October 19 - 25, 2024

CHEP 2024

Aksieniia, Carmelo, Daniel, Daniele, Enrico, Luca, Lucia, Matteo

Codice interessi

- ✨ AI
- ☁️ Distributed systems
- 🚜 Farming
- 🏢 Infrastruttura
- 🛜 Rete
- 🔒 Security
- 🧑‍💻 Sviluppo software
- 💾 Storage
- 🧑‍🔧 Supporto utenti


CHEP
2024
Computing in High Energy
& Nuclear Physics

Plenary Session

October 19 - 25, 2024
CHEP 2024
Conference on Computing in High Energy and Nuclear Physics

Hardware technology trends in HEP computing

Speaker: Andrea Sciabà

- **Technology tracking at CERN and in HEPiX**
- **WLCG evolution and technology**
  - WLCG needs a good understanding of technology evolution for its medium-long term planning
  - A five years gap to prepare for HL-LHC - difficult to make sensible predictions on this timescale

WLCG CERN HEPiX

CPU cost extrapolation at CERN

- CPU price decrease rate relatively stable over time
- **AMD becoming again competitive gave a boost**
- Stagnant recently
  - Is Intel going to be competitive again?
  - When (if) will Arm start having an impact? Or RISC-V?
  - When (if) will GPUs will make CPUs less relevant?

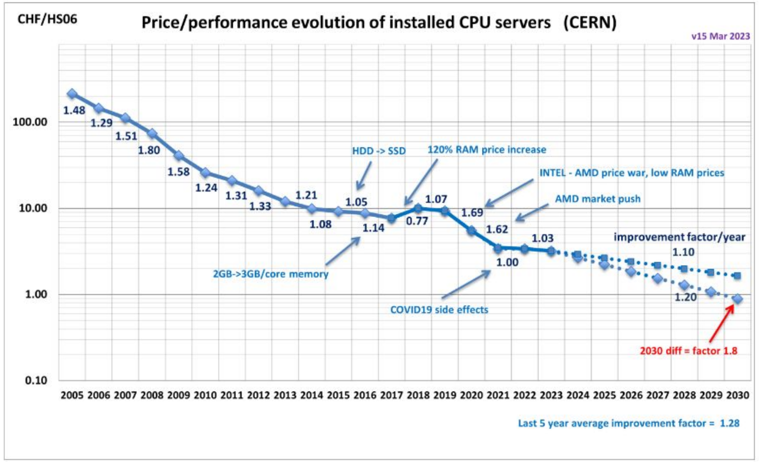

Source: B. Panzer-Steindel

Disk cost extrapolation at CERN

- **Stable price decrease but also flattening out recently**
- **What about SSDs?**
  - Not expected to completely replace HDDs, but usage will certainly increase and impact the overall cost of storage

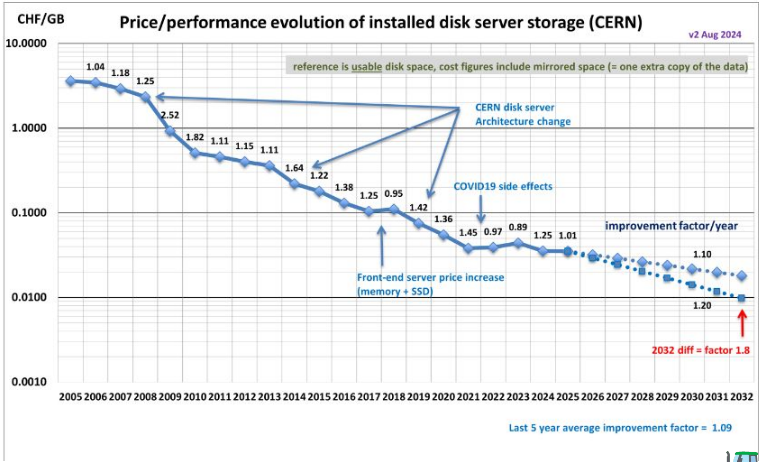

Challenges in resource requirements

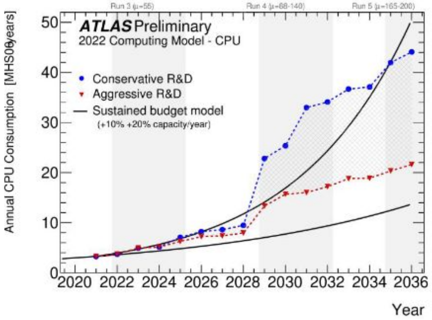

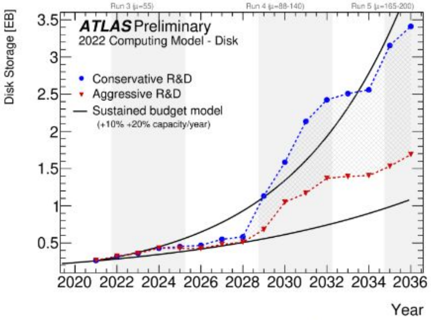

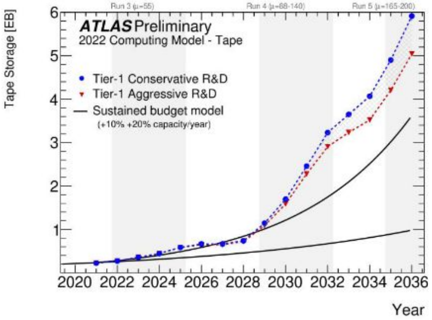

Sources: ATLAS and CMS

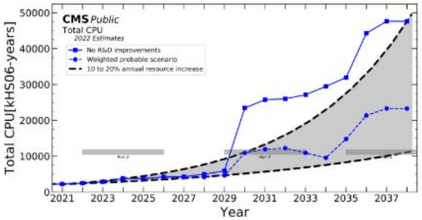

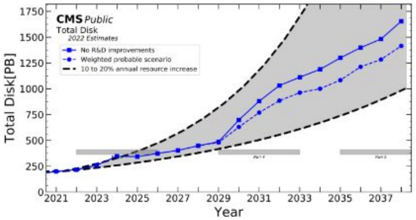

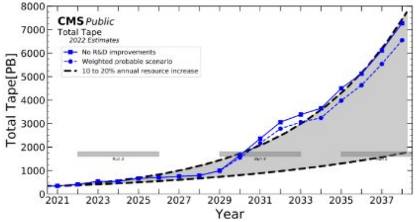

- From the experiment plots, with a flat budget, a 15%/year price reduction is required even in the most optimistic case (for CPU and disk)
- **Stressing the importance of technology awareness**
  - i.e., how to get the most value out of the existing technology

