27th Conference on Computing in High Energy and Nuclear Physics (CHEP2024) summary

Tania Korchuganova (Pitt)

SDCC-NPPS meeting, 14 Nov 2024

~~ATLAS weekly, 5 Nov 2024~~

University of Pittsburgh

Introduction

The **CHEP conference** series addresses the **computing, networking and software issues** for the world's leading data-intensive science experiments that currently analyse **hundreds of PB of data** using **worldwide computing resources**.

Hosted by the AGH University of Kraków, Institute of Nuclear Physics Polish Academy of Sciences and Jagiellonian University in **Krakow, Poland**

470 participants

351 talks & 143 posters

Scientific program: 9 tracks

Track 1 - Data and Metadata Organization, Management and Access

Track 2 - Online and real-time computing

Track 3 - Offline Computing

Track 4 - Distributed Computing

Track 5 - Simulation and analysis tools

Track 6 - Collaborative software and maintainability

Track 7 - Computing Infrastructure

Track 8 - Collaboration, Reinterpretation, Outreach and Education

Track 9 - Analysis facilities and interactive computing

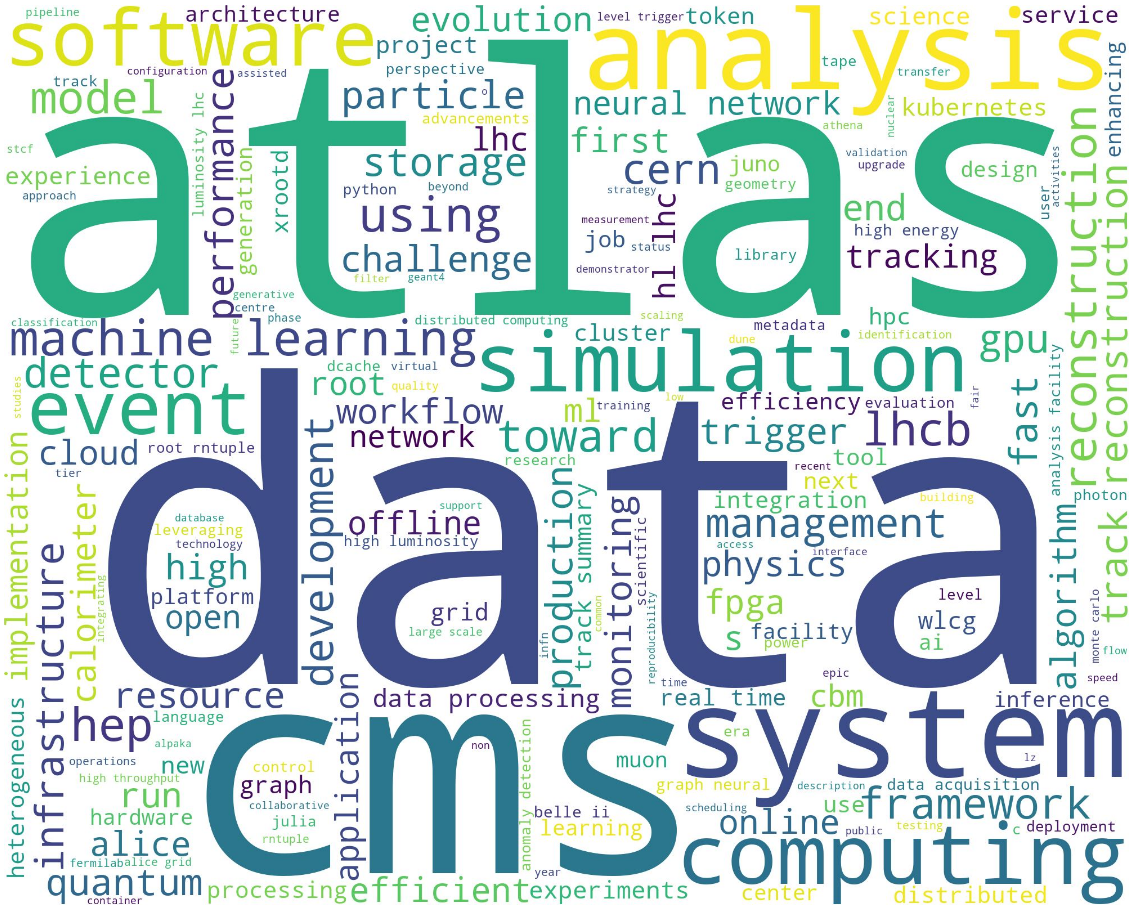

Scientific program: 9 tracks

Track 1 - **Data and Metadata Organization, Management and Access**

Track 2 - Online and real-time computing

Track 3 - Offline Computing

Track 4 - **Distributed Computing**

Track 5 - Simulation and analysis tools

Track 6 - Collaborative software and maintainability

Track 7 - Computing Infrastructure

Track 8 - Collaboration, Reinterpretation, Outreach and Education

Track 9 - Analysis facilities and interactive computing

7 parallel sessions every afternoon

⇓

it is impossible to attend everything and one need to run between rooms for interesting talks

⇓

<u>very subjective</u> summary of the conference

University of Pittsburgh

ATLAS contributions

65 contributions (~13%):

- 2 plenary talks:
  - The First Release of ATLAS Open Data for Research
  - Total Cost of Ownership and Evaluation of Google Cloud Resources for the ATLAS Experiment at the LHC
- 44 session talks
- 19 posters
  - 2 winners of "Poster of the day":
    - Performance of the ATLAS GNN4ITk Particle Track Reconstruction GPU pipeline
    - Evolution of Regional, Age and Gender Demographics in the ATLAS Collaboration

