CHEP24 highlights (PoV)

Napoli, 19-20 novembre 2024

L. Rinaldi

CHEP 24 in numbers

470 participants

351 talks (XX morning plenary) & 143 posters

9 tracks (afternoon parallels)

Many interesting topics, but impossible to follow everything…

The First Release of ATLAS Open Data for **Research**

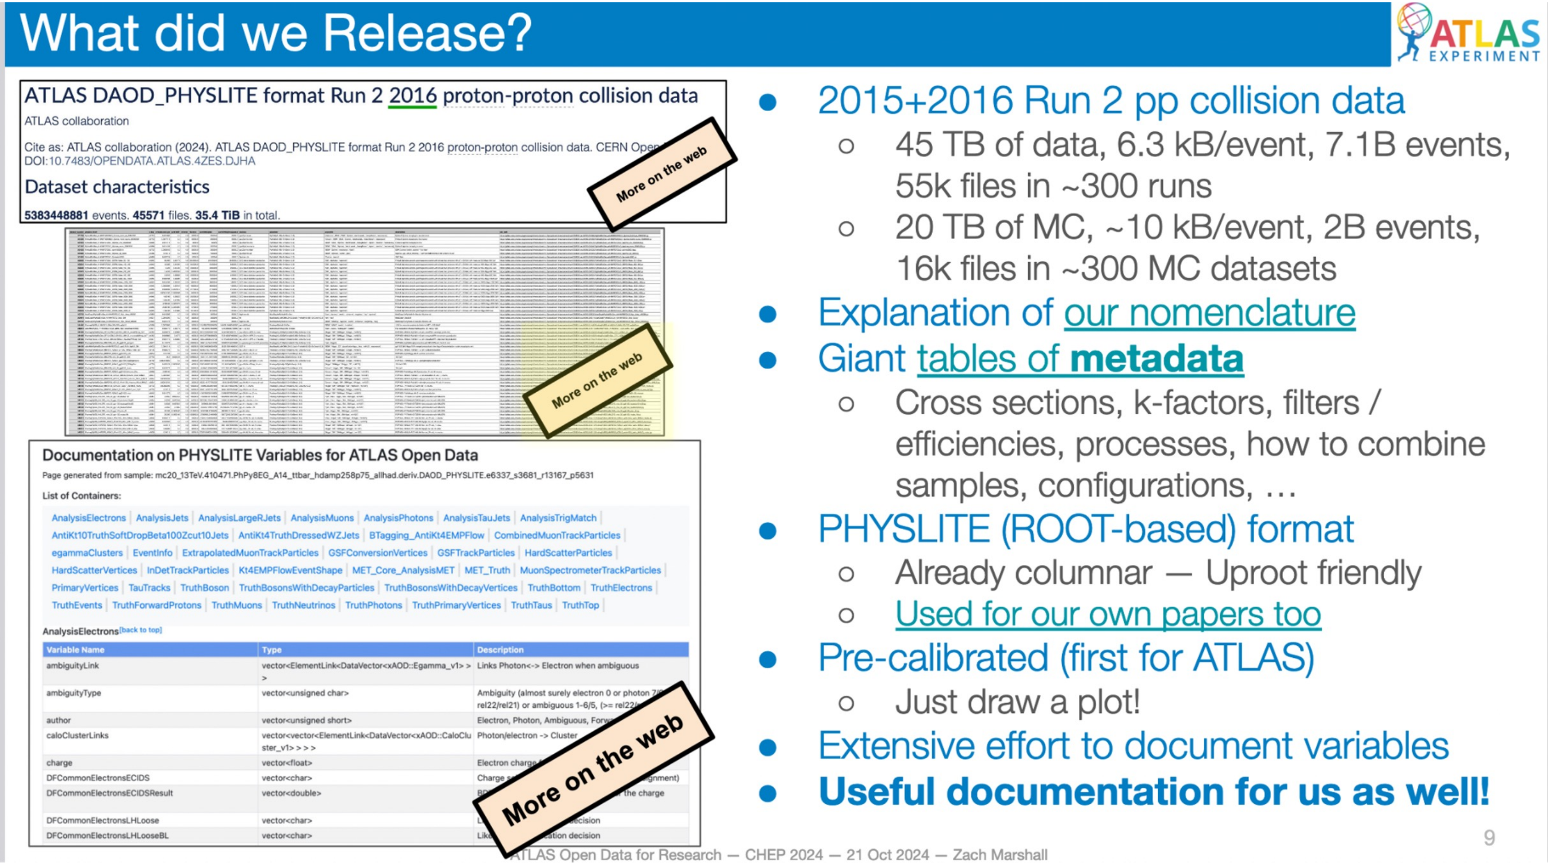

Simulating the Carbon Cost of Grid sites

- **Sustainability:** Can the CO2e per unit work be reduced?
  - Yes, tune hardware to optimal frequency, particularly when forecast Carbon Intensity is high.