Power consumption

- **Data centers are proliferating due to the AI boom**
  - AI models increase by 1000x every three years
  - A single GPU uses ~4 MWh per year, Nvidia sold ~4 million in 2023! And soon there will be 2 kW GPUs…
  - Power infrastructure is a big constraint, strong incentive to energy efficiency
  - Data centers used ~2% of global electricity in 2022, estimated twice as much in 2026
- **The global IT market is increasingly focusing on sustainability**
  - This is becoming a hot topic also in our community
    - Dedicated WLCG workshop in December

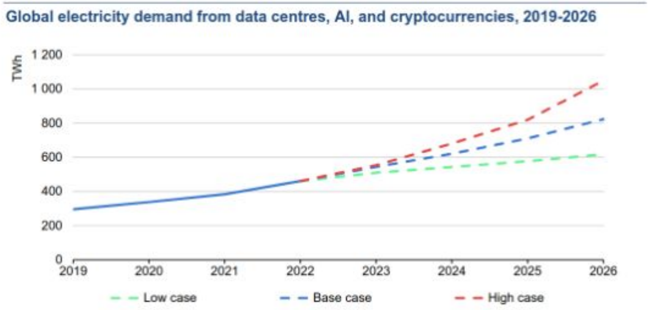

Source: IEA Report 2024

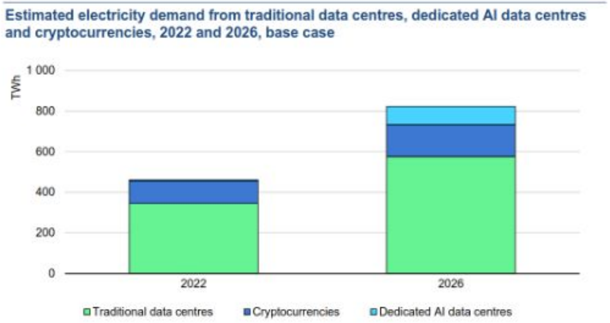

Many server hardware components are rising in price due to the AI boom

- **Revenues steadily increasing in the last few years, with a few exceptions**
  - Intel dropped quite a bit and lost most of its value, future uncertain
  - Nvidia skyrocketed in 2023
- **AI is the main driver and Nvidia has a practical monopoly**
  - AMD might increase their share, as demand is very high and have competitive products

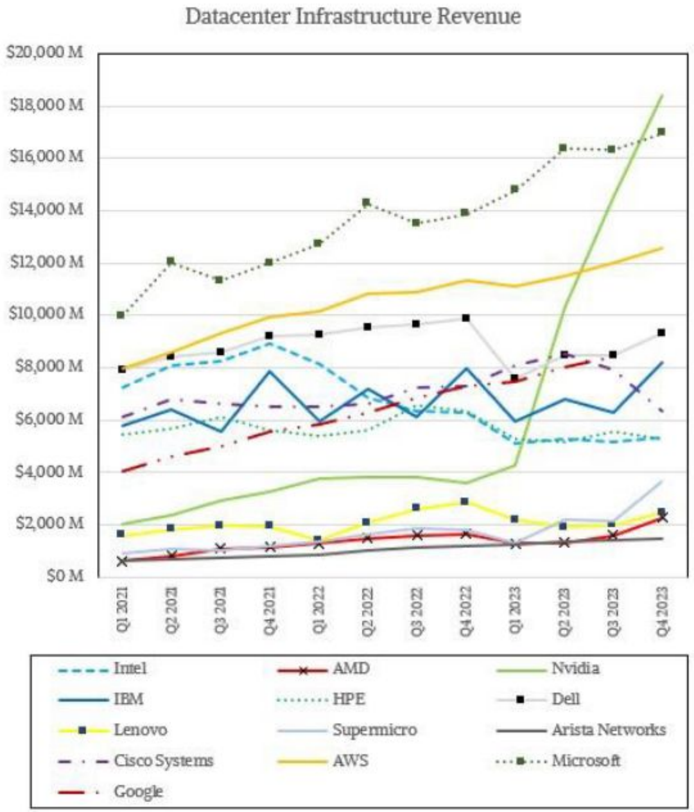

Server Market

- **Server shipments are expected to slightly increase in 2024**
  - AI servers are 12% of the total and increasing at a much faster rate than general purpose servers
- **AMD quickly gaining ground**
  - 34% CPU server revenue market share in 2024
- **Arm also increasing, 3x in 3 years**
  - Best suited for hyperscalers and cloud providers
  - But Ampere's revenues are a tiny fraction of the total, and Nvidia Grace is very expensive
  - Still early for us to heavily invest on it

Server designs

- **High core counts make single socket servers a very interesting option**
  - Simpler, cheaper, use less power
- **Arm servers are becoming a viable alternative**
  - Better power efficiency than x86 (more on this later)
- **Liquid cooling destined to become mainstream**