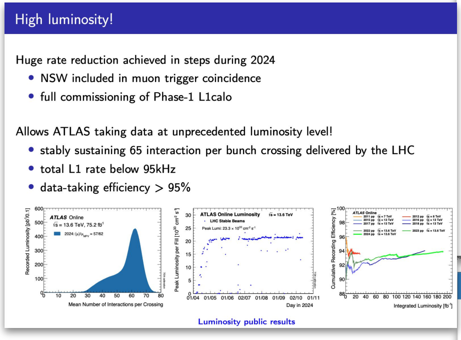

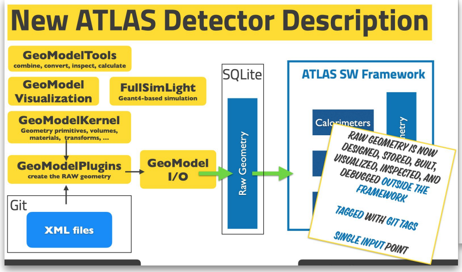

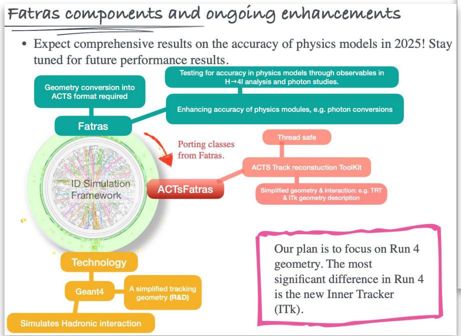

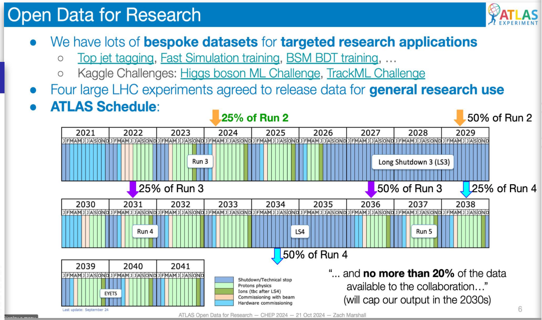

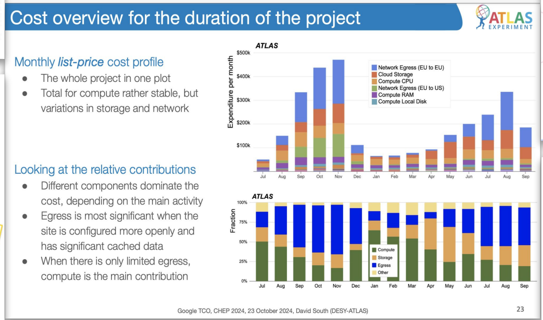

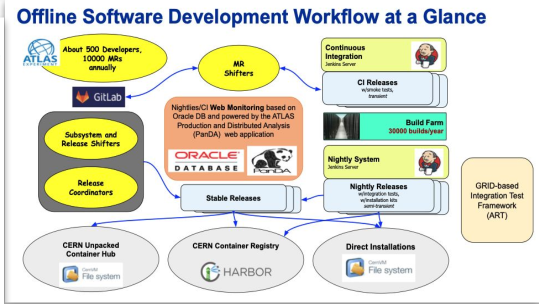

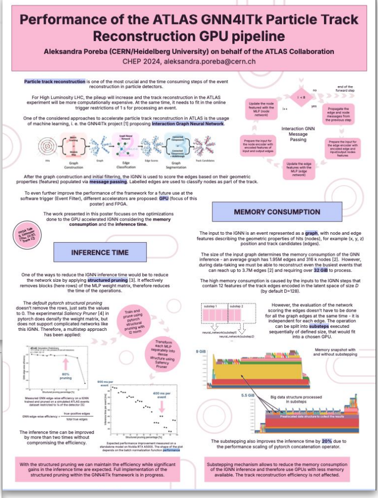

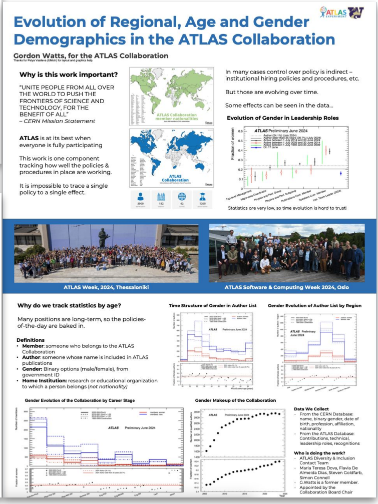

Track 1 - Data and Metadata Organization, Management and Access

- WLCG Data Challenge 2024 (25% of HL-LHC)
  - Observed bottlenecks, planning the next steps, 50% challenge in 2026
- Tape and Archive:
  - CTA evolution, smart writing, which data formats to store on tape, ATLAS Data Carousel

DC24

- **WLCG Data Challenges**: established in 2021 to prepare for HL-LHC data rates. DC24: ATLAS, CMS, ALICE, LHCb, BelleII (poster) and Dune

