Track 9 Summary

Analysis Facilities and Interactive Computing

Enric Tejedor Saavedra
Marta Czurylo
Nick Smith
Nicole Skidmore

25.10.2024, Kraków, Poland

In short

- 3 parallel sessions on Wednesday and Thursday
- 18 talks in total
- Great atmosphere, many questions asked, follow up discussions…
- Here: quick summary split into topics - links to talks provided

Analysis Facilities overviews

Evolution of national AFs

- Chinese HEPS
- German DESY NAF
- Spanish CIEMAT
- Spanish Tier1 and Tier2

In general - recurring themes:

- Notebook-based interfaces are widespread
- Scale out with Dask + Condor

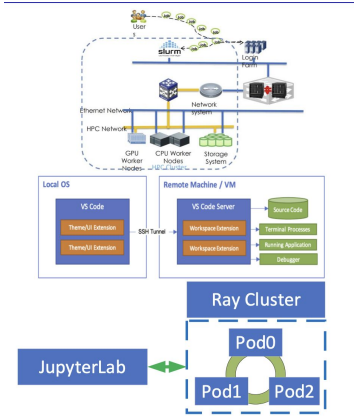

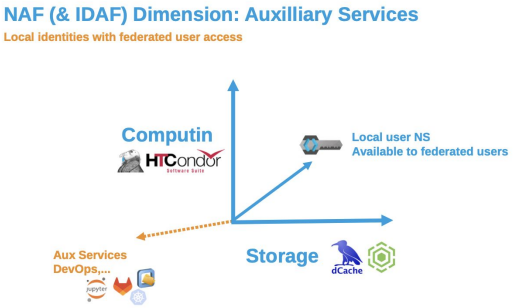

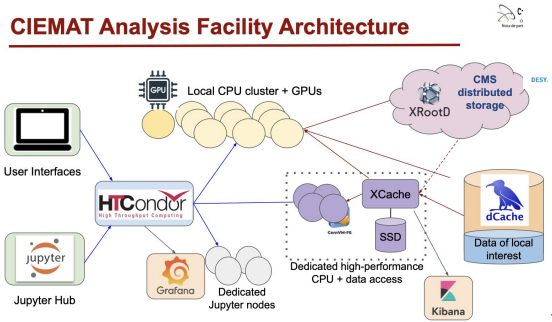

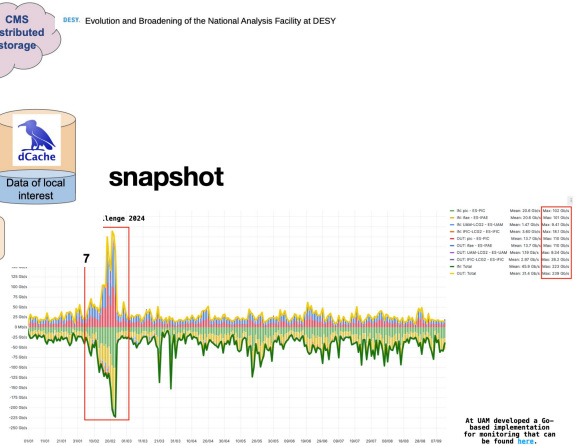

CERN Analysis Facility Pilot

- Available for **scale out of interactive analysis**
- Positive feedback from early testers

The pilot

- Focus on **scale out of interactive analysis**
  - On **already existing** CERN Batch system resources
  - Via RDataFrame / coffea + Dask