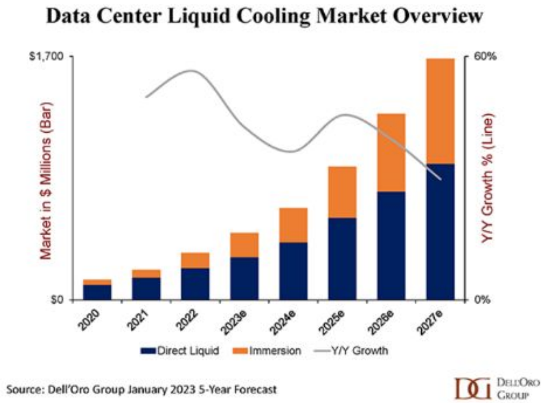

Current and upcoming x86 CPU generations

	AMD 4th gen EPYC "Genoa"	AMD 4th gen EPYC "Bergamo"	AMD 5th gen EPYC "Turin"	Intel 5th gen Xeon (Emerald Rapids)	Intel 6th gen Xeon "Sierra Forest"	Intel 6th gen Xeon 6 "Granite Rapids"	Intel "Clearwater Forest"
Launch	2022 Q4	2023 Q3	2024 Q4	2023 Q4	2024 Q2	2024 Q3	2025
Node	TSMC N5	TSMC N5	TSMC N3	Intel 7	Intel 3	Intel 3	Intel 18A
Max Cores	96 Zen4	128 Zen4c	128 Zen5 192 Zen5c	64	144 (288 next year) E-cores	128 P-cores	288?
Max L3 cache	384 MB	256 MB	384 MB	320 MB	108 MB	504 MB	?
Max TDP	360 W	400W	500 W	350 W	500 W	500 W	500 W?
Memory	12 ch DDR5 up to 4800 MHz	12 ch DDR5 up to 4800 MHz	12 ch DDR5 up to 6400 MHz + CXL 2.0	8 ch DDR5 up to 5600 MHz	Up to 12 ch DDR5-6400 CXL support	Up to 12 ch DDR5-6400 MCR-DIMM and CXL support	?
I/O	Up to 160 IO lanes of PCIe-5	Up to 160 IO lanes of PCIe-5	Up to 160 IO lanes of PCIe-5	Up to 80 IO lanes of PCIe-5	Up to 96 IO lanes of PCIe-5	Up to 96 IO lanes of PCIe-5	?

- **Choosing a specific CPU model is not easy**
  - WLCG has HEPSCORE23 as primary tool to measure performance
  - It will be interesting to determine if E-cores can be a viable option for HEP
  - Intel CPUs feature several built-in accelerators (AVX, AMX, IAA, DSA, DLB, QAT, …), probably not very useful to us

CPU	Best HS23/phys. cores
AMD Genoa	43
Intel Sapphire Rapids	40
AMD Milan	36
Nvidia Grace	32
Ampere Altra	16

What about Arm?

- **Arm entered the server market in 2018 and sells CPU designs**
  - Neoverse N1 in 2019: Ampere Altra (80 cores), AWS Graviton2
  - Neoverse N2 in 2020: Microsoft Azure Cobalt
  - Neoverse V1 in 2020: AWS Graviton3
  - Neoverse V2 in 2022: AWS Graviton4, Nvidia Grace, Google Axion
  - AmpereOne in 2024: up to 192 cores (custom design)
- **Nvidia main competitor with the Grace CPU**
  - Combined with the Hopper GPU or two CPUs in a "superchip" (144c)
- **Ampere CPUs already deployed at a few WLCG sites**
  - Notably Glasgow, published several efficiency measurements and comparisons with x86
  - Clearly better power efficiency!
- **Not quite yet a valid alternative for WLCG**
  - Not all experiment workloads are validated on Arm
  - Ampere is a very small company, Grace is very expensive
  - Both Intel and AMD are pushing strongly on power efficiency
- **And RISC-V??**
  - Not yet a viable platform for WLCG, but it is ramping up fast!

GPUs and accelerators

- GPU usage in HEP quickly gaining momentum, but so far mostly on dedicated facilities (HPCs, HLT farms, analysis facilities, …) or for R&D
- Currently, almost an Nvidia monopoly, but AMD is gaining ground
- **Performance evolution is not going in a direction we like**
  - **FP32/FP64 performance will not increase much, or at all**

Storage evolution summary

- **To summarize:**
  - AI boom drives volume increase for all types of storage
  - HDD shipments will soon be almost only nearline, but increasing in capacity
  - SSDs will not replace HDDs in data centers anytime soon
  - Our usage of SSD will probably increase to cope with the performance bottlenecks of HDDs
  - Tapes are not going anywhere either

FIGURE 7. PROJECTION OF DRIVES BY MARKET NICHES (1,000'S-UNITS)

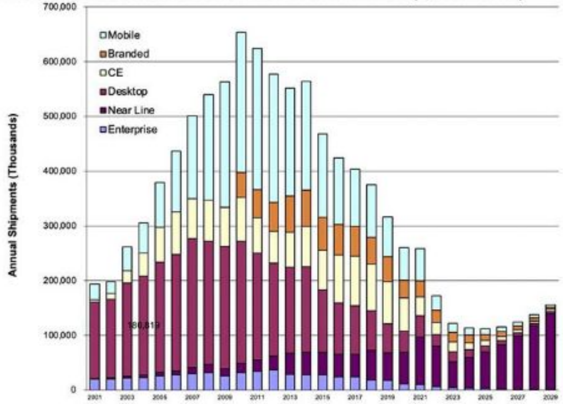

Source: [Blocks and Files]

FIGURE 11. CAPACITY SHIPMENTS FOR LTO TAPE, SSDS AND HDDS

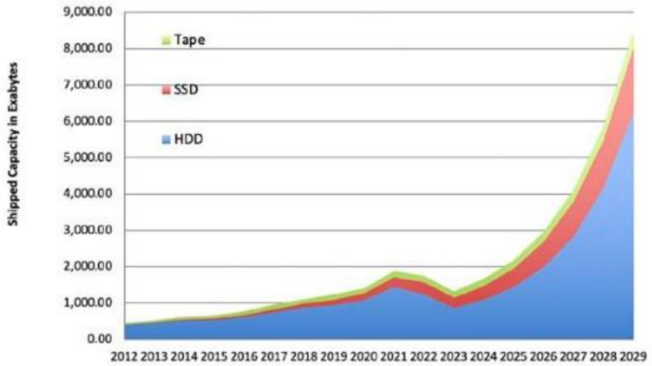

Network

- **Network bandwidth correspondingly increasing on LAN and WAN**
  - To cope with increase in cores and storage/server
  - For LHC, driven by HL-LHC data rates: LHC network traffic exponentially increasing, will need Tb/s links on major routes by 2029
    - **Aggregate network traffic from ATLAS + CMS will be O(10 Tb/s)**

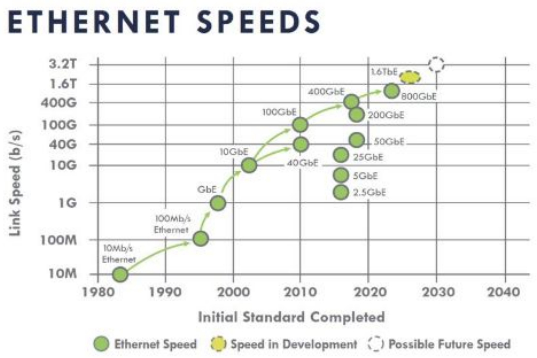

Track1: Data and Metadata Organization, Management and Access

- Conveners
  - Lucia Morganti, INFN-CNAF
  - Tigran Mkrtchyan, DESY
  - Sam Skipsey, University of Glasgow, UK
  - Ruslan Mashinistov, BNL
- 41 talks, 1 plenary, 15 posters

