Analysis at FAIR / GSI

M. Al-Turany, V. Friese, D. Kresan, T. Stockmanns

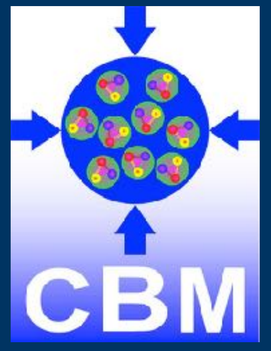

HSF DAWG Meeting, May 5 2021

FAIR Facility under construction

https://fair-center.de/

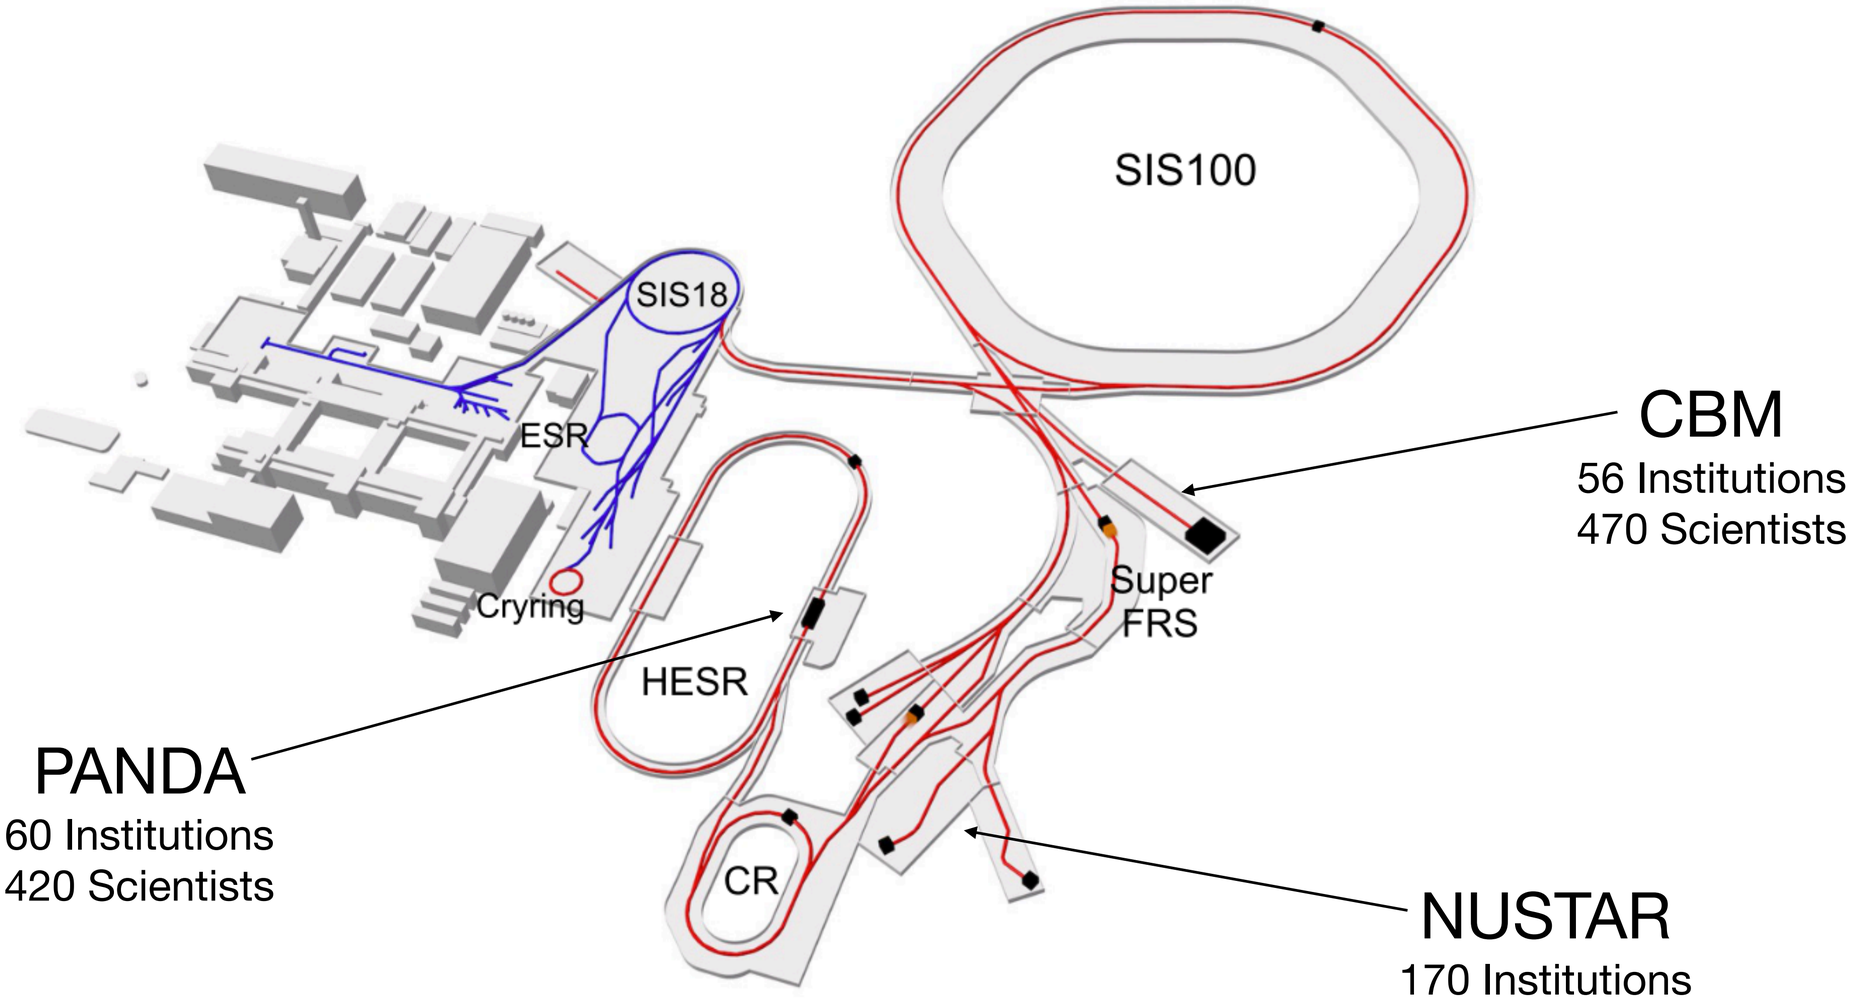

Four scientific pillars of FAIR