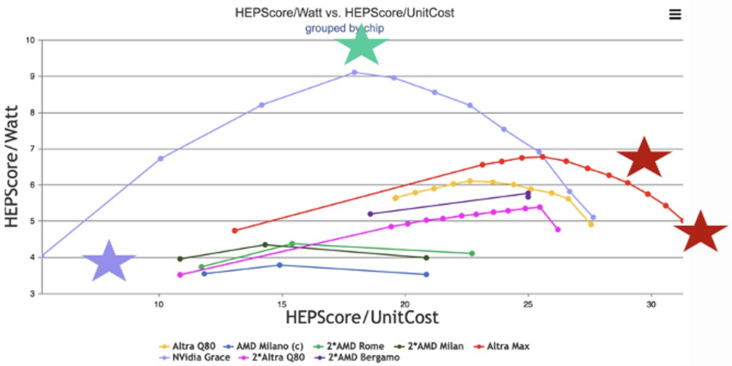

- **Costs:** What hardware requires the minimum CapEx (most HS23/£) and the least OpEX (least power per unit work)?
  - CapEx decisions informed by the Smart Procurement utility.
  - OpEx can be reduced by running at lower frequencies (compromise!)

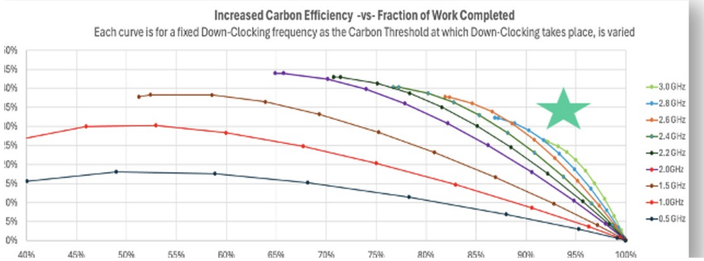

- **Power Supply:** In a crisis, how much power could be saved at peak times?
  - With ARM: 40-75%
  - With x86, 20-50% (Slide-11)