Year	% of HL-LHC
2021	10
2024	25
2026?	50
2028?	100

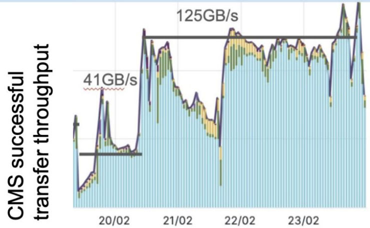

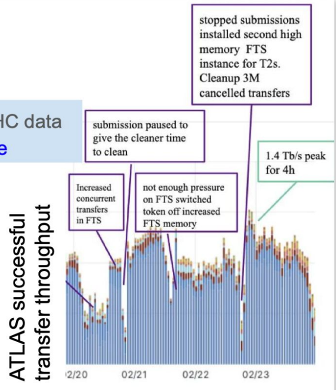

- Observed bottlenecks (ATLAS, CMS): scalability of FTS, Rucio and token handling, risk to overload storage systems
  - Data Movement Manager prototype interface between Rucio and the software defined networking (SDN) service SENSE by ESNet. Dynamic provision of network path
- DC26: 2xDC24 rates (50% HL-LHC), tokens established and battle tested, tape. Intermediate mini challenges
  - FTS3 Token Support for a Proxy-less WLCG world
- ALICE and LHCb run a program mostly focused on CERN to Tier-1s traffic; Belle II and DUNE also participated