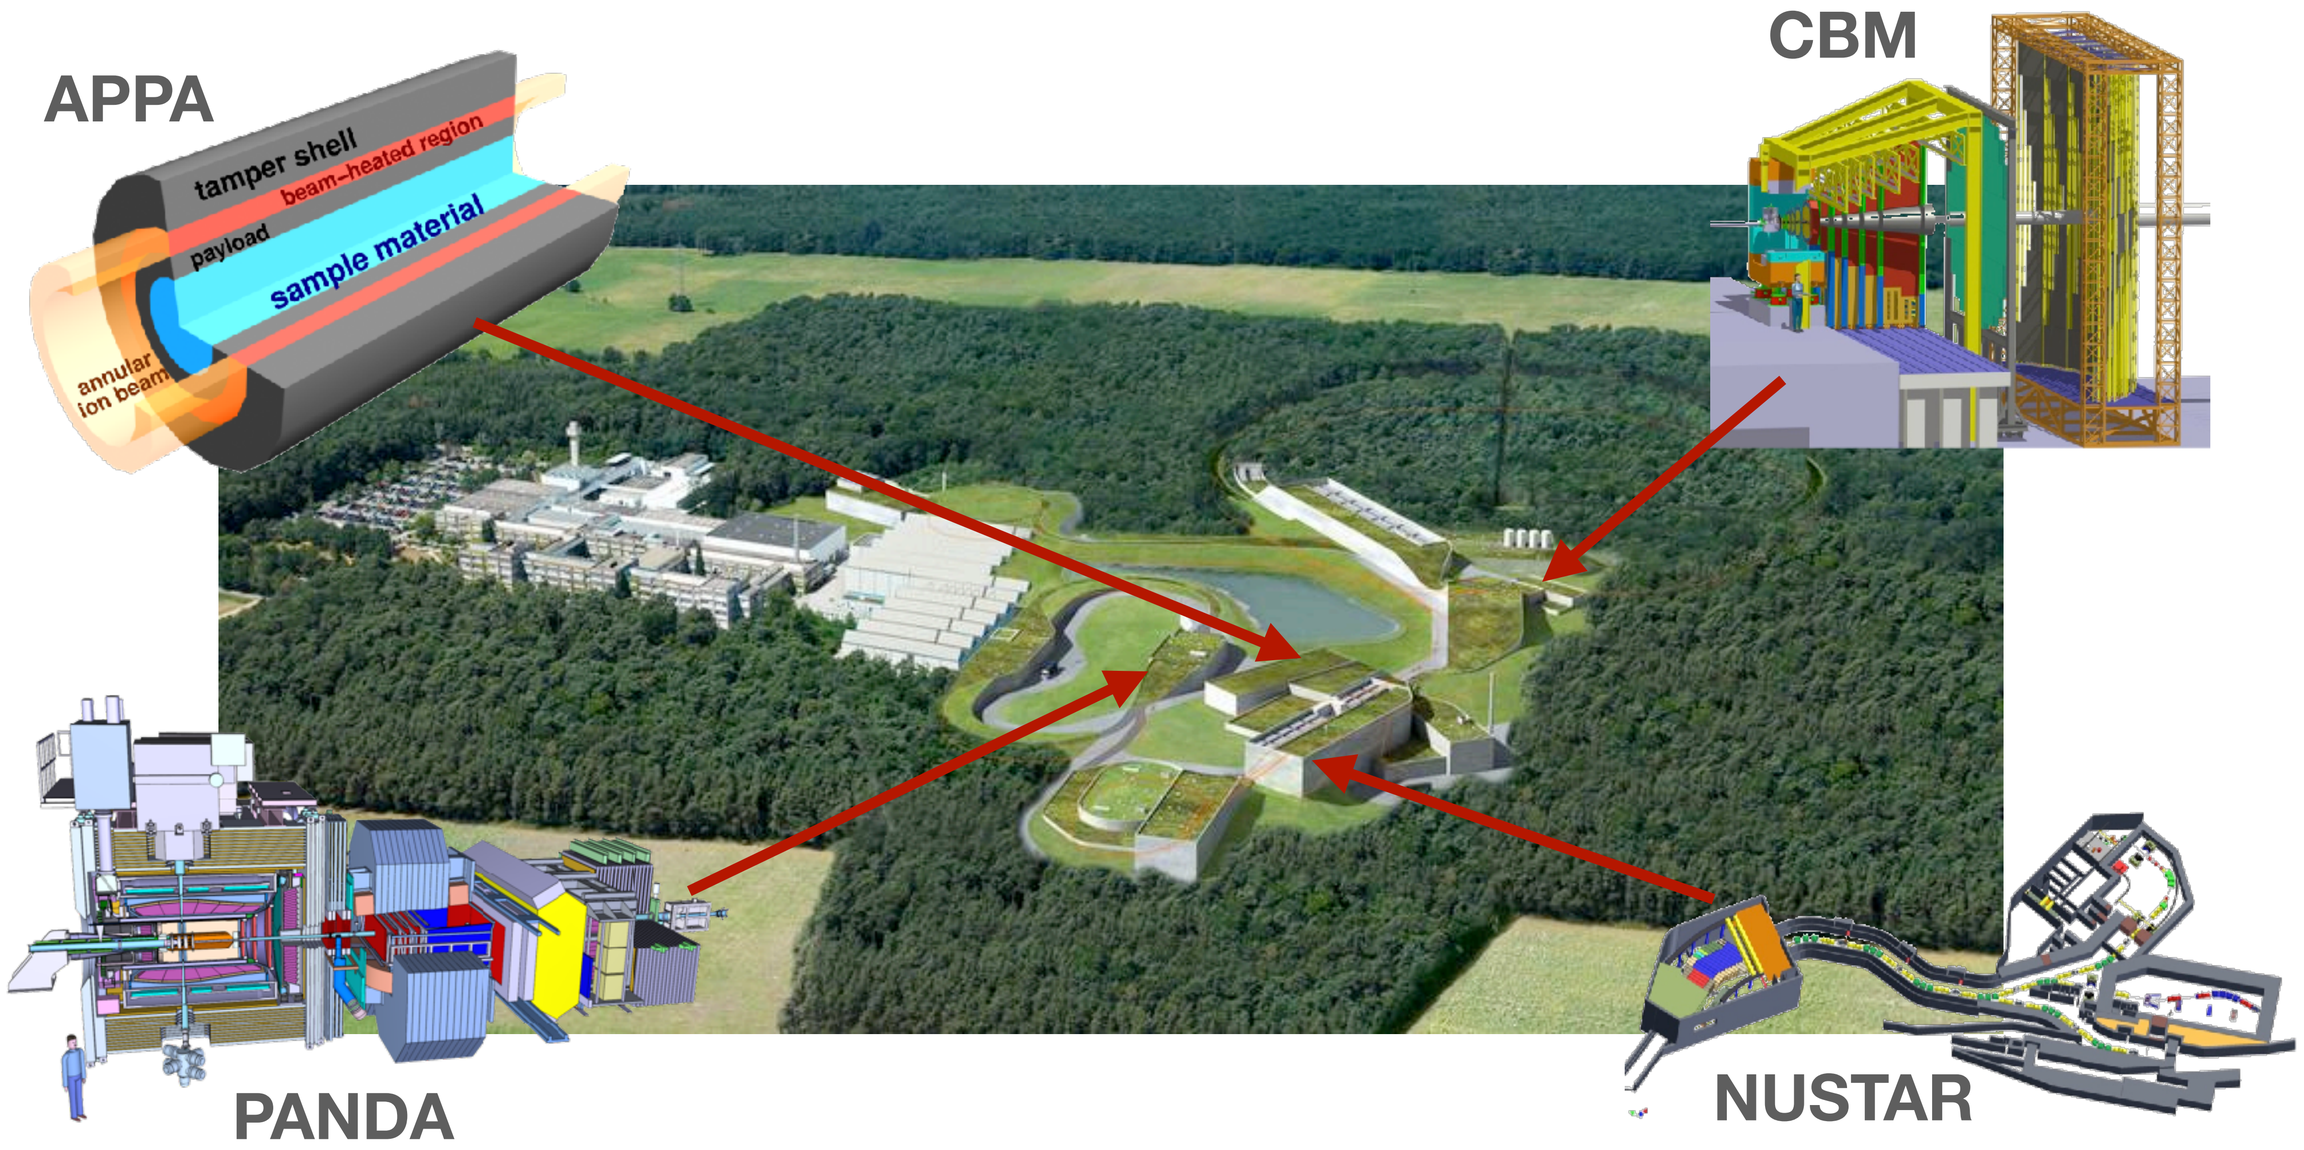

Computing at FAIR / GSI

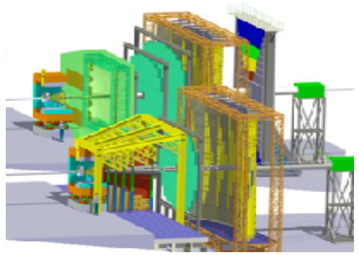

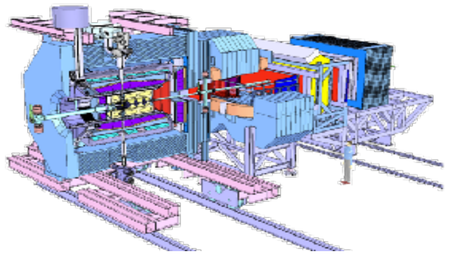

Analysis Facilities

ALICE

- Provide computing infrastructure for different experiments
- ALICE Analysis facility