Simulating the Carbon Cost of Grid sites

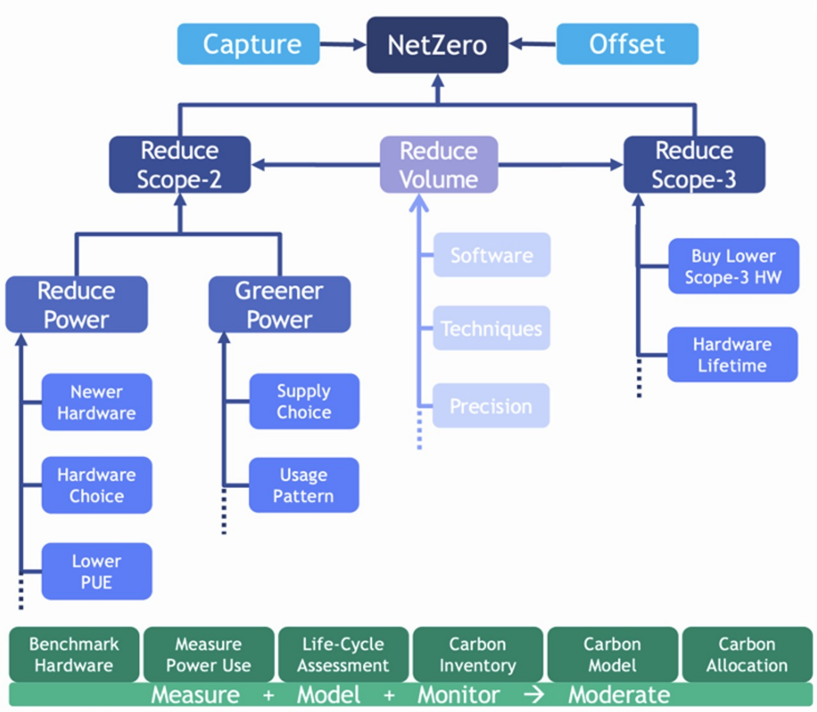

- Many things will need to contribute to the journey to NetZero.
- Those thing depend on time and place.
- Money should be spent where it does the most good.
- There can be clear wins: In 2023/4 the adoption of ARM gave us more compute, for less money, *and* better carbon efficiency.
- There are also compromises: Our investigations of down-clocking suggest that for a modest 5-10% reduction in work-rate, we can reduce the carbon/job by up to 30%.
- There is a lot more to be done.

Carbon, Power, and Sustainability in ATLAS Computing

(Computing) Carbon footprints

- Much more attention to carbon footprints recently (CERN reports in 2021 and 2023; see bkup)
  - Outside groups also looking: EE HPC WG (recommendations), Open Compute Project, GreenDiSC
- Computing is ~5% of CERN's footprint when the LHC runs (accelerator cooling ~80%)
  - CERN mostly draws power from the French (nuclear, quite low-carbon) grid
- About 10–20% of ATLAS computing is at CERN, much of the rest is on less green power
- Most studies focus on *power* (operational carbon); embodied carbon >15% of the total (1, 2)
  - Relevance of embodied carbon (Scope 3) will increase as western power grids decarbonize
- Extrapolating to HL-LHC, computing could be a **large** fraction of the ATLAS carbon footprint!

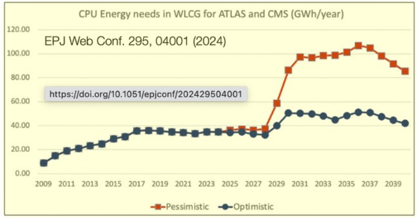

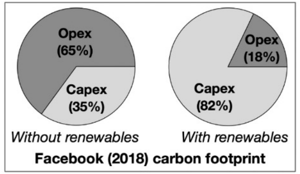

arXiv:2011.02839 Opex: Operational carbon
Capex: Embodied carbon

Waste, loss, and unused data

- *Using* carbon for important science is "allowed" — *wasting* carbon is never ok!
- Constantly monitoring *unused data* in the production system
  - Requests made, bug found, reproduced before the data were looked at
  - Processing done based on a too-inclusive pattern (mc*)
- Steady progress to improve CPU/wall efficiency to over 90% (below for 8-core production jobs)
  - Impact of failed jobs is visible; errors on copying output (after all the CPU has been consumed) are killer!
  - Constant effort to reduce serial portions of many-core jobs as well (wasted CPU and power)

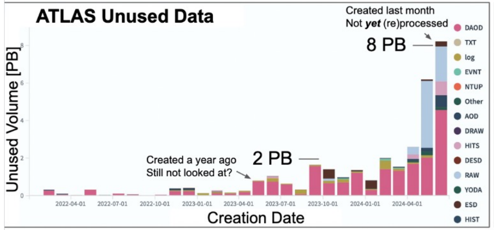

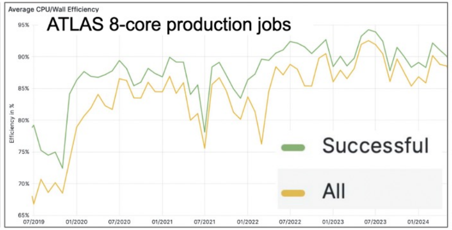

Heterogeneous Computing and Power Efficiency in HEP

ARM Physics Validation

Most LHC experiments (**ATLAS**, **CMS**, **ALICE**) have done a first round of extensive Physics Validation campaigns against our ARM cluster @ Glasgow:

- **ATLAS**: Full simulation and Reconstruction are physics validated.
  ATLAS is ready for pledged ARM resources!
- **CMS**: Physics validation on ARM mostly successful, but not conclusive.
  CMS is not in a position to use ARM processors in production!
- **ALICE**: Extensive test of MC simulation jobs, no analysis workflows.
  Recommends ARM segregation or mixed queue with enable/disable!
- **LHCb**: Groundwork & test samples done, full physics validation not done.
  Production use of ARM unlikely before end of 2024!