Data challenge (s)

- Scitags ready for production (and partially adopted during DC24)
- Data Movement Manager for SENSE-Rucio:
  - Current model of push-now-worry-later for data transfers will not be feasible in the near future
  - What if we could have controlled data-flows for the largest datasets which dominate network usage?
  - SDN (Software Defined Networking) decouples the control plane (how data is routed) from the data plane (how data moves), allowing dynamic, programmable networks.
  - SENSE is a SDN-based service enabling automated and dynamic provisioning of network paths

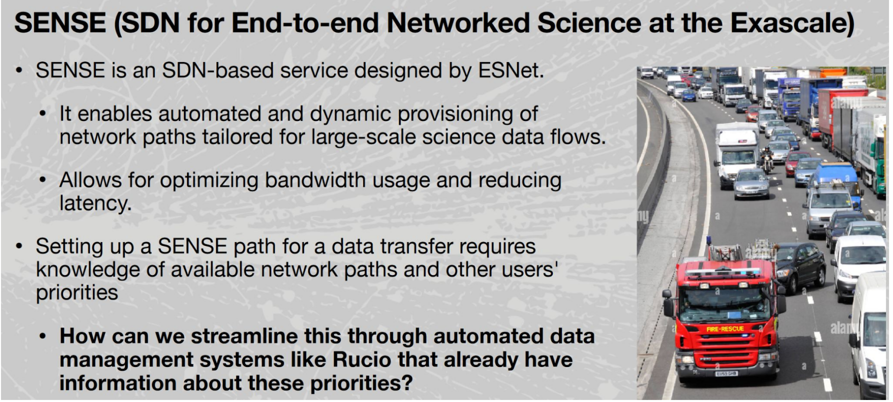

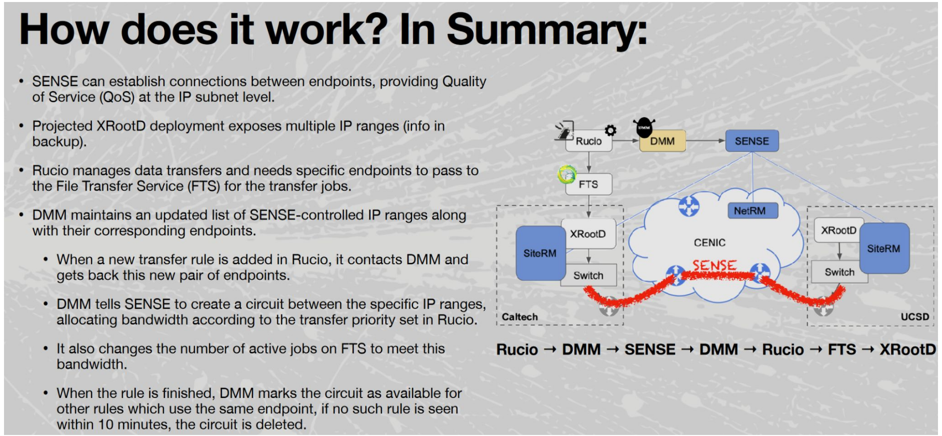

Networking

Packet and flow marking with SciTags (StoRM WebDAV has opt-in support)

IPv6-only grid

SDN for LHC experiment and NOTED

INFN Roma 1 developed a FPGA-based NIC

- ALICE implemented [EOS O2] (Run4 prototype), the largest single storage system at CERN
- [RAL] tried (and failed) to achieve 100 Gb/s data transfers with XrootD
  - Connected storage (CephFS) likely to be the bottleneck (also a [talk] on slow ops with CephFS from GridPP)
- RAL presented a new XrootD [load balancer] (+ DNS round robin limitations)

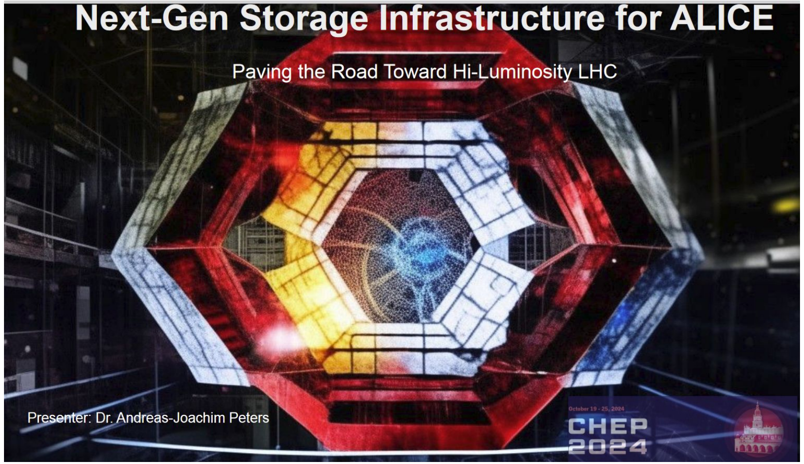

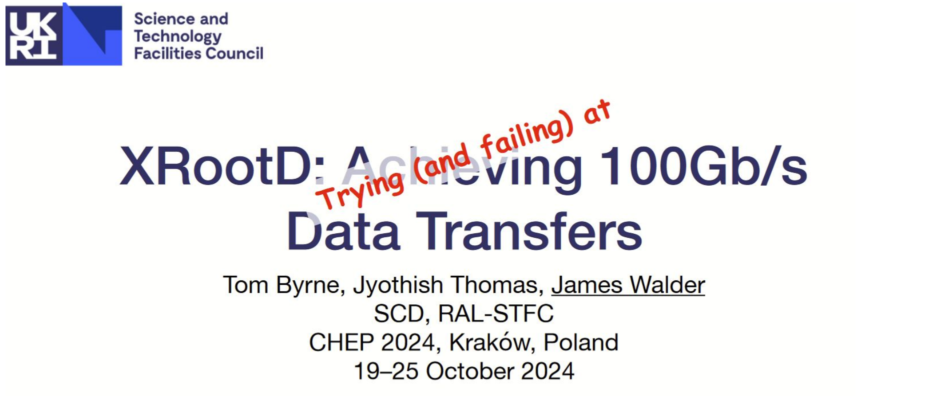

- Plenary talk di Katy, brand new
- Talk Track 1 di A. Forti (ATLAS)
- Talk Track 1 di Wissing (CMS)