Generic batch farm for GSI/FAIR Users

CBM

R3B

panda

HADES

AGATA
ADVANCED GAMMA TRACKING ARRAY

BIOMAT

APPA

CBM

Compressed Baryonic Matter

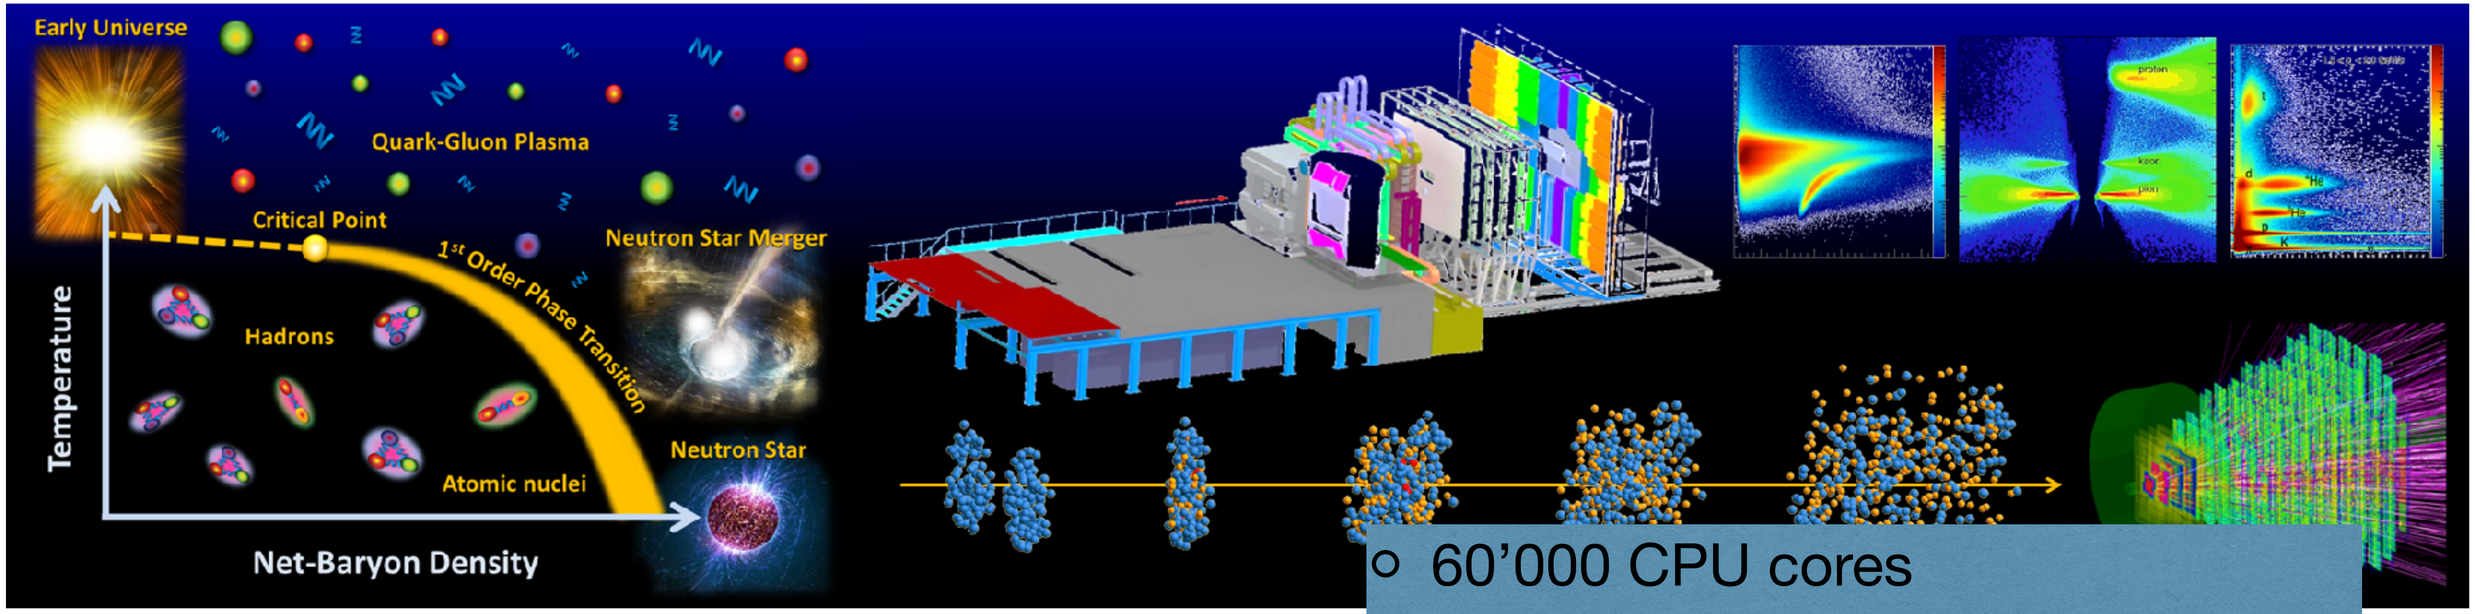

- Extensive detector simulations with Geant3 / Geant4
- Online reconstruction in trigger-less mode