Latest reports from **GDB** (June 2024 @ CERN): https://indico.cern.ch/event/1356135/

It's time for **VOs** to start sending ARM jobs our way … we have over 4k ARM cores !

Taking on RISC for Energy-Efficient Computing in HEP

Benchmarks Comparison

The histogram below compares the results of the benchmarks so far, normalized to the AMD server.
On the two servers (**AMD & ARM**), the HEP-Score results are also shown.

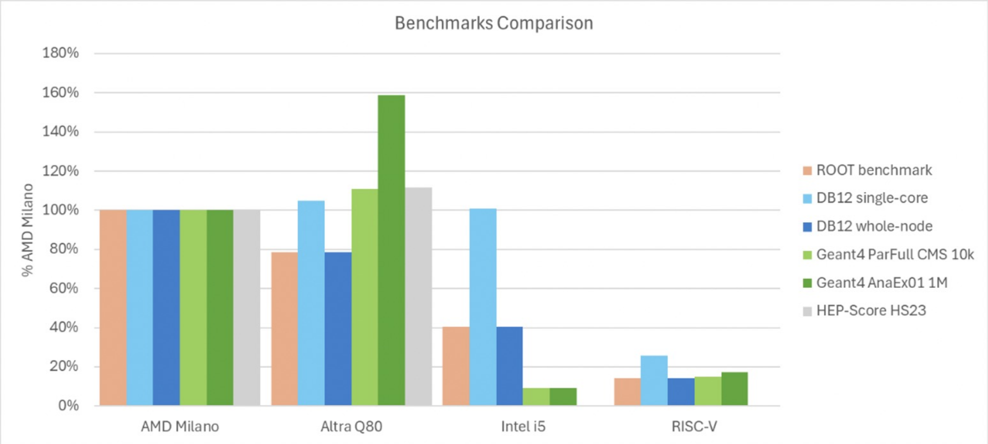

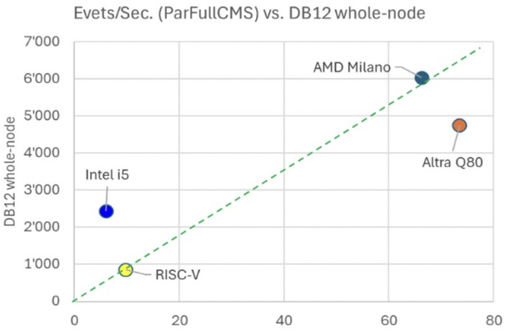

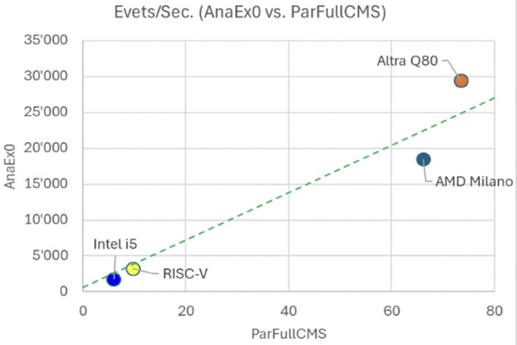

As mentioned, the first two benchmark (**ROOT** & **DB12 single**) are single-threaded, the others are multi-threaded (**DB12 whole**, both **Geant4** simulations, and the **HEP-Score**).

The scatter plots on the right show how compatible are their results.

Note: the green line is not an actual fit.

