



# **GSI Analysis Facility**

D. Kresan

GSI Helmholtzzentrum für Schwerionenforschung

Darmstadt

# Facility for Antiproton and Ion Research in Europe



3000 scientists from  
50 countries


First experiments  
expected in 2028




# Highlights from FAIR construction site – installation started 2024



First power supply unit




SIS100 Tunnel – Technical Building Installation




First thermal cycle of the SIS100 string

# Green IT Cube



In operation since 2016



PUE < 1,07  
4 MW cooling  
Capacity for 768 racks on 6 floors

Foto: G. Otto, GSI Helmholtzzentrum für Schwerionenforschung

# Computing resource requirements

## CPU

	NUSTAR	CBM	PANDA	APPA
Number of cores (a)	9 k	45 k	68 k	11 k
Number of cores (b)	7 k	45 k	34 k	-


+400 GPUs – if requested

(a) Resources for simulations

(b) Resources for online data reconstruction

## Storage

	NUSTAR	CBM	PANDA	APPA
Disk total (TB)	34.250	103.000	60.680	7.037




- No dedicated / fixed hardware for an experiment
- Will not take beam all at the same time
- Computing resources will be shared dynamically

- Resources reserved on a shared cluster
- ~16k logical cores (hyperthreading) ↔ 8k physical cores
- 7 PiB disk storage under a Lustre distributed file system
- 400 GPUs

Nodes	CPU / node	Total CPU	Memory / node
169	96	16.224	192 GB


- Network connection
  - Internally 100 Gb/s HDR InfiniBand
  - 10 Gb/s LHCONE, 2 Gb/s DFN
- Memory limits imposed by Slurm via cgroups
  - 4.4 GB per physical core



- **GSI\_4core** and **GSI\_8core**
- Almost all jobs are queued on 8-core VObox
- JobAgent fills the 8-core queue with single- and multicore jobs



Jobs per n\_cores at GSI::8core




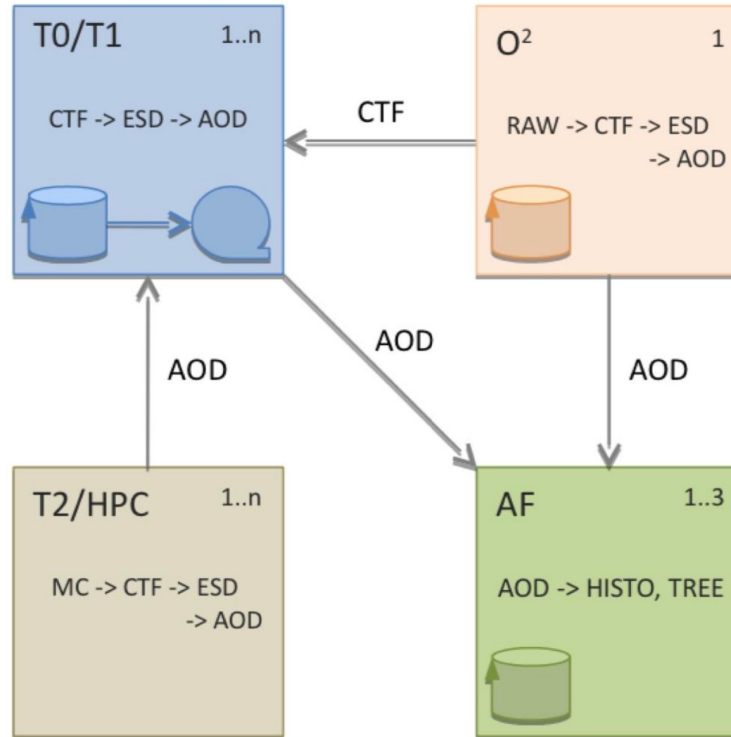
Average in the last year:

- ~3.5k utilized cores
- ~1.1k running jobs

# Analysis Facilities in ALICE computing model

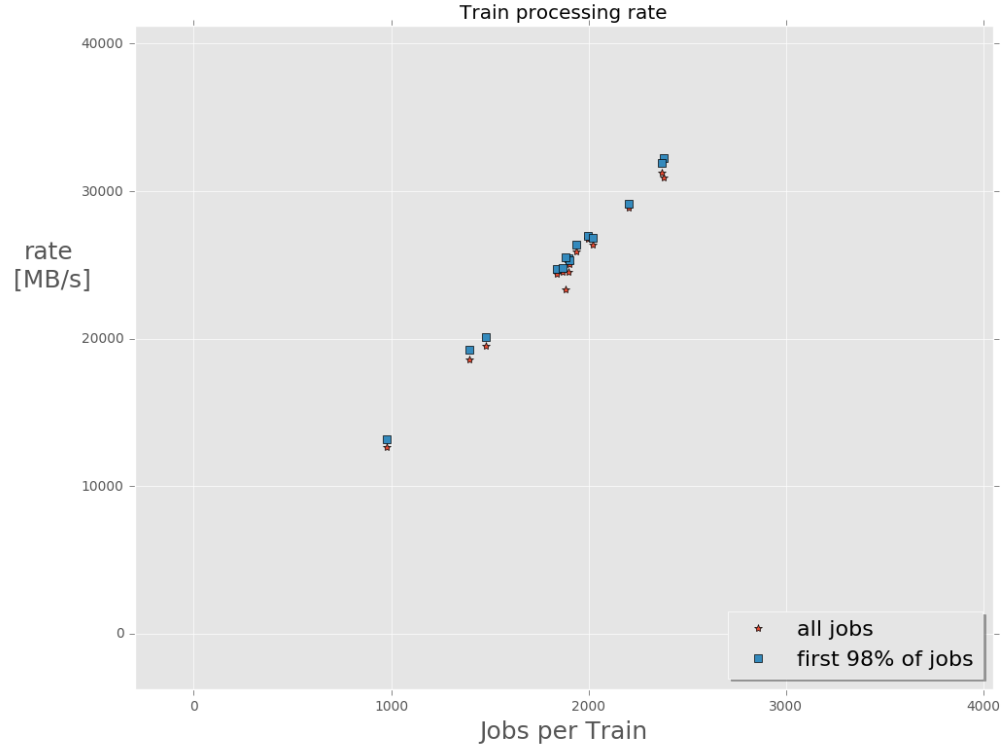

## Requirements:

- Be able to digest > 5 PB of AOD in a 12 hours period
- Analysis trains need 5 MB/s per job slot
- Serve 20.000 job slots at an aggregate throughput of 100 GB/s




# Increasing I/O performance

- XRootD with plugin (developed at GSI)
- Plugin for **local redirect** – optimizes I/O throughput of the analysis jobs



# Containerized environment

- Host: minimal RHEL 8 (Rocky Linux 8.9)
- Jobs run within Apptainer
- Taken from [/cvmfs/alice.cern.ch/containers/](https://cvmfs/alice.cern.ch/containers/)
- GSI was the first WLCG site running with containers



- Supports individual experiments and groups
  - Heavy Ions, ALICE, FAIR Experiments, Nuclear Structure
- Still very much batch oriented
- Increasingly towards AI-based analysis



- Towards web-based, interactive deployment of Python and ROOT analysis using JupyterHub
- To consider for production deployment:
  - Considering authentication, security, port-forwarding



Thank you for attention!

