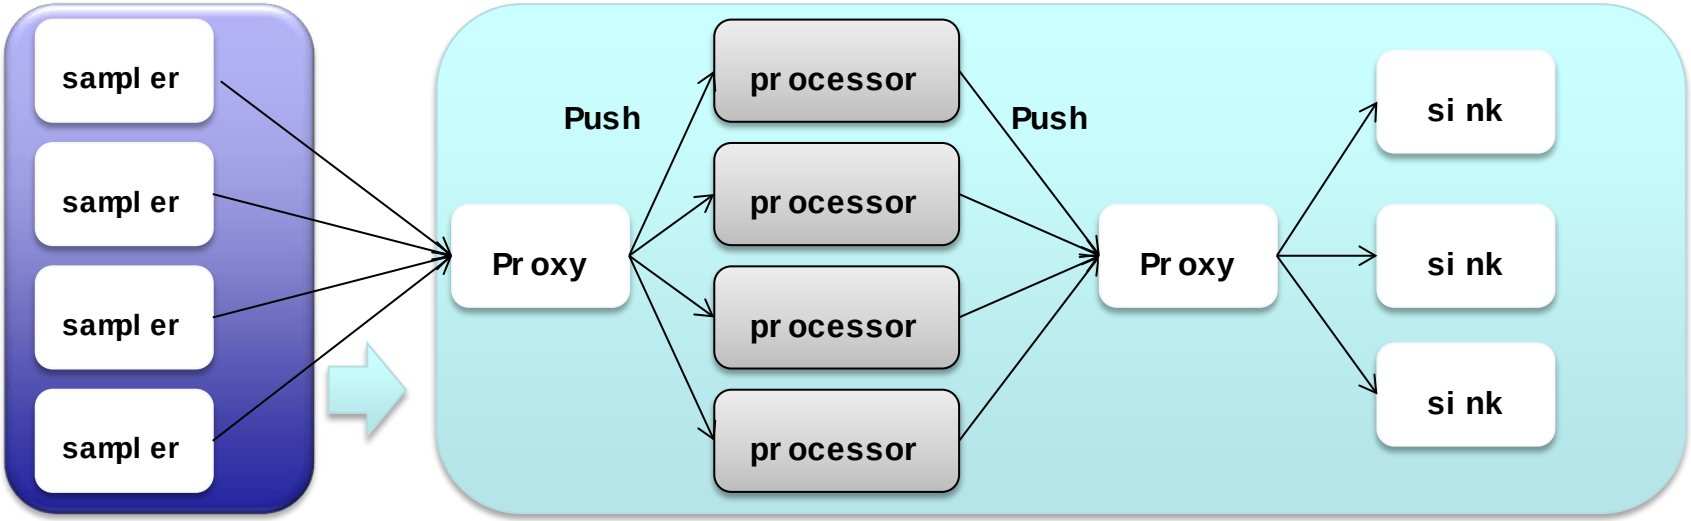

Design Goals

- Scalability, Maintainability, Reliability
  - efficient use of multi-core architectures
- Reusable with common data processing components
  - Reduce cost of new developments, agile development

FairMQ

Looking at the IT landscape: shift towards

- Microservices
  - Unbundled, decentralized modules
  - Organized around specific capability
- Containers
- Algorithm Economy

These are at the heart of the
„cloud/app“ business model/economy

- driven by scalability and reliability demands
- based on multi-process and message exchange
- development cost advantage

FairMQ uses many of these technologies under the hood;
replacing custom code (e.g. ALICE HLT framework)