Tape and Archive

- New **GridKa** Tape storage systems, switching from SP to **High Performance Storage System**
  - **aggregating** files from the same directory/creation time and computing+saving **checksum**

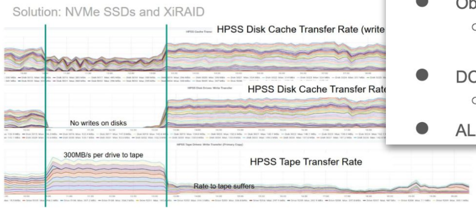

- **ATLAS** HL-LHC demonstrators with **Data Carousel** to improve tape bandwidth utilization
  - data-on-demand and **tape smart writing**, tests with ATLAS Tier-1s **FZK** and **RAL**
- A **Tape Remote Storage Element** in support of **LSST** backup challenge @SLAC (8 PB/year)
  - a specialized xroot server that orchestrates datasets backups, automatically creating a copy of a **Rucio** dataset as a single indexed file; works also for active restores

- **HTTP Archive Metadata** header added to **http data transfers** and received by tape endpoints to improve data **collocation** on tape and **tape scheduling**

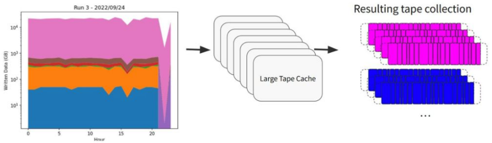

- Data movement and management:
  - Rucio & FTS is a standard now
- Data formats & metadata:
  - RNTuple is a hot topic, efficiency studies wrt TTree by ATLAS and CMS
- DB and filesystems:
  - AI in Oracle soon, CEPH is heavily used

Track 4 - Distributed Computing

- Moving from X.509 to tokens
  - The infrastructure is almost ready
  - Focusing on operation models
  - Improving IAM
- Monitoring & Operations
  - A lot of different technologies in use, dashboards-as-code is promising
  - CMS is trying to stick to MONIT stack & UEM
  - ATLAS HammerCloud with automatic recovery of blacklisted sites
- Optimizations
  - HEP score benchmark cross-check of sites core power
  - ALICE optimizing job placement & pushing sites to provide dedicated nodes for whole node scheduling
  - CMS improving CPU efficiency through Pilot overloading