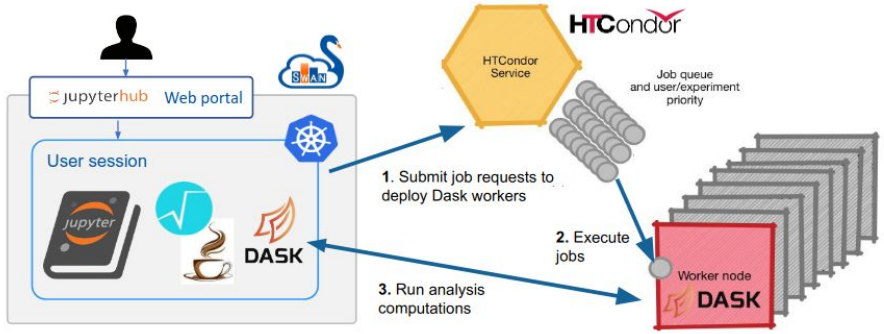

Distributed columnar analysis and its performance

- Distributed RDF used for real analysis for ATLAS and CMS

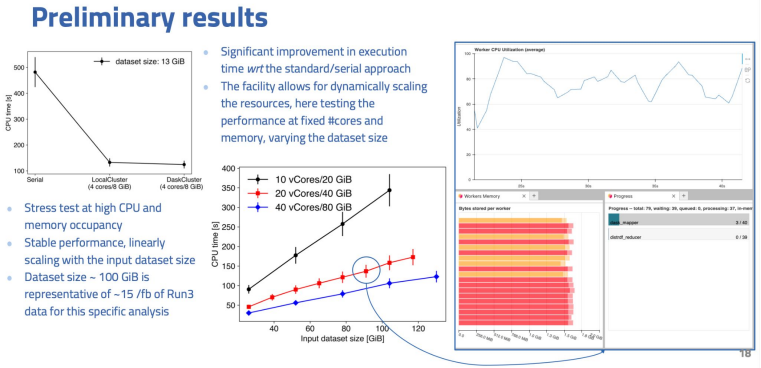

- With coffea: Performance optimization using TaskVine scheduler

- Distributed RDataFrame - current status and production readiness

Can we get out of Experimental ?

- User interface: **stabilised** and **unified**, with **easy inclusion of C++ code**
- Data input sources include **TTree, RNTuple** and **RDatasetSpec**
- The more **Pythonisations** in ROOT, the better the distributed RDF
- Distributed RDF is performant in many different **Analysis Facilities**

Performance optimization

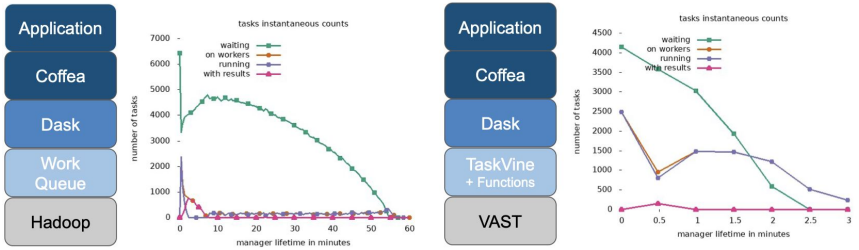

Significant performance improvement within same application by swapping out bottom two layers (task/data scheduler and file system

Talk at SC24 ← Barry Sly-Delgado

US ATLAS AFs federation

- Effort to **federate Analysis Facilities** for US ATLAS (user session allocation happens across sites via network overlay + kubernetes)

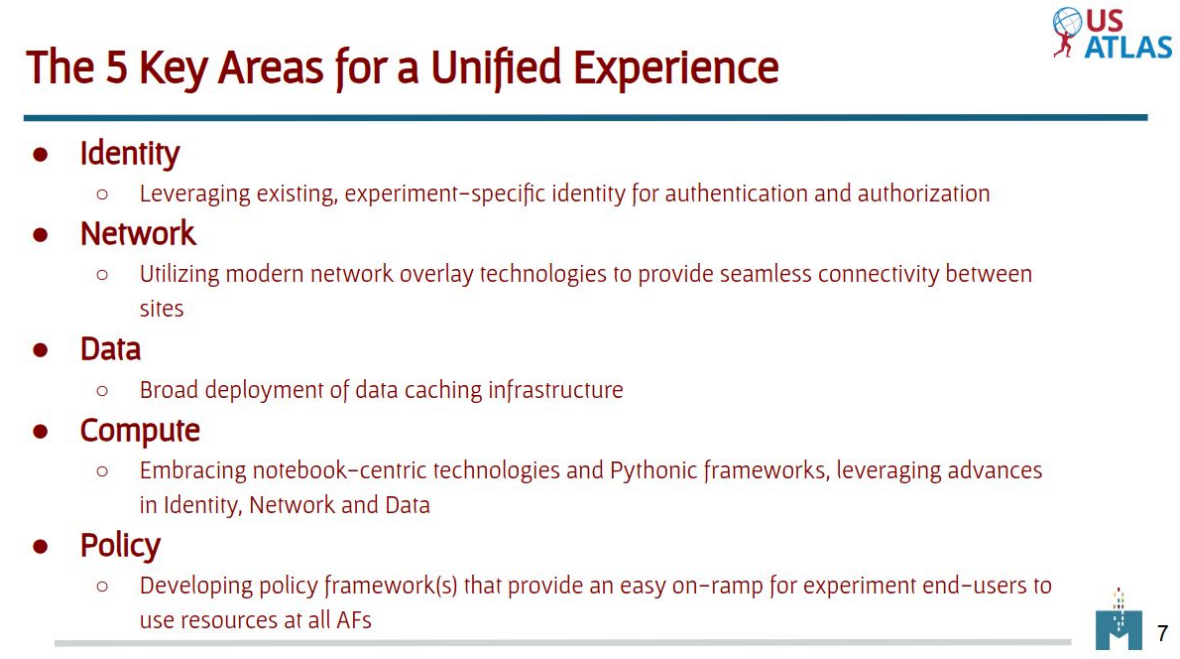

ML and GPU computing in Analysis facilities

AI_INFN
Artificial Intelligence technologies for INFN research

Machine Learning Training Facility

- Training of the models → reproducibility, scalability, efficiency
- MLFlow as the training framework