Heterogeneous computing at INFN-T1

Experiment point of view: ATLAS

INFN CNAF

ATLAS has been running jobs on ARM@CNAF for one year

- initially, only test jobs
  - ATLAS Software already Physics-validated on ARM
  - Technical validation performed at CNAF (HTCondorCE, pilot/PanDA, containers)
- workflows:
  - Full and Fast MC Simulation
  - MC reconstruction
  - Group production
  - User Analysis

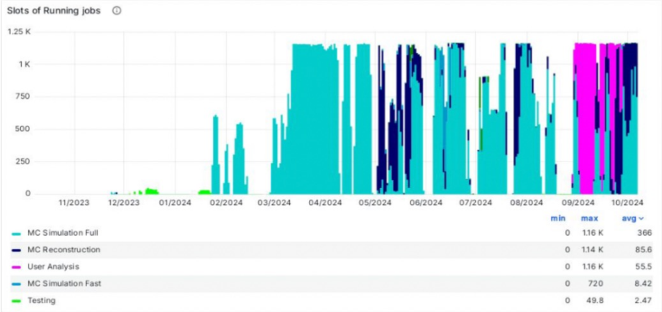

Very good performance observed

- Steady use of available resources. On peak: ~12% of ATLAS-dedicated resources at INFN-T1

12

AtlFast3

- Machine learning methods are now fully integrated into the standard simulation of many experiments
  - Hybrid usage next to full Geant4 simulation and parametrised modeling

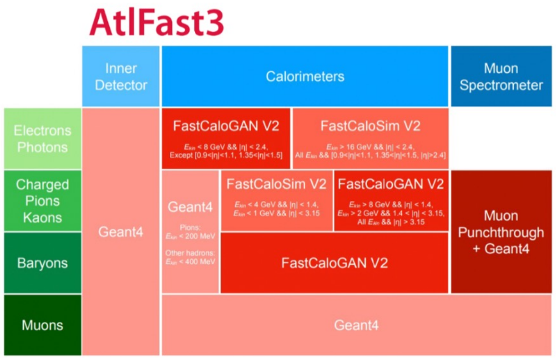

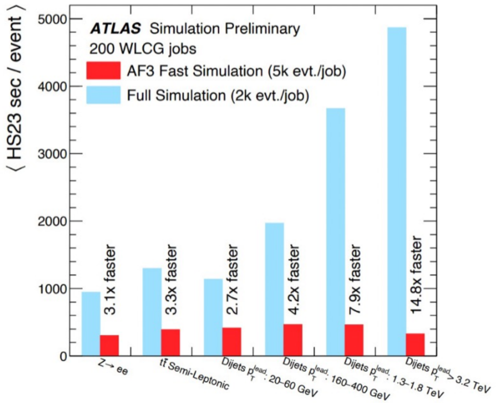

- Training runs **3-4x faster** on Leonardo (**A100** GPUs) and 2-3x faster on CNAF-HPC with respect to LXBATCH at CERN. **Supercomputers give a great performance boost!**
- **To do**: run on other resources (also cloud), architectures (ARM) and for more particle types, code optimisation (both general and to take even more advantage of multi-CPU/GPU nodes).

Quantum Computing?

QC topics

- IBM technology and roadmap
- CERN QTI provided examples of current implementation and use cases: integral acceleration with hybrid QA, quantum anomaly detection to detect BSM, quantum generative models to simulate calorimeter images, for parton shower, Hamiltonian moments calculation for elastic scattering
- Jet discrimination with QGNN, Jet reconstruction (annealing)
- Track reconstruction (HHL)
- Particle identification
- Quantum event generation