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

Intermediate mini data challenges focused on sites, technologies

MANCHESTER 1824

GridPP UK Computing for Particle Physics

Results explained

- FTS orchestrates **transfers per link over many links**
  - **Doesn't orchestrate throughput**
  - To increase throughput we had to increase the number of allowed parallel transfers by an over an order of magnitude
- Has a concept of fair share per activity
  - Doesn't have a concept of links priorities within an activity, i.e. all links are equally treated T0-T1 same level as T2-T2
    - Could prioritise faster transfers or more important channels
- Testing also new authz system with tokens put further load on the system

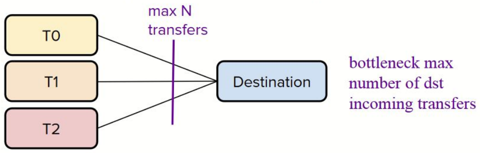

- **Agreed with FTS** problems to solve first for next challange

Comparison: Rates achieved vs. targeted – Tier 1s

		JINR		FNAL		IN2P3		RAL		PIC		KIT		CNAF	
Day	Scenario	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC	DEST	SRC
1	T0 Export	1.42	N/A	1.13	N/A	1.09	N/A	0.76	N/A	1.18	N/A	1.16	N/A	1.17	N/A
2	T0 Export	1.46	N/A	1.12	N/A	1.10	N/A	0.50	N/A	1.17	N/A	0.94	N/A	1.17	N/A
3	T0Export, T1Export	1.31	0.62	1.08	0.88	1.33	1.03	0.72	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	1.28	0.93
4	T1 Export	N/A	0.37	N/A	0.91	N/A	1.12	N/A	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	N/A	1.00
5	T1-Export, Prod-out	1.18	1.72	1.15	0.87	1.25	0.89	0.98	1.01	1.21	1.09	1.23	0.77	1.17	0.77
6	T1-Export, Prod-out	1.14	2.42	1.18	0.88	1.47	0.88	0.72	0.81	1.17	1.03	1.19	0.76	1.18	0.95
7	T1-Export, Prod-out	1.19	2.19	1.15	0.87	1.22	0.87	0.81	1.04	1.20	0.98	1.21	0.73	1.16	1.02
8	AAA	1.30	N/A	N/A	1.10	1.39	N/A	1.31	N/A	1.31	N/A	1.70	N/A	1.32	N/A
9	All	0.38	0.34	0.87	0.84	0.57	0.57	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	0.56	0.65	0.25
10	All	0.70	0.34	0.98	0.74	0.58	0.65	0.56	[illegible]	0.70	[illegible]	[illegible]	0.98	0.63	0.28
[illegible]	[illegible]	0.63	0.33	0.91	0.73	0.43	0.76	0.77	1.05	1.09	0.84	0.91	1.09	0.69	0.24
12	All	0.40	0.54	0.92	0.86	0.89	1.00	0.85	1.15	1.21	0.87	1.13	0.89	0.78	0.29

LHC OPN link to RAL disturbed

Tuning of FTS

Typically storage overloaded

Typically storage overloaded

Ratio = observed/targeted

Green - ratio > 0.9

Yellow – 0.9 > ratio > 0.7

Orange – 0.7 > ratio > 0.5

Red - ratio < 0.5

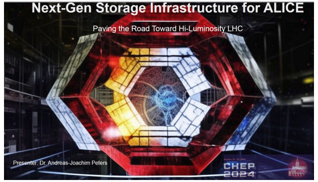

ALICE O^2

- **ALICE O^2 is**
  - the largest single storage system at CERN - **150 PB**
  - the most stable physics storage system at CERN - **100% availability June to October**
  - the only 100% erasure coding physics storage with the smallest overhead **+20%**
  - the most performant storage system at CERN **380 GB/s wr/rd + 700 GB/s rd/**direct
  - the most cost-effective storage system at CERN - still performant when 98% full
  - **really great!**
- **Lessons learned**
  - in a large distributed storage system a single hardware issue can lead to significant and unpredictable performance loss - you always learn something in production!
- **Future Outlook**
  - O^2 is a prototype for **Run-4** storage systems online & offline environments
    - the performance capacity ratio has to be followed with care with ever growing HDD capacities
      - Possible mitigations
        - 400G or bonded 200GE
        - Less disks per front-end [96 => 60 HDDs] at the cost of having more front-end nodes

EOS as Storage System for ALICE O^2 - Value & Failure Rates

Virtual value created if **O^2** would be market price cloud storage **97 M$**

Cost-Matrix	TB*Years	Avg-Drive-Hours	Tot-Drive-Hours	Cloud$-Replica	Cloud$-Erasure
eosaliceo2	466.41 K	23.59 K	281.63 M	58.30 M$	97.17 M$

Model: 20$ / TB*Month
(with 3 year subscription Azure is cheaper!

Azure Calculator   Google Calculator   AWS Calculator

64 disks removed in population of **12k** in **10 month** = **Failure rate** 0.64%

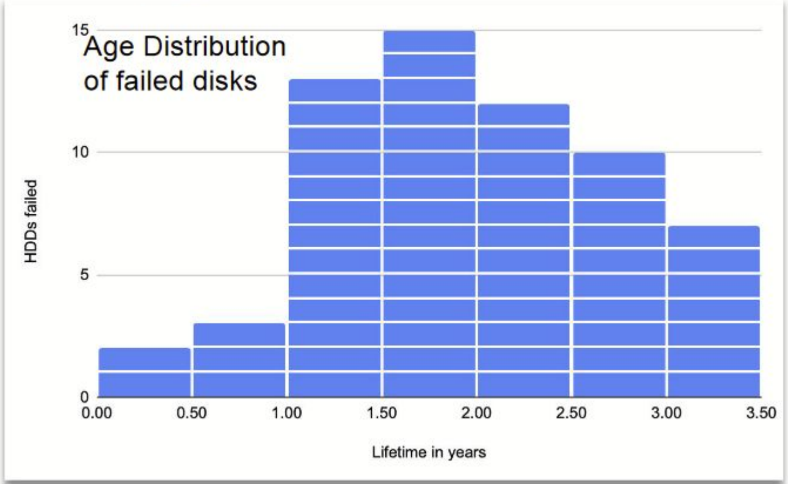

Also a plenary talk: The ATLAS Google project
Running a cloud site (europe-west1 in Belgium)
in production from July 2022 to September 2023