On-grid GPU development

AI_INFN's primary activities:

- facilitating access to **HPC and GPU resources**
- organizing **educational and training events**
- fostering the **AI community** within INFN
- conducting **R&D** to integrate new technologies (*e.g.*, FPGA, quantum computing) into the platform

GridPP
UK Computing for Particle Physics

University of Glasgow

On-Grid Interactive GPU development

Many problems to solve

Problems vaguely fall into 2 categories:

GPU problems

- Variety of cards (what to target)
- Variety of software tools
- Large dependancy issues
- Because of their nature user jobs tend to be small enough to not need to scale out to the Grid
- The idea is to facilitate on-Grid development and reduce the overhead in submitting Grid jobs via a submission engine

Interactive job develop problems (Analysis Facilities?)

- User authentication
- Flexible development environment
  - —> allow users to install packages
  - —> maintain site security
- Data storage / integration
- Scalability, i.e. easily scale out to the rest of the Grid

Preparation of AFs for 200 Gbps challenge

- AF **administrator perspective** on how the challenge was prepared and operated

CMS

Dask + HTCondor stats running over 1300 cores:
Rate over time and runtime to access each file

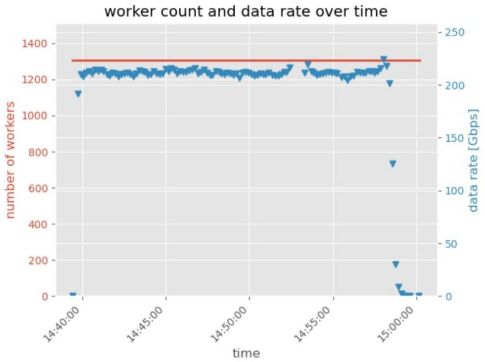

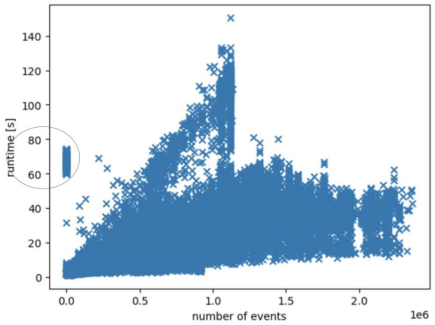

event rate (aggregated time spent in function): event rate (aggregated time spent in function): **27.66 kHz**

11

ATLAS

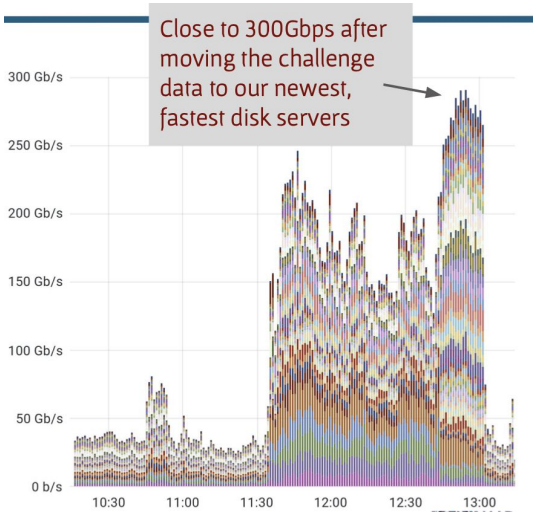

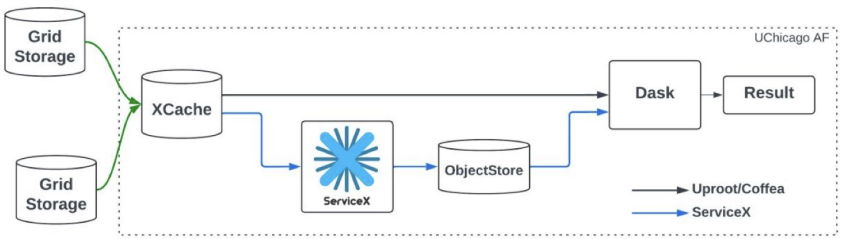

Reproducibility (REANA)

VRE at CERN

- Virtual research environment
- Open Source bottom up, modular approach towards AF
- Check out the demo under this link