Methods and applications

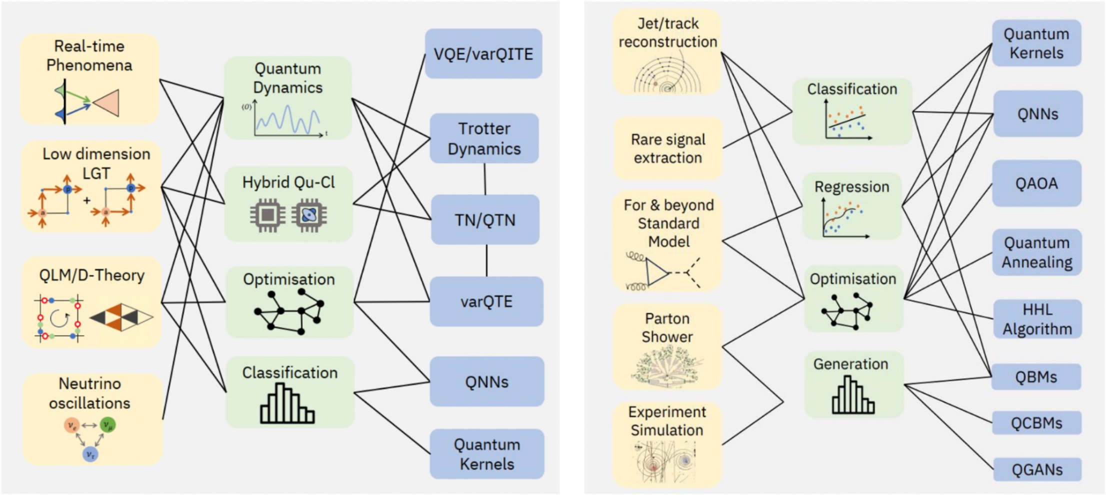

Data Challenge 24

- **Preparing for HL-LHC data rates:**
  - DC24: 25% HL-LHC
  - DC26: 50% HL-LHC (planned)
- **Experiments:**
  - ATLAS, CMS, ALICE, LHCb, Belle II, DUNE
- **FTS is a pilar technology for DC**
- **Observed bottlenecks:**
  - Scalability of FTS and Rucio, token handling, storage systems overload, network tuning
- **Positive outcome: improved production rates!**
- **Future DC26 priorities (2x DC24 rates):**
  - Network orchestration and monitoring (scitags); token support; focus on tape
  - Intermediate mini-challenges

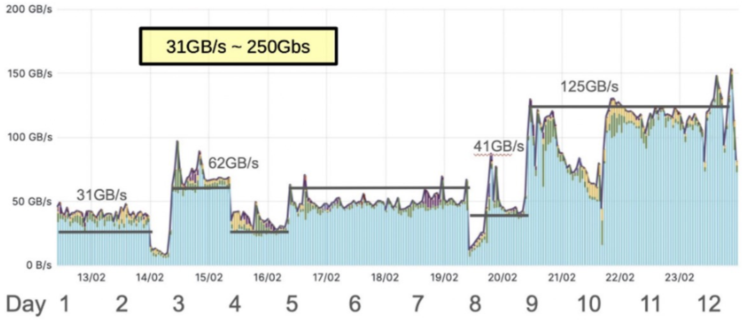

CMS successful transfer throughput

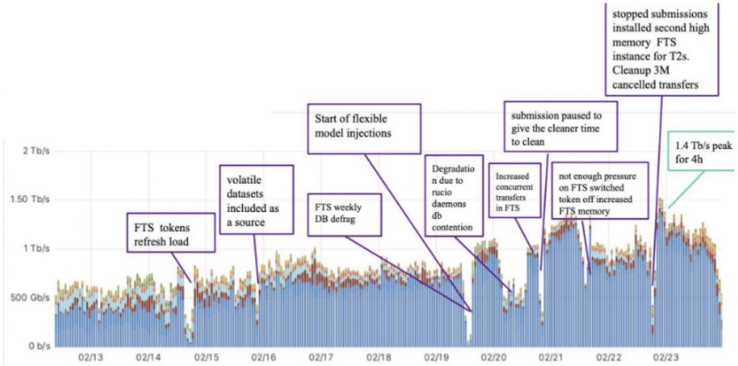

ATLAS successful transfer throughput

Rucio highlights

Rucio was **omnipresent** in CHEP (mentions in plenary and side tracks).