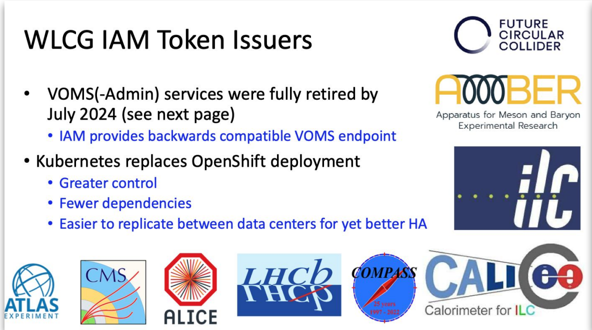

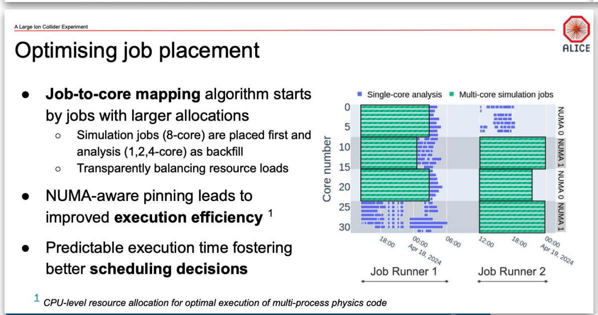

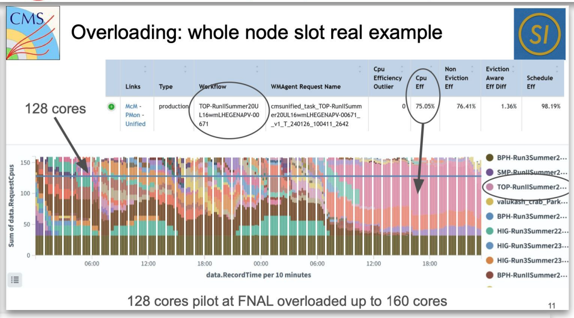

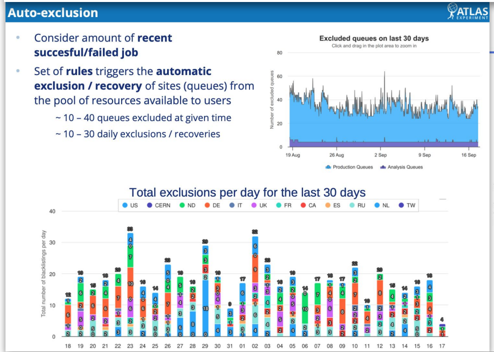

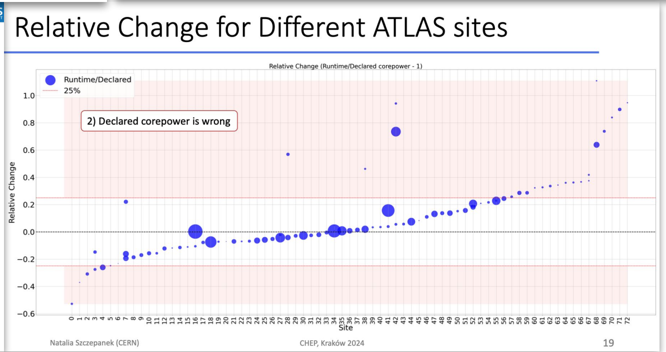

Track 4 - Distributed Computing

- Workflow management systems:
  - ATLAS PanDA is under refactoring & modernization
  - DIRAC is being rewritten from scratch, first release of DiracX is in 2025
  - ALICE JAliEn is evolving, started to support ARM, and has integrated HPC
- Outside HEP:
  - **SKA** (all purpose radio telescope) expects 400PB/year by 2030, using common tools like IAM, FTS
  - **HERD** (High Energy cosmic-Radiation Detection), start expected in 2027, using DIRAC & Rucio
  - **DUNE** (neutrinos) starts in 2029, expects 30PB/year, using CMS stack (HTCondor + GlideinWMS) + production system justIN and Rucio
  - **Vera Rubin Observatory** is in final steps of construction, start is in 2025, uses PanDA, Rucio, IAM