Scientific analysis stages can be performed in the VRE

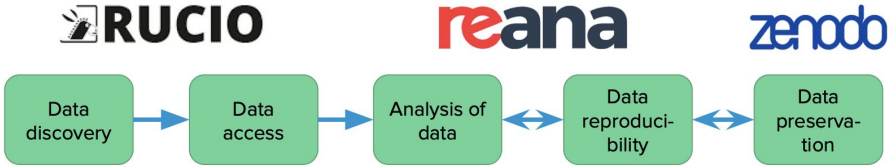

Data discovery, analysis and reproducibility in VREs - CHEP 2024 - 24 Oct - Enrique Garcia

AGC with Snakemake and REANA

Setting up REANA test cluster

Kubernetes v1.30 cluster with 53 nodes

- one master node
- three REANA infrastructure nodes
  - web server, database, message broker
- one workflow orchestration node
  - Snakemake
- **48 job-running nodes**

Node flavour

- 8 vCPUs
- 15 GB RAM

Final results

Runtime reduced from 14m26s to 3m54s (3.7x faster) when tested with 48 nodes

Reached 323 peak concurrent jobs, from initial 102 jobs (3.2x more) when tested with 48 nodes

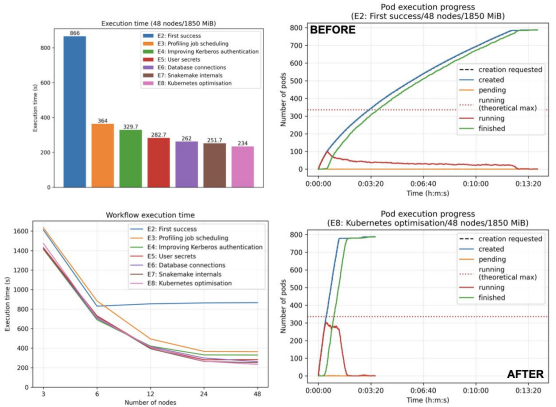

Scientific interactive computing – languages and tools

- Python, Julia and awkward array
  - Combining Python and Julia
  - Find areas at which Julia could outperform other tools

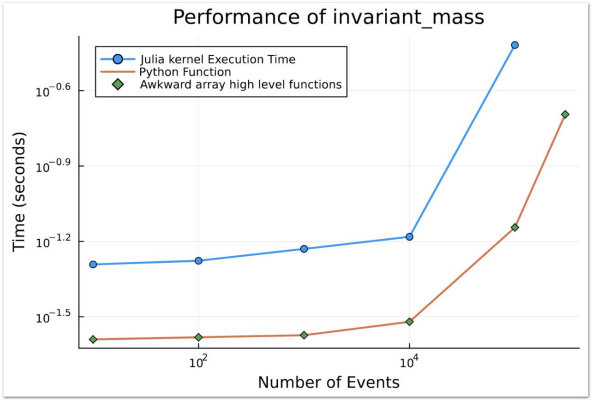

Ianna Osborne, CHEP 2024, Krakow, Poland

- Web based ROOT graphics
  - JavaScript ROOT that is great for interactivity with added functionalities and improved performance.
  - Introduced a large scale CI testing within ROOT

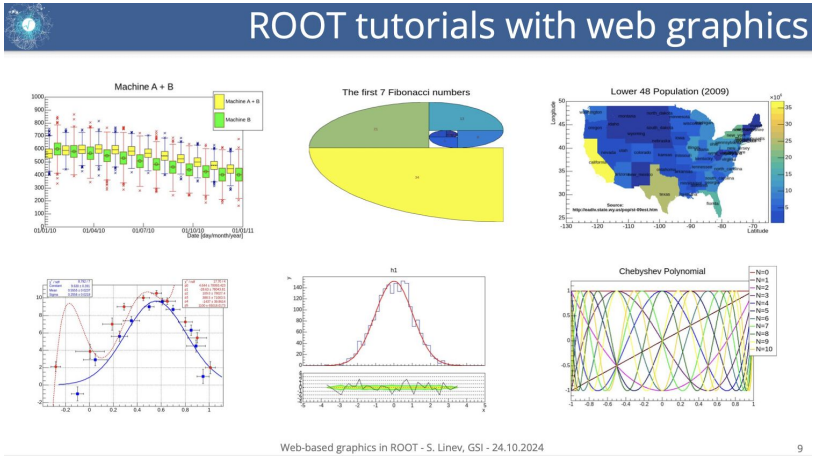

- In the ROOT session
  - root --web=chrome hsimple.C

We would like to thank the speakers, participants and the organisers!