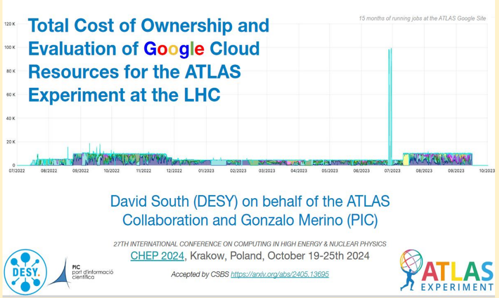

ALICE O^2 running 93% full with 10% of HDDs at 100%

December 2023

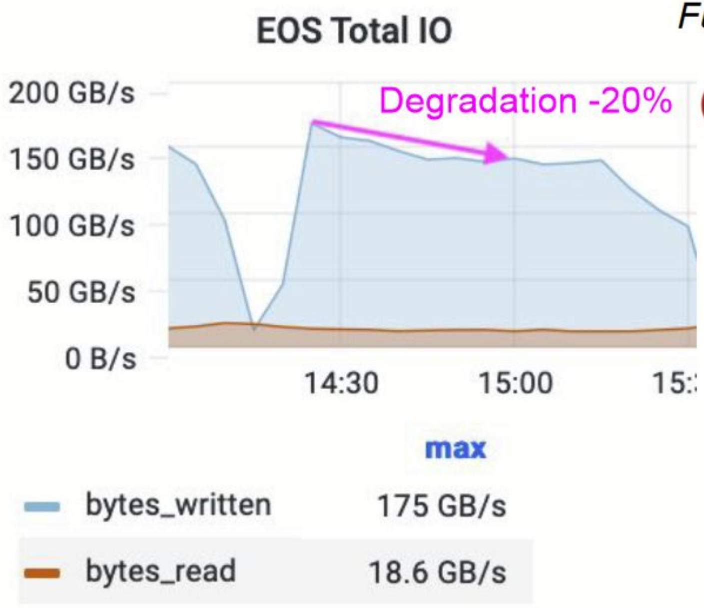

Full chain test from EPNs to O^2

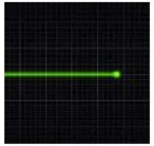

Rebalancing

170 GB/s stable for 1h

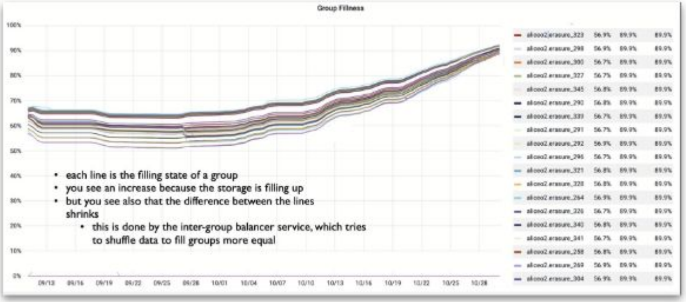

- Running with a very full instance bears the risk, that part of the disk population is full at close to 100%, **write bandwidth decays** due to unusable scheduling groups
- **Rebalancing** full instances **stabilizes performance** because even at 99% all HDDs can still write data

Tape archival

- CTA: new scheduler (and tape storage numbers for CERN)
- Interesting talk on the challenges of repack
- GridKA migrated from SP to HPSS
- Archive metadata (CERN)

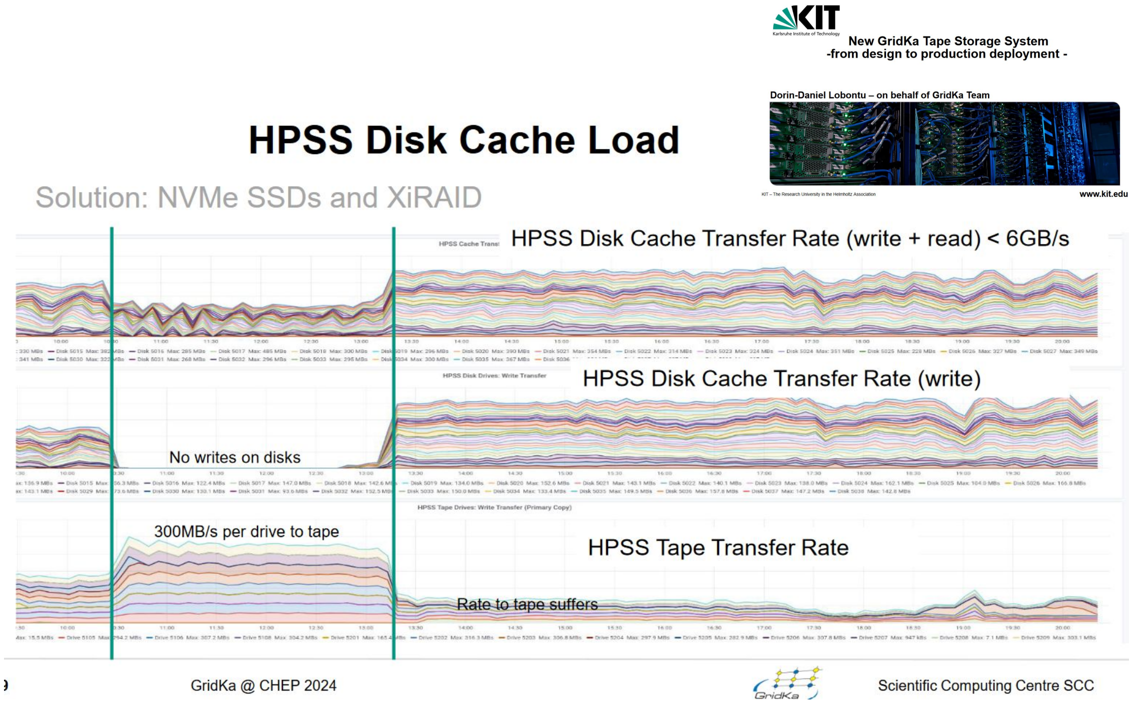

Passing Archive metadata

- **Archive metadata travels through the full experiment data management stack to reach tape storage endpoints**

- **ATLAS + Rucio example:**
  - Rucio generates Archive metadata
  - Add it to every FTS transfer to tape
  - Tape endpoint receives ArchiveMetadata HTTP header **before the first byte of data is received**

Tape software can decide what to do with upcoming file content

CERN tape collocation constraints

- **T0 legacy tape collocation mapped on experiment directory structure**
- **T0 tape is low latency very high throughput**
  - 1 tape family for RAW using time based collocation
- **At T0 strict separation or RAW data by tape family by dataset would add too many constraints**