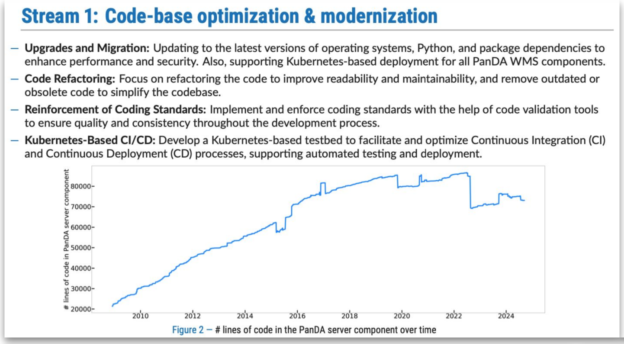

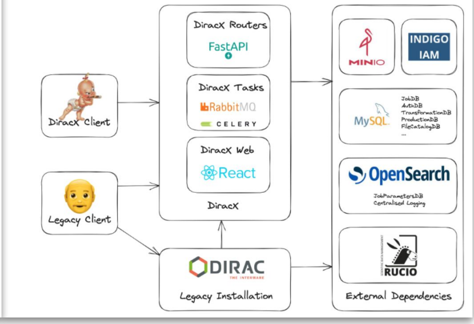

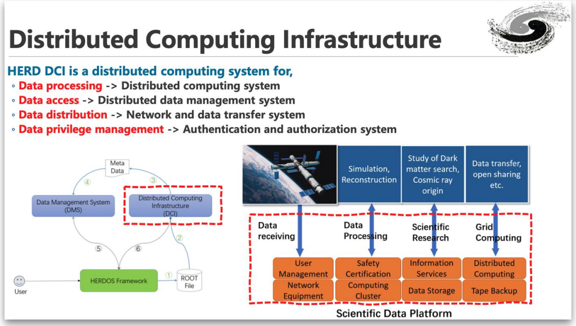

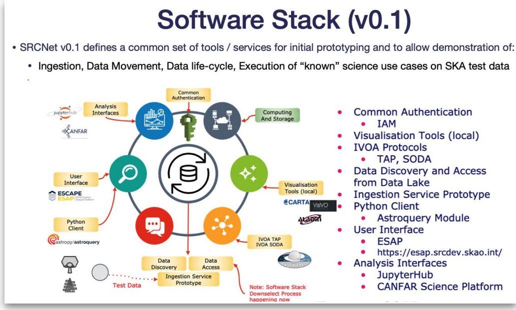

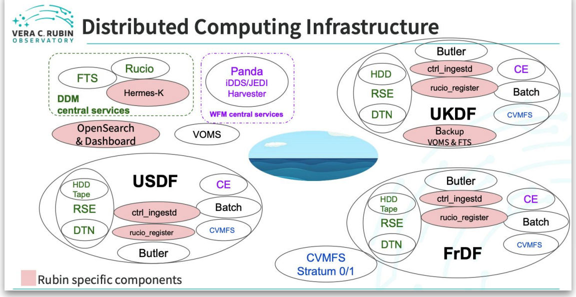

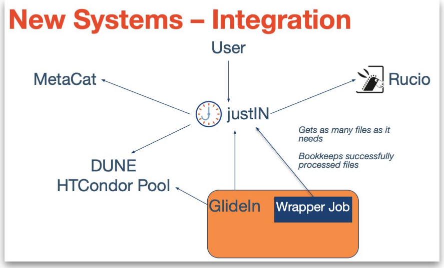

Other Hot Topics

- ML/AI and LLMs (~**40** contributions) for:
  - code review, operations, user support, anomaly detections, teaching & training , navigating conference abstracts, fast simulation in LHCb, CMS FlashSim, reconstruction for FCC detectors, photons classification in ALICE, simulating CMS HG Calorimeter, tracking in GlueX, triggers, AthenaTriton, MC tuning in JUNO, particle flow in ATLAS detector
- Quantum computing
- RNTuple format
  - Improves I/O performance and data compression