- 60’000 CPU cores
  - To perform online a full event reconstruction on the 1 TB/s input data stream
- ? GPUs
  - To speed up the reconstruction

PANDA

Proton ANti-proton in DArmstadt

- Highlights:
  - Mixed background and signal simulation
  - Time-based simulation (overlapping events + pile-up)
  - Reconstruction of low momentum tracks
  - Full online event reconstruction

- Current development activities
  - Online tracking with GPUs
  - ML based tracking and PID
  - Online event building

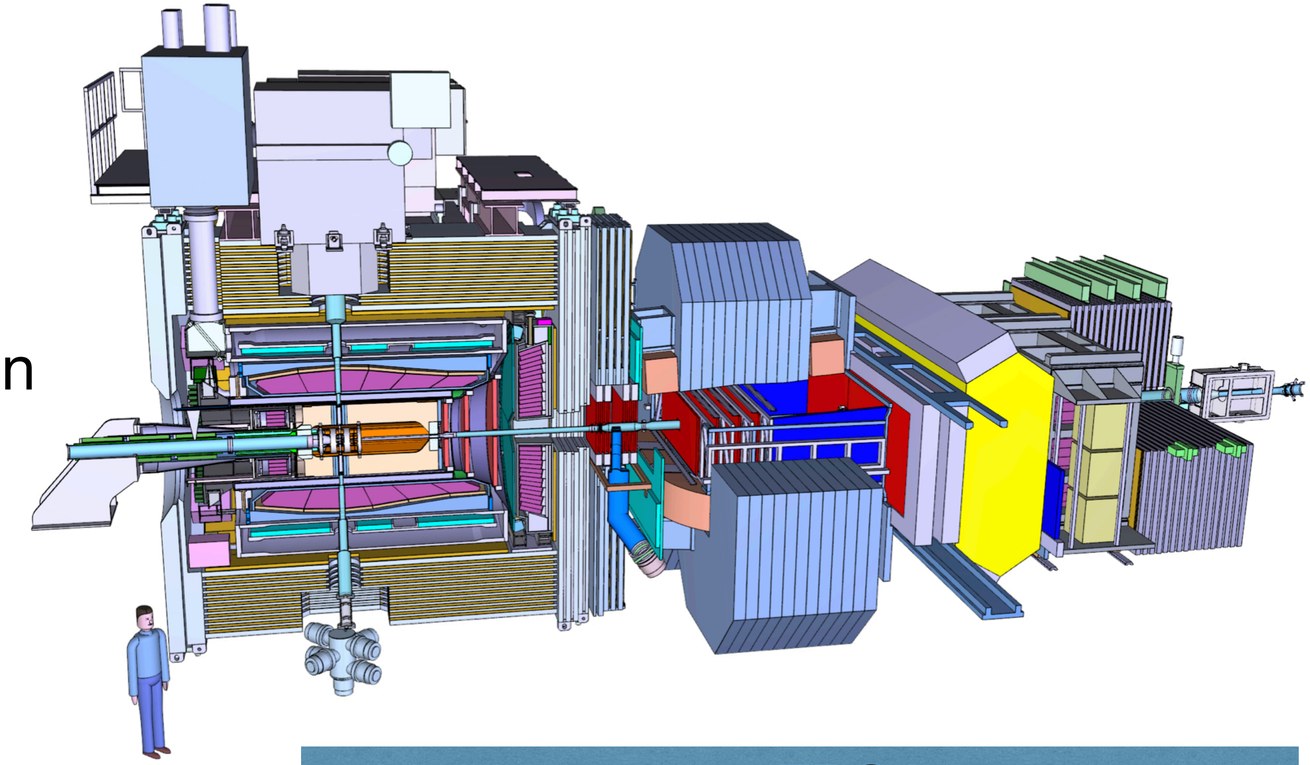

- 20,000 cores + FPGAs
  - For full online event reconstruction at 200 GB/s data stream *
  - Usage of GPUs under study
- 10,000 cores for offline reconstruction

* High luminosity run

NUSTAR

Nuclear Structure, Astrophysics and Reactions

- HISPEC/DESPEC (High-Resolution in-flight Spectroscopy/Decay Spectroscopy)
- R3B (Reactions with Relativistic Radioactive Beams)
- MATS (Precision Measurements of very short-lived nuclei using an Advanced Trapping System)
- LaSpec (Laser Spectroscopy)
- ILIMA (Isomeric Beams, Lifetimes and Masses)
- AIC
- ELISe (Electron-Ion Scattering in a Storage Ring)
- EXL (Exotic nuclei studied in light-ion induced reactions at the NESR storage ring)
- Super-FRS Experiments
- SHE (Super-Heavy Element Research)

R³B as part of NUSTAR

Reactions with Relativistic Radioactive Beams

- Operational full-size detector — nuclear physics experiments at SIS18 (FAIR Phase-0)
- Moderate data flow
- Requirement for distributed computing in calibration / tracking / alignment stages