Became de-facto platform used in HEP and Astrophysics organizations for scientific data management (**20+ sites and > 2 exabytes of data**)

1st IT department talk on behalf of the Rucio project: Advancing Large Scale Scientific Collaborations with Rucio

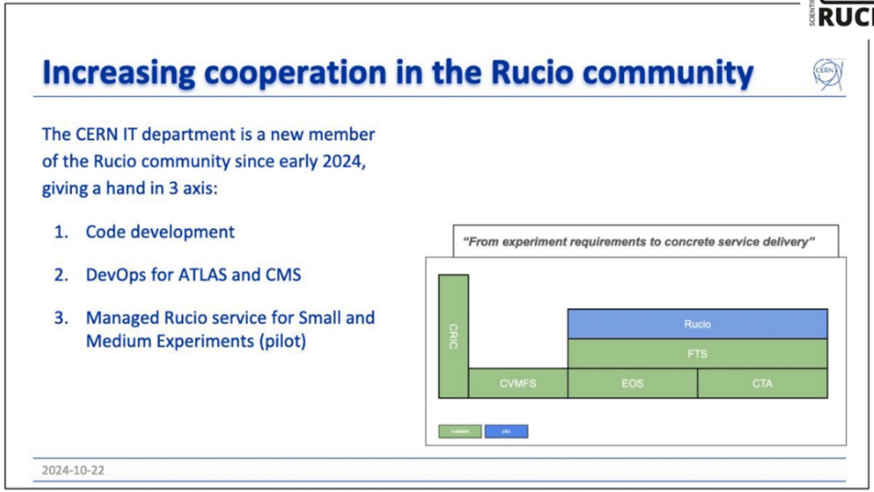

Public recognition of IT work on Rucio for experiments

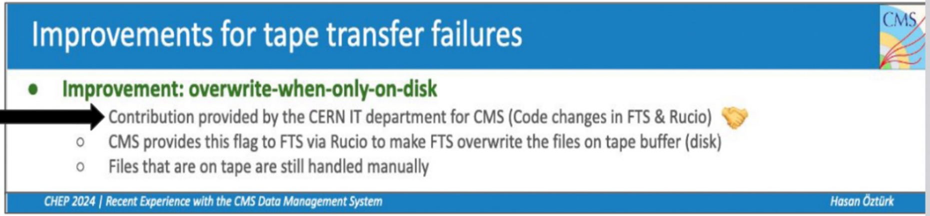

Distributed Computing

Track 4 - Main Themes

30 Talks
19 Posters

- Tokens
  - Tokens being widely used and lessons are being learned
  - Current production setups
  - Development of best practice
  - Satellite services development (e.g. IAM)
  - Adoption of tokens outside WLCG - early stages
- Operations
  - Grid computing adapts to changing circumstances
    - Incorporation of modern resources such as ARM, HPCs, GPUs and Cloud.
    - Optimizing use of available resources
    - Monitoring
  - Security: It's not just technology, people matter, too. The pDNSSOC package was suggested as a lightweight way for smaller sites to get the benefits of a SOC
- Distributed computing as part of non-WLCG computing models:
  - Gaining popularity especially in Astronomy: SKA, LSST, Einstein Telescope, CTA, HERD, but also DUNE (not astronomy)
    - Predicted SKA data volumes easily comparable to WLCG, building on WLCG experience for large scale operations!

Tokens

- Tokens are now a reality! The infrastructure is almost token ready- time to focus on the operational models!
- DC24, a major milestone: millions of transfers with tokens!
- Token implementations of middleware need to improve: FTS/Rucio/Dirac workflows, IAM: all doing a lot of work:
  - Since August ATLAS has been running tokens to 15 sites: 1-2Hz with 5Hz spikes!
  - CMS also has most CEs and quite a few transfers running with tokens.
    - Will rely to the CERN-IT Vault instance
- Performance must be stress-tested
- The balance between operability, security and performance needs to be found:
  - Audience, lifetime, scopes are the three orthogonal Parameters that one needs to tweak to meet the operational needs without compromising too much security.