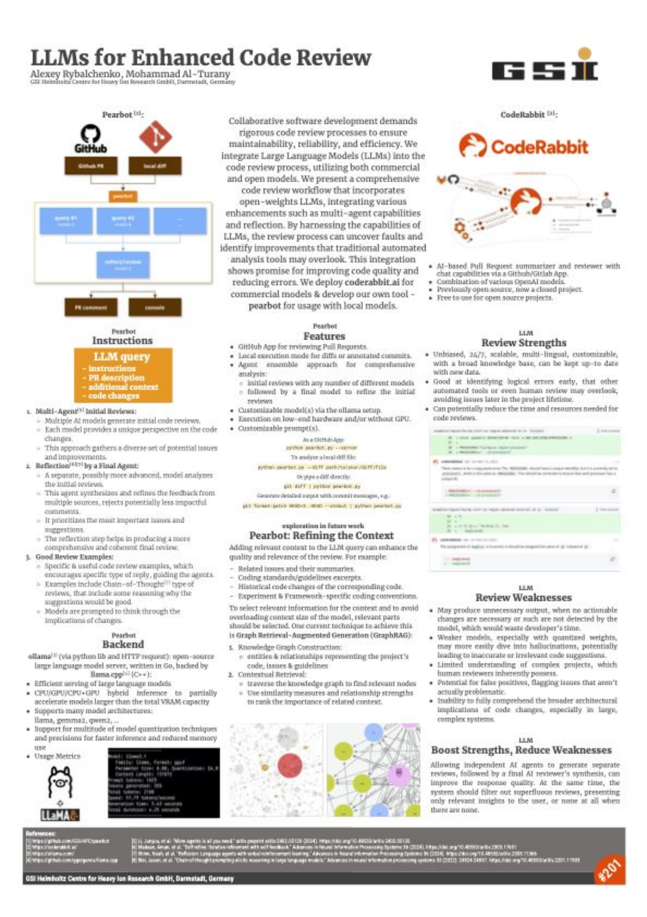

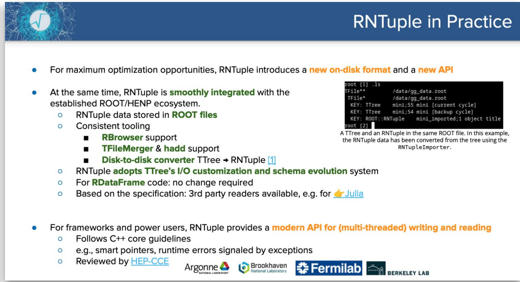

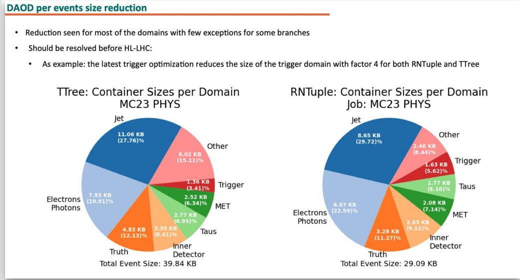

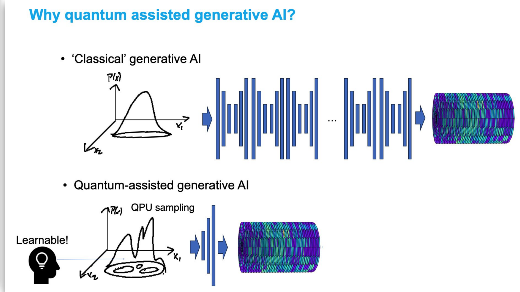

Other Hot Topics

- Kubernetes:
  - A lot of examples of performance improvement by moving from VMs to k8s-based deployment for both online and offline computing, and services like CERN SSO, dCache
- Analysis Facilities
  - Several national AFs provide notebook-based interfaces with possibility to scale out with Dask + Condor
  - Reproducibility with virtual research environment in REANA
  - Effort to federate AFs for US ATLAS with user session allocation across sites via network overlay + kubernetes
  - Successful 200Gbps challenge as part of Analysis Grand Challenge targeting HL-LHC scale analysis