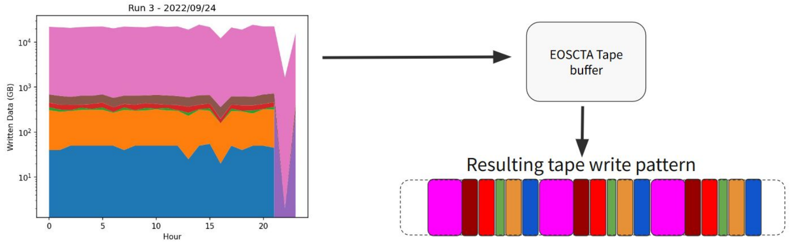

Tier1s tape collocation constraints

- **T1s receive out of order delayed transfers from T0**
- **T1s rely on strict tape families to demultiplex streams**
  - Many more tape families needed than T0
  - Logically grouping data transferred over multiple days requires
    - large and expensive tape caches in HSM
    - additional hints that signal logical set completion to trigger flush to tape

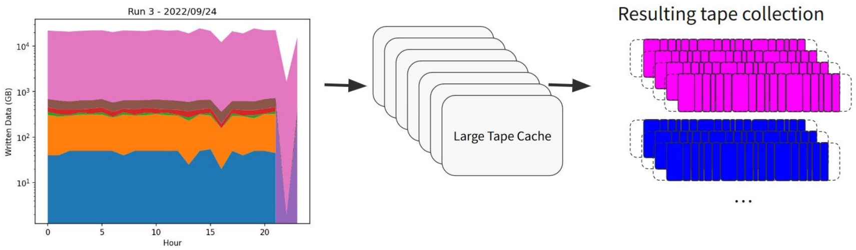

Improve tape scheduling

- **CMS:**
  - RAW must go ASAP to tape
  - MonteCarlo can wait for beam dump

Would allow to move large chunk of total T0 traffic outside of peak

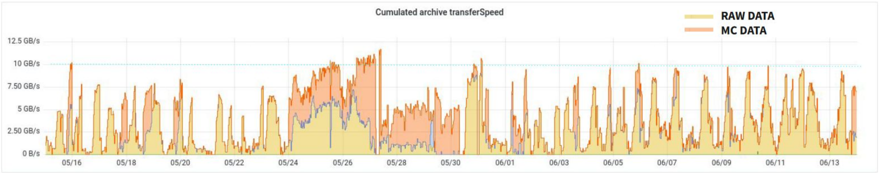

Toward a common solution for Archive Backpressure?

There is no point accepting files in tape buffer/cache if their time to tape is expected to exceed agreed Service Level Agreements or compromise site tape operations

- **scheduling_hints**
  - **archive_priority: "0" to "100"**
    - "0" is lowest priority, 100" is highest
    - Rucio policy deduces value from activity
  - If bandwidth to tape is too constrained
    - Exceeding allocated experiment pledge
    - Sudden loss of bandwidth (tape hardware failure on site,...)
  - Allow to apply *backpressure* on archive transfers
    - Protect RAW data transfers

```
"scheduling_hints": {
    "archive_priority": "100"          # highest priority
},
```

Software, data movement and management

Track 1: Data and Metadata Organization, Management and Access

- dCache, StoRM (Luca), Rucio
- Several problems for XrootD shoveler: too early for roll-on over many sites
- Proxyless WLCG world, Mihai
- In general, wide scale adoption of Rucio, FTS, tokens: Dune, SKA, ET, LSST, Belle II

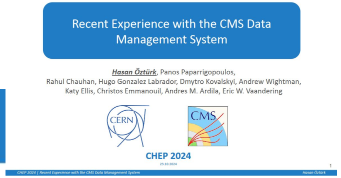

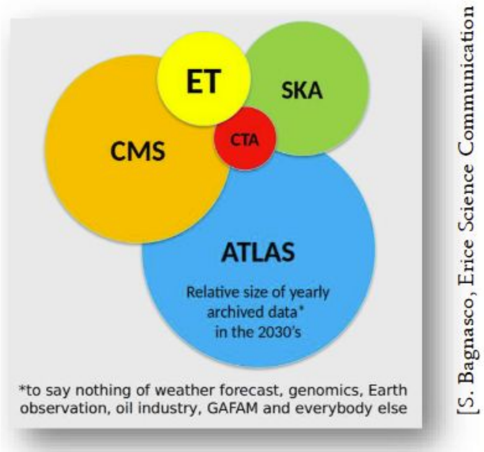

[S. Bagnasco, Erice Science Communication and Journalism International School 2023]

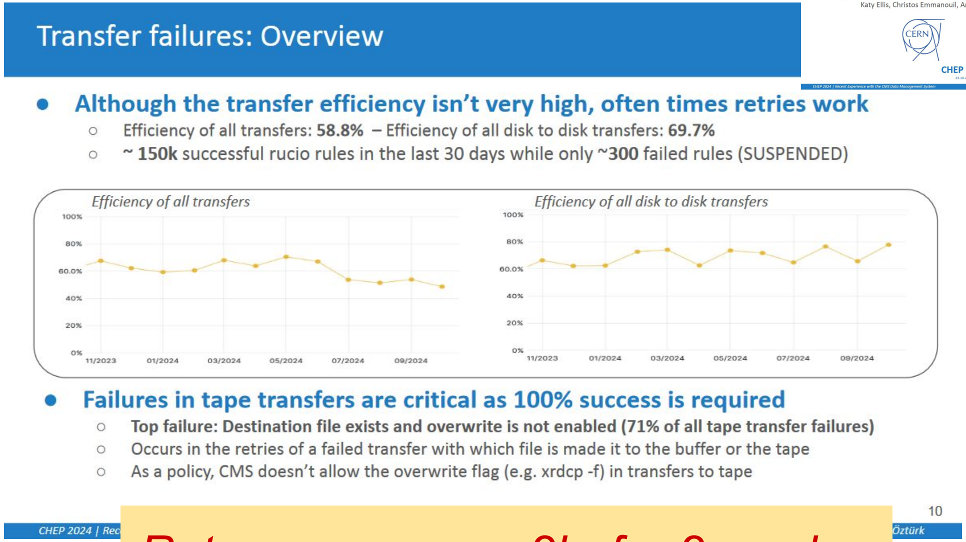

Retry: once every 2h, for 2 weeks

dCache Project Updates

Tigran Mkrtchyan for the dCache collaboration

Protocols and Instances

- Data is going to grow… A lot…
  - High ingest data rates.
  - More movements between sites.
- Shared Computing Resources
  - Analysis Facilities.
  - Grid Farms.
  - HPC.
  - Cloud resources (CPU&Storage).
- Standard analysis tools
  - ROOT.
  - Jupyter Notebooks.
  - Non-ROOT tools.
- Competing Tape Operations

