

# Dress Rehearsal Review: FAIR Perspective

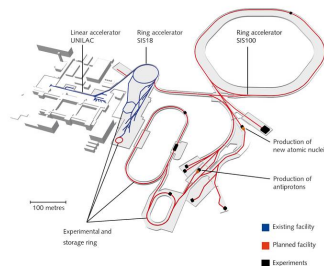
Marek Szuba / GSI

2020-12-09, 2nd ESCAPE WP2/DIOS Workshop

# Outline


- 1 FAIR and Its Pillars
- 2 FAIR Data-lake Work Flows
- 3 November Dress Rehearsal
- 4 Summary

# Introduction to FAIR



- Facility for **Antiproton and Ion Research**
- Upcoming international accelerator centre in Darmstadt, Germany
- Extension of acceleration facilities of GSI Helmholtz Centre for Heavy-ion Research
- Four experimental pillars:
  - PANDA — antimatter studies
  - CBM — heavy-ion physics
  - NUSTAR — rare/exotic beams for astrophysics
  - APPA — Atomic, Plasma Physics and Applications


# PANDA



# $\bar{P}$ ANDA

- **anti-Proton AN**ihilations at **DA**rmstadt
- General-purpose detector with large acceptance
- First fixed-target, later collider mode
- Hadron spectroscopy, hadrons in matter, nucleon structure, hypernuclei
- Collisions:  $\bar{p} - p(A)$ 
  - *high event rates but relatively low event size*

# CBM



# CBM

- **Compressed Baryonic Matter**
- Modular detector — different configurations for different runs
- Fixed-target
- Phase diagram of strongly interacting matter, phase transitions, quark-gluon plasma, neutron stars
- Collisions: up to and including heavy-ion  $A - A$ 
  - *high event rates AND large events*

# FAIR Data-lake Work Flows

## Current State

- None yet!
  - still a few years till FAIR Phase 1 launch
  - work ongoing on experiments' computing models
  - hardware not ready for target requirements
- Likely first candidate: Monte-Carlo reconstruction for CBM
  - January 2021???
- Meanwhile, we can approximate



# FAIR Data-lake Work Flows

## Operational Assumptions

With CBM raw-data ingestion as upper bounds:

- 3-month runs every year
- $\sim 18$  PB/run  $\rightarrow \sim 200$  TB/day
- ROOT files — binary, possibly compressed data
- file size:  $\sim 1$  GB
- near real-time reconstruction — no ESD, no persistent AOD  
— raw data as analysis input
- QoS classes:
  - RAW\_HOT — reasonably reliable, reasonably fast —  
 $\sim$  CHEAP\_ANALYSIS
  - RAW\_COLD — highly reliable —  $\sim$  SAFE

## FDR #1 “FAIR” Work Flow

- mock CBM raw-data ingestion at  $\sim 1/500$  expected data rate
- randomly generated 1-GB files uploaded every  $\sim 10$  minutes
- source system: workstation on GSI office network (Gigabit Ethernet)
- uploads to “home” RSE (GSI-ROOT; QOS=CHEAP\_ANALYSIS)
- QoS tagging: 1x CHEAP\_ANALYSIS (FTS no-op) + 2x SAFE
- steering: Rucio CLI client + bash scripts
- authentication: long-lived X.509 proxy cert with periodic keyless renewal of VOMS extension

# Run Progress

- before 12:30 UTC: all uploads failed due to GSI-ROOT problems
- afterwards: nearly 100 % success for both uploads and replication
  - 7 uploads failed when the Rucio server had to be restarted
  - replication of last few uploads delayed due to FTS issues, eventually succeeded
- upload frequency doubled around 13:30 UTC, no problems

# Summary and Outlook

- so far so good ...
  - ...for a mock work flow and  $\sim 1/250$  nominal data rate
- next step: continuous uploads over 10GE
  - highest we can go on current hardware
  - theoretical maximum:  $\sim 105$  TB/day
- Rucio server as potential bottleneck
  - local buffer needed to prevent data loss on outages
  - minimise client-server communication: subscriptions for QoS rules, ...
  - asynchronous renewal of Rucio auth tokens?

**THANK YOU**