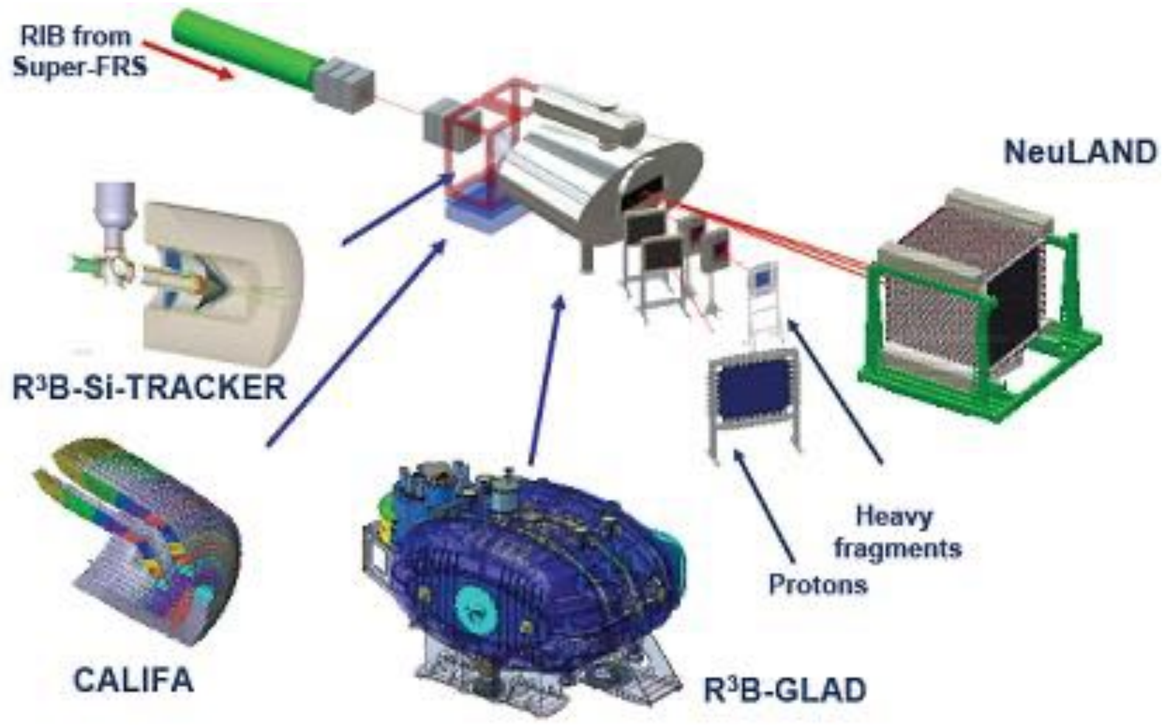

(Partial) Online reconstruction

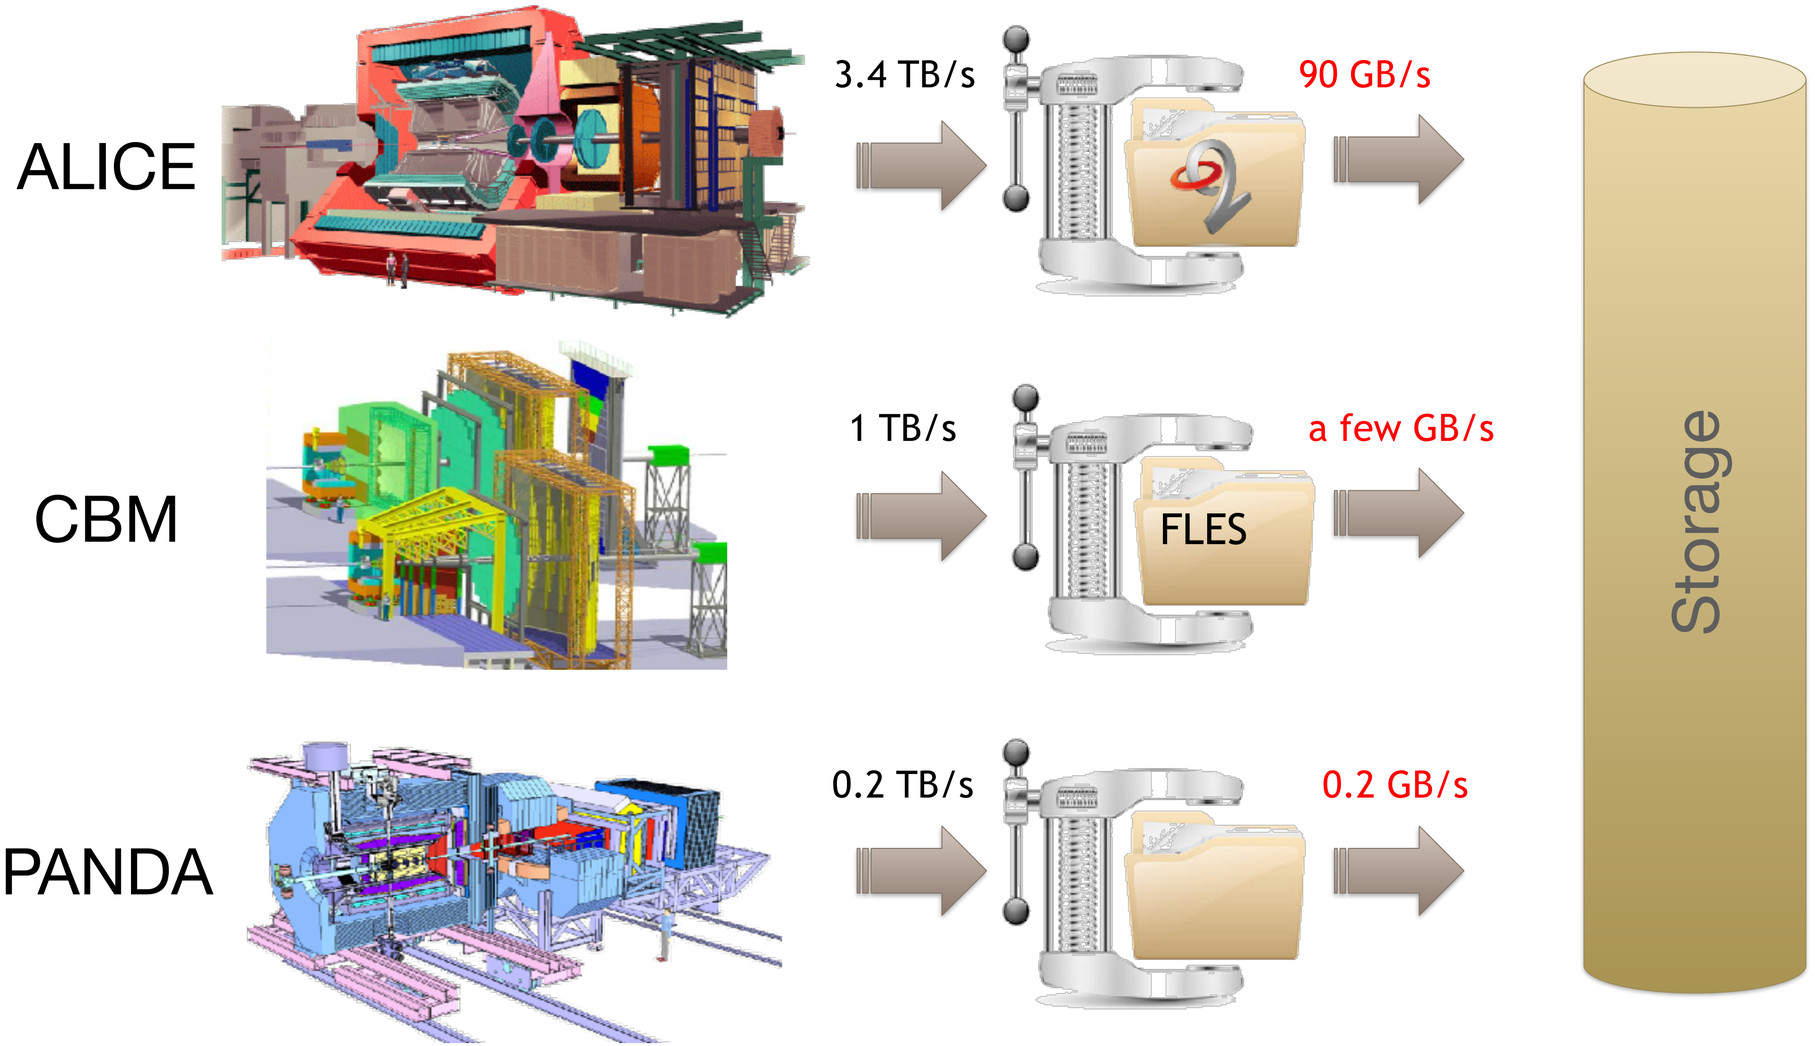

Different experiments — similar requirements

Develop **common libraries** for high throughput distributed data processing

FairRoot https://github.com/FairRootGroup/FairRoot

FairMQ https://github.com/FairRootGroup/FairMQ

DDS https://github.com/FairRootGroup/DDS