October 19 - 25, 2024
CHEP 2024

Overview of national Analysis Facilities

AF evolution, operational experience at:

- Chinese HEPS
- German DESY NAF
- Spanish CIEMAT

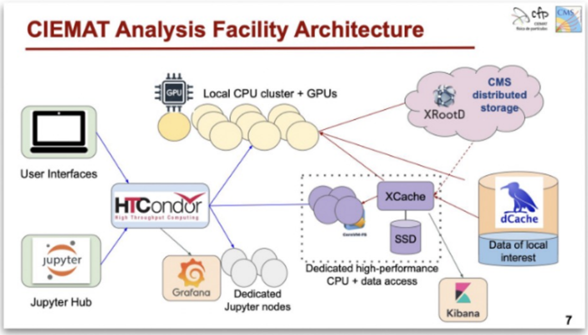

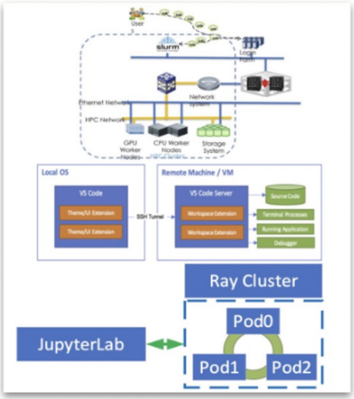

Recurrent themes:

- **Notebook**-based interfaces are widespread
- **Scale out** with Dask + HTCondor

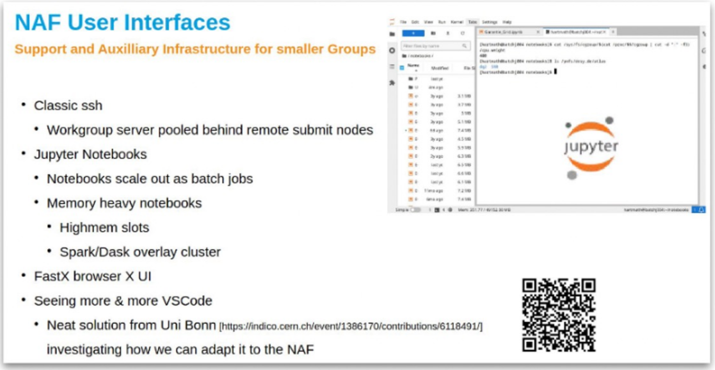

INFN Data Stewards

Skills 4 eosc

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

INFN

CNAF, LNGS, LNF, Ferrara, Torino

SO FAIR SO GOOD

the INFN strategy for Data Stewardship

Stefano Bianco, Daniele Bonacorsi, Concezio Bozzi, Luca Dell'Angello, Luciano Gaido, Francesca Marchegiani, Irene Piergentili, **Lorenzo Rinaldi**, Stefano Dal Pra

The INFN Data Steward team

Stefano Dal Pra
Senior technologist @ INFN-CNAF
Master degree in Electronic Engineering

Francesca Marchegiani
Technologist @ INFN-LNGS
Master degree in Chemistry

Irene Piergentili
Fellowship for technological research @INFN-LNF
Master degree in Archive and Library theory and management

Lorenzo Rinaldi
Associate Professor @ UNIBO and affiliation @ INFN
PhD in Physics

XXX YYY

User support plan

The main objective is to support small research groups
- Data Management Plan drafting, according to a precise check-list
- publication (which Open Access level?)

Target:

Researchers with no or few knowledge of FAIR principles

Different levels of support

Researchers (highly) familiar with FAIR principles

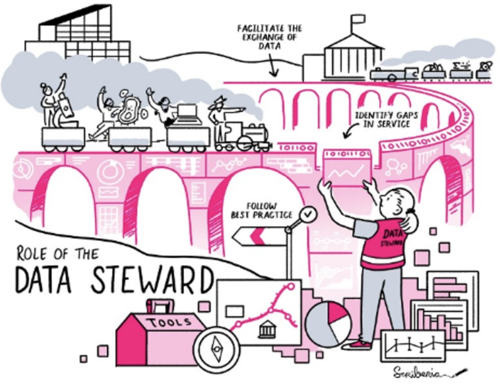

https://openworking.files.wordpress.com/2022/04/data-stewards.jpg?w=1024

28th Conference on Computing in High-Energy and Nuclear Physics (CHEP 2026)

25-29 May 2026
Bangkok, Thailand

Phat Srimanobhas
Department of Physics, Faculty of Science,
Chulalongkorn University

https://indico.cern.ch/event/1471803/