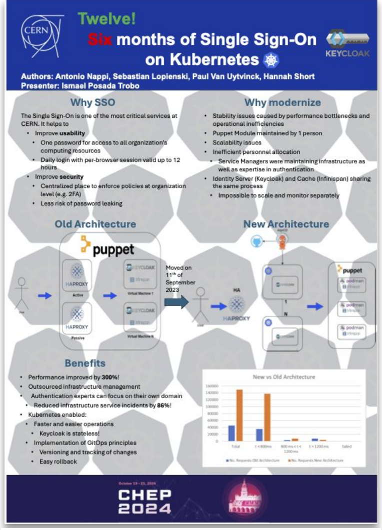

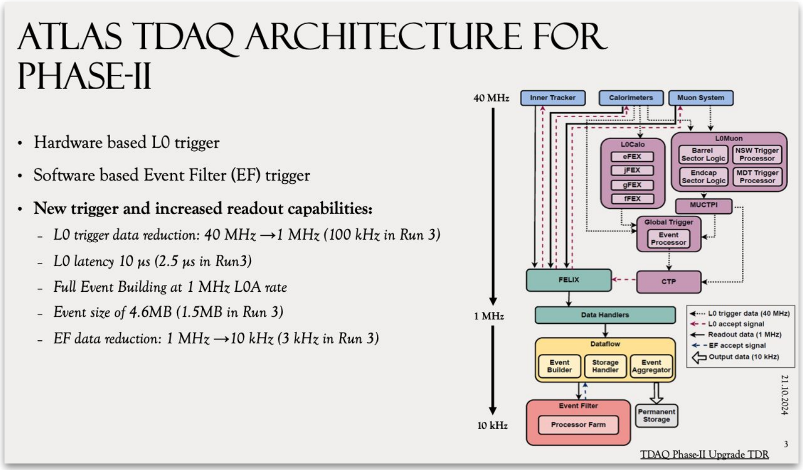

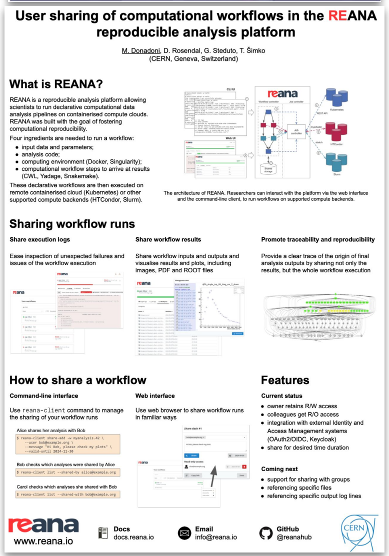

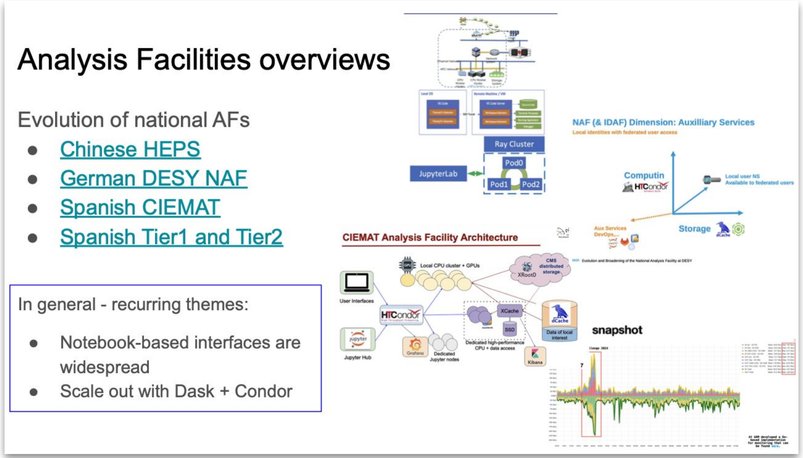

Summary

Very dense schedule

A lot of interesting talks and posters

For more please check timetable and contribution list

CHEP in numbers

5000 cups of coffee,
2500 cups of tea
1500 liters of water

176 kg of potatoes
470 kg of meat
140 kg of deserts

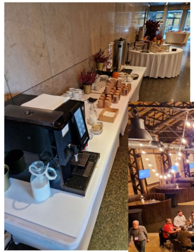

28th Conference on Computing in High-Energy and Nuclear Physics (CHEP 2026)

25-29 May 2026
Bangkok, Thailand

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

Questions?

Notes

True cache in Oracle 23ai - looks interesting

Indico IAM is also using AngularJS, they will move to ReactJS

Alice wants to use only whole node scheduling, they pushing sites to move to only this type of resource allocation

Dashboards-as-code with grafonnet (still alpha, to be released soon) for Grafana, UEM at CERN IT uses it