ODC https://github.com/FairRootGroup/ODC

ALFA Framework

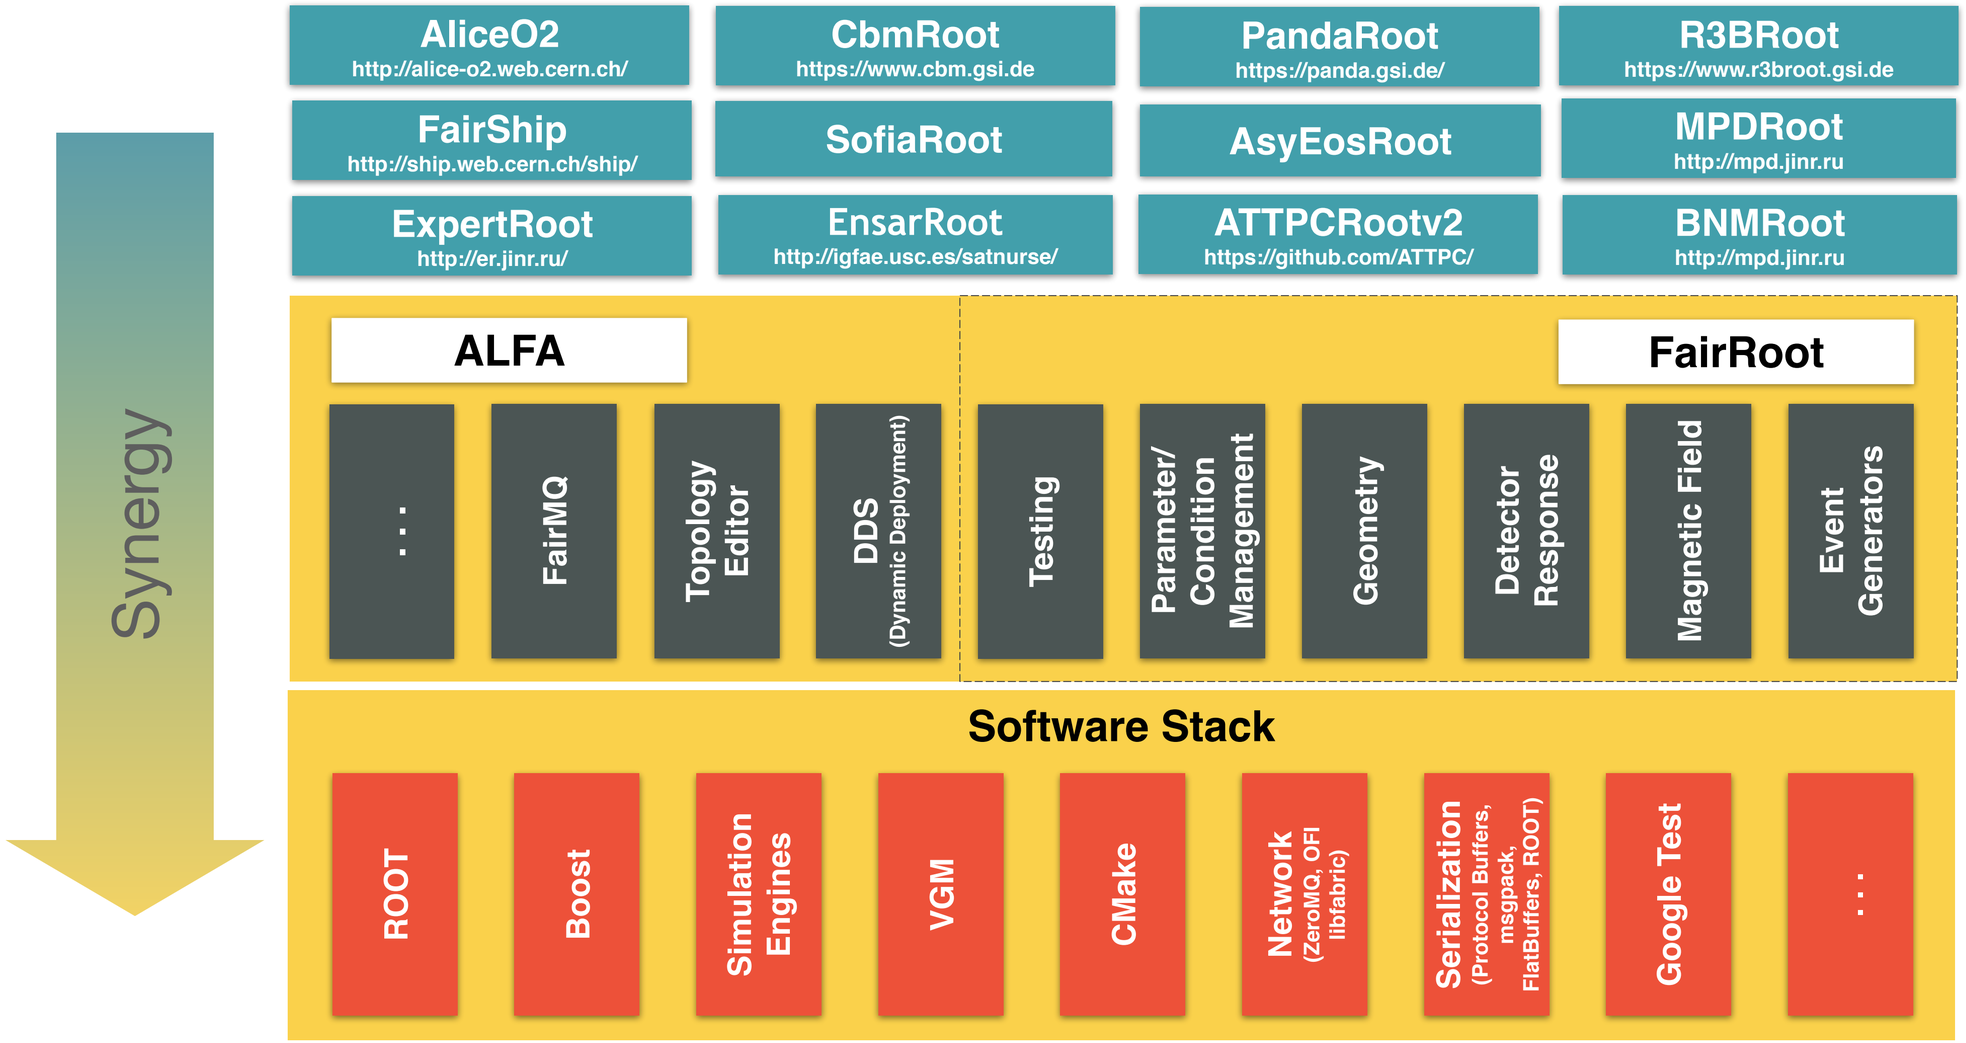

Concept of distributed computing

- High throughput online computing
- Standalone processes
- Message passing (over network or using shared memory)
- On heterogeneous hardware (CPU, GPU, FPGA)
- Independent of programming language
- Multithreading support