Example Application: PANDA MVD Pixel Detector Prototype

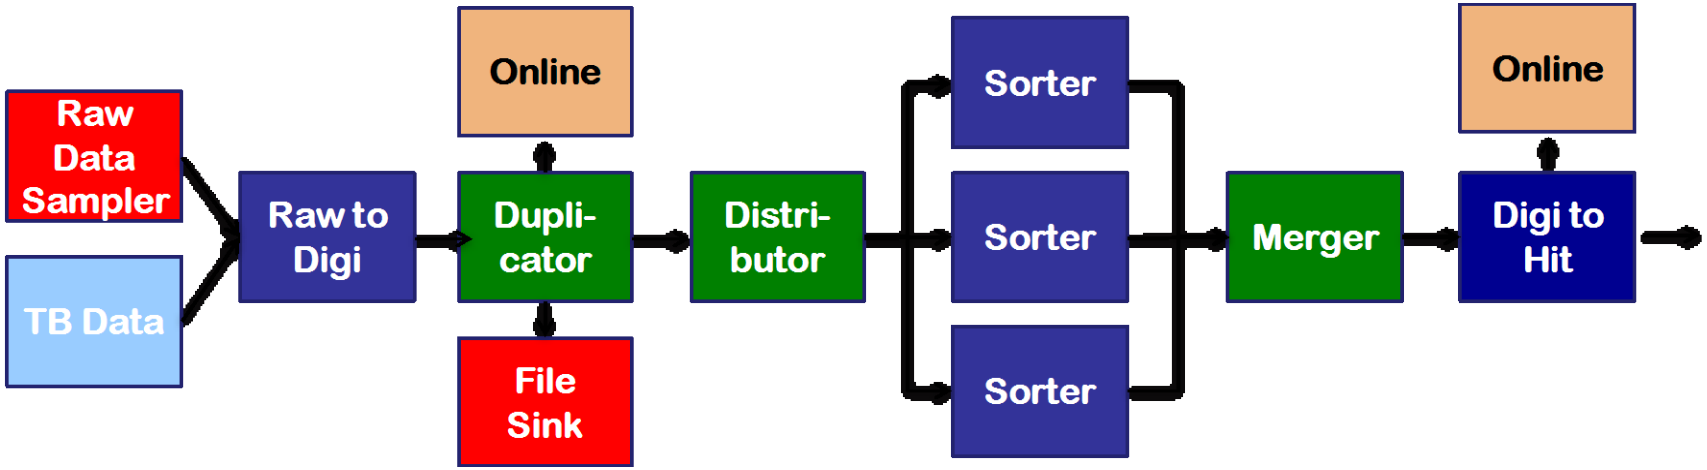

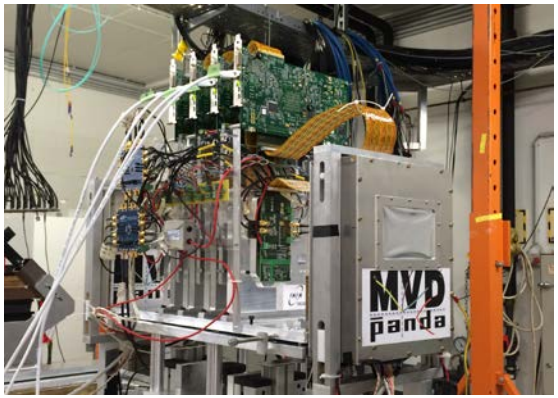

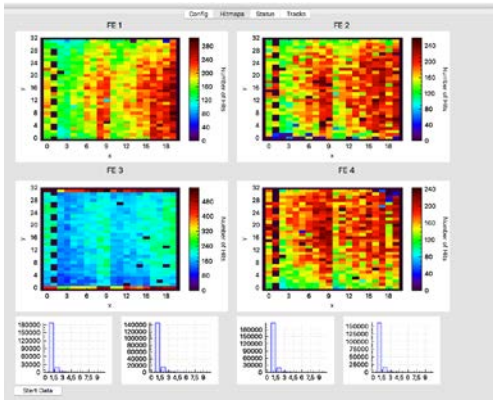

Tobias Stockmanns (FZJ, PANDA)

FAIR Computing

Driven by needs of FAIR experiments
for online reconstruction

- ~1 TByte/s into online farms
- ~300.000 cores (majority on-site in common compute center)
- 35 PByte/"year" disk
- 30 PByte/year tape

Algorithms and software development

- common frameworks & libraries, e.g. FairRoot