ALICE at CERN

EPN Workflow — running example

D. Rohr, G. Eulisse

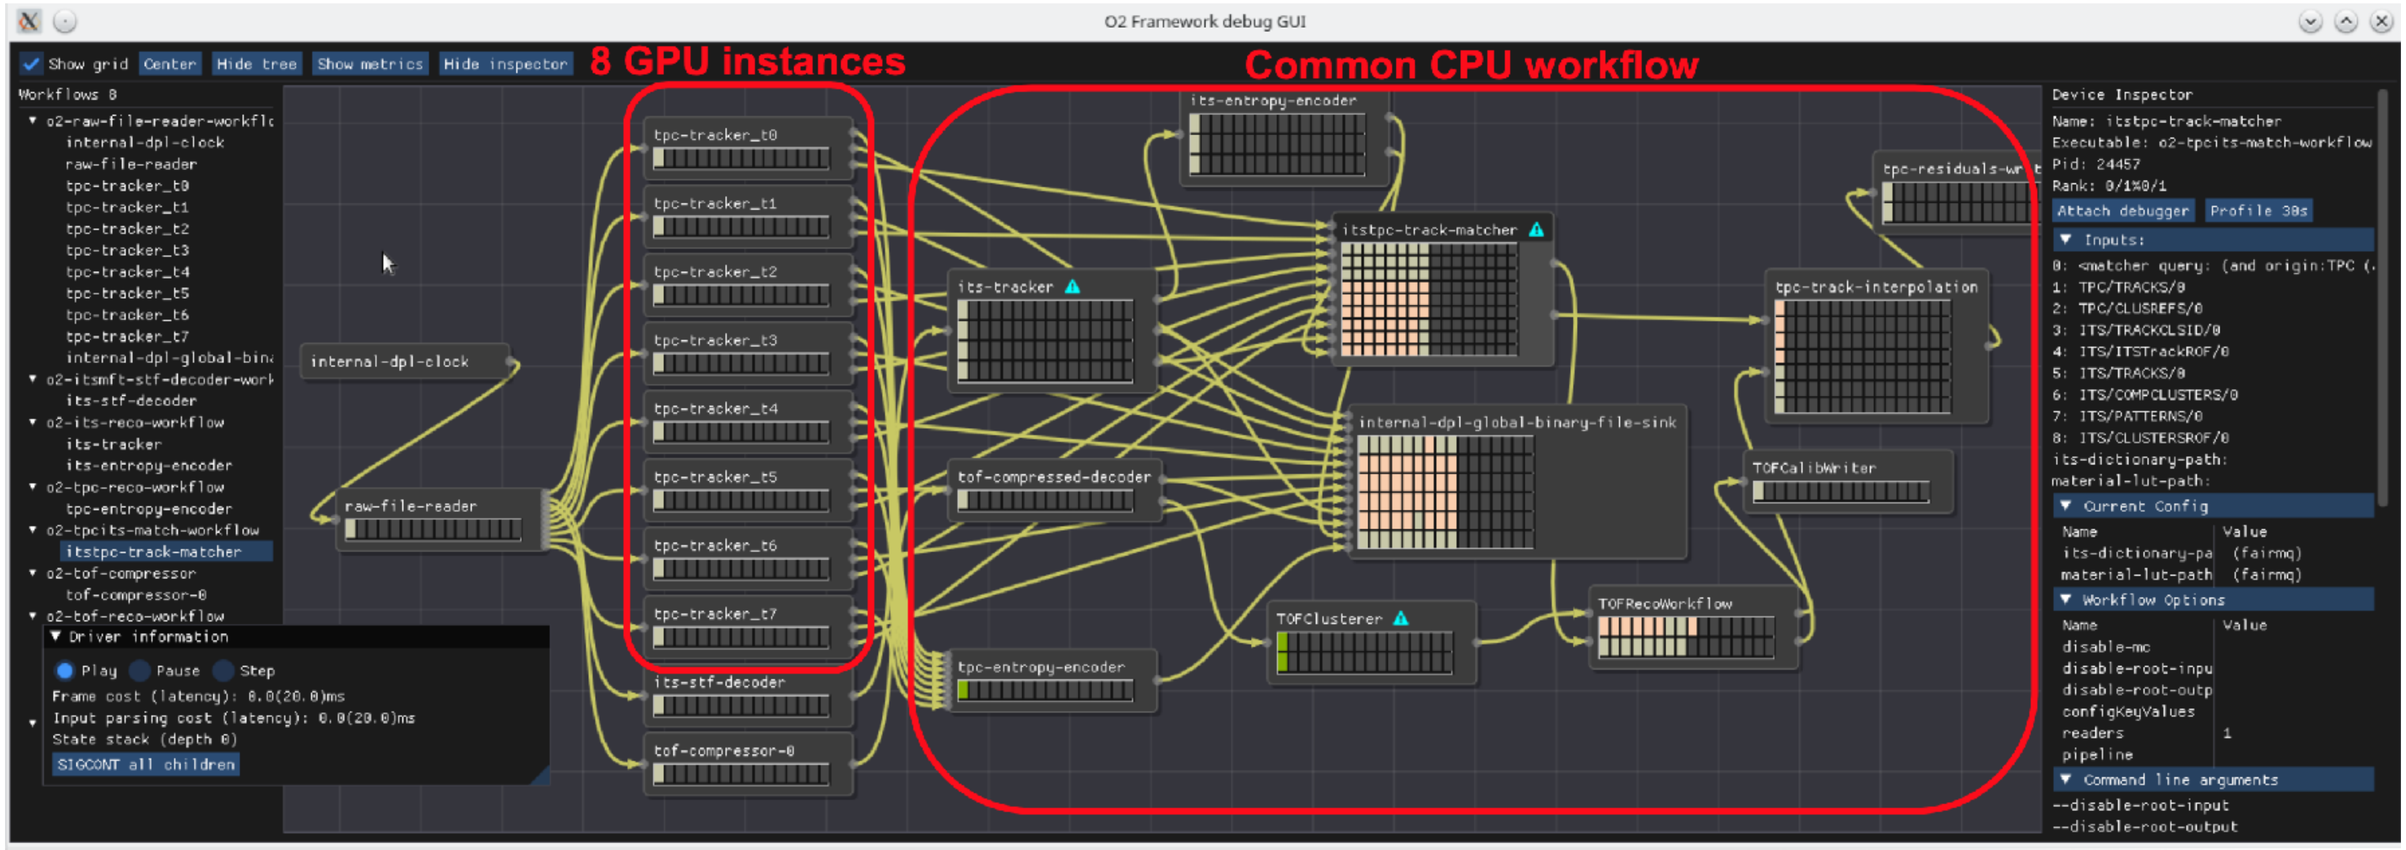

IT Infrastructure - Green Cube

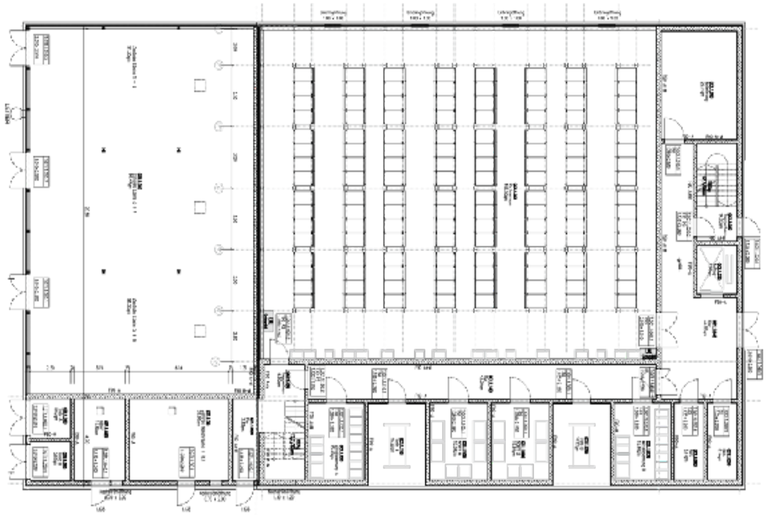

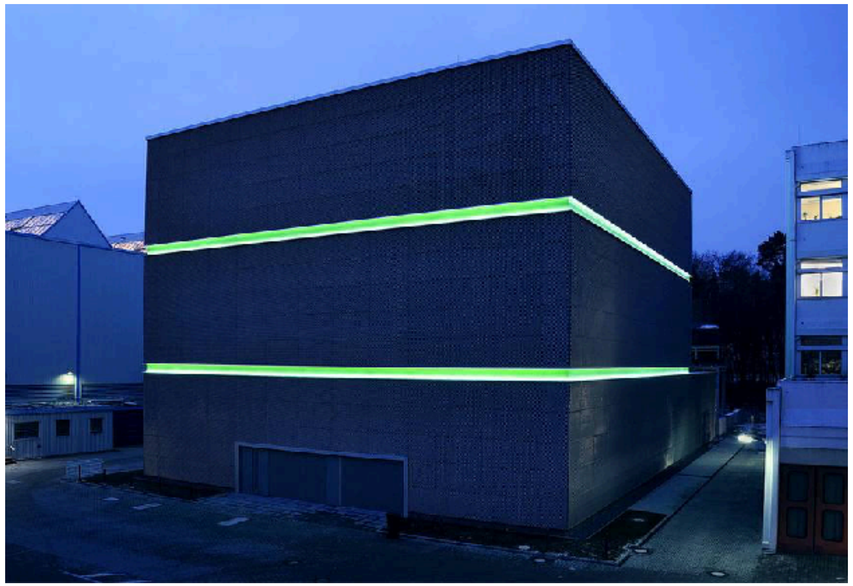

- Space for 768 19” racks (2,2m)
- 4 MW cooling (baseline)
- Max cooling power 12 MW

- Can be used for any commercial IT
- PUE <1.07
- In operation since Feb. 2016

Virgo Computing Cluster — hardware

- 530 batch execution nodes in total, 27k cores
- Infiniband network
- Lustre for high performance storage and CVMFS for read-only software distribution

- GPU cluster in operation
    - Will be upgraded with new hardware this year (order ongoing) and integrated into Virgo cluster (about 400 GPU’s)

Virtual Application Environment at Virgo

- Thin OS layer on the host - low maintenance cost
- System packages + user software in VAE (CentOS 7, Debian 10)
- Based on singularity containers
- Build for multiple architectures using Spack
    - distributed using CVMFS
- CI / CD using GitLab
- Users can start own container

Online / Offline operation

Computing at FAIR: The resources in the Green Cube will be shared between the different FAIR/GSI Partners

Dynamically allocated resources for exclusive usage and limited time

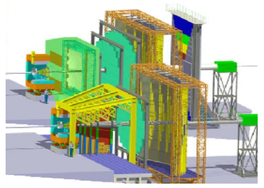

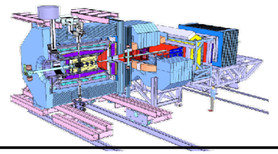

Analysis Facilities

ALICE

No separate hardware for the online compute clusters of CBM and PANDA

Generic batch farm for GSI/FAIR Users

CBM

R3B

panda

HADES

AGATA
ADVANCED GAMMA TRACKING ARRAY

BIOMAT

APPA

Summary

- Similar requirements of experiments — common libraries for message-queue based distributed data processing
- Dynamic orchestration of computing resources between many experiments
- Moving to purely containerised approach for Virgo computing cluster

Backup slides

Data Center State of the Art

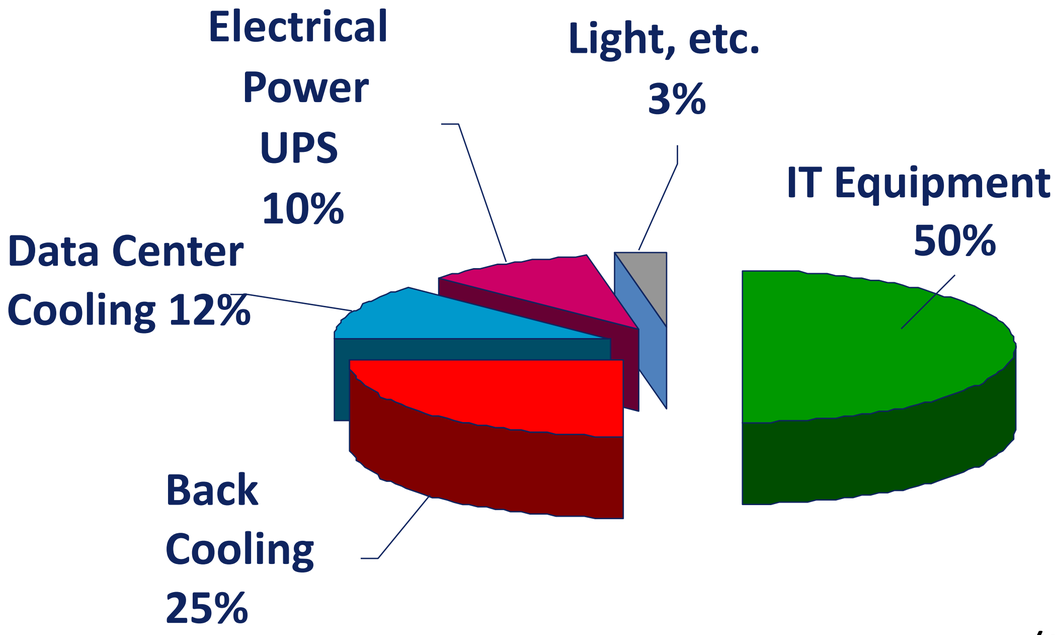

the green grid™

$$\textbf{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

Power **U**sage **E**ffectiveness

PUE = 2

(typical PUE = 1,6 to 3,5)

Note: 12C/kWh → 1W = 1,05€/a

Source: EYP Mission Critical Facilities Inc., New York