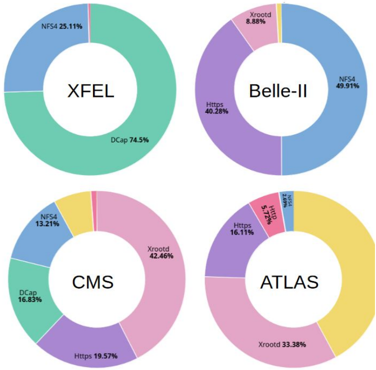

CTA Integration

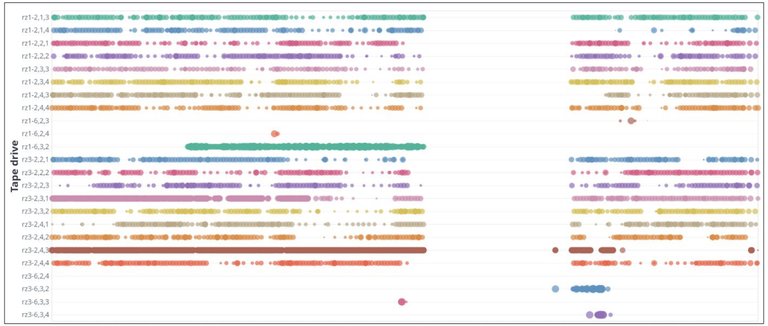

For details, visit our Poster: *dCache-CTA Integration: One Year in Production at DESY, Oct 22*

CI-Pipeline

- Migration from Jenkins to Gitlab-CI.
- The build/test process is ‘documented’ in repo.
- Helm charts to deploy dCache in k8s.

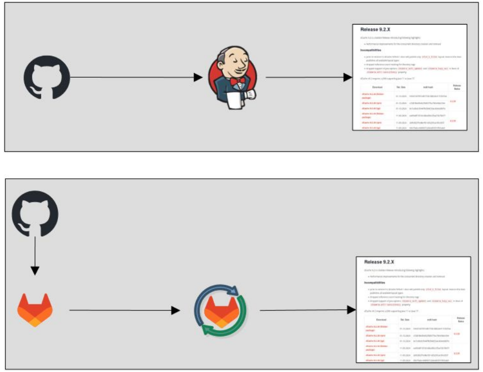

dCache Built-in Java Flight Recorder

- A low overhead (~2%) profiler built into JVM
- Simple admin commands to start/stop recording
  - More sophisticated configs can be enabled with java tools.
- The collected recordings can be shared with developers for later analysis
- Detailed information in the *Debugging* section of “The dCache book”

More from Track 1

- *“>70% of researchers have failed to reproduce another scientist’s experiment, >50% have failed to reproduce their own”* (https://www.nature.com/articles/533452a, quoted in a talk about AMI ATLAS Metadata Interface)
- With respect to data formats, several talks about RNTuple for ATLAS and CMS, and a plenary talk
- Cache and ML, J. Flix
- Ceph@CERN: backup, cross-region consistent storage
- Multi-tiered storage systems: data placement according to data temperature
- Estimate of carbon cost of storage. Following up, plenary: Carbon compromises: minimising the Carbon-Cost and Power-Use of Grid sites
    - WLCG Environmental Sustainability Workshop

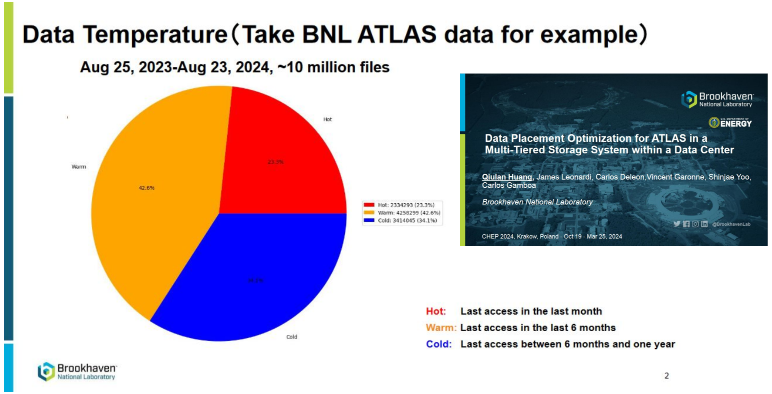

https://dataviz.boavizta.org/serversimpact

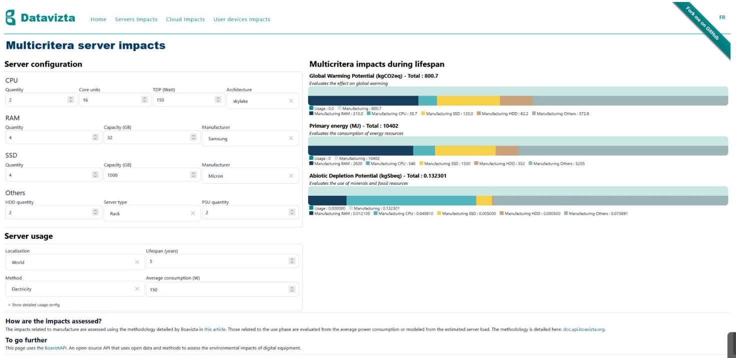

New programming languages for HEP

Much interest about using Julia in HEP:

- Easier than C++, more performant than Python
- Many wrapper to library used in experiments, e.g. Fast JET and AwkwardArray

Gaudi has integrated Rust in the C++ codebase

Improved multithreading in Python 3.13 (wrt GIL), UpROOT leverages this to improve performance

Energy efficiency between languages (but also ATLAS, UKRI, ARM and RISC-V)

🧑‍💻 ☁️

Next RUN

Transition to tokens of Fermilab and Tier-1 (Carmelo) and obviously IAM (Enrico)

100 GB/s with XRootD and load balancing

Rucio

Condor usage is ATLAS and Belle II

Miscellanea

https://spack.io

dCache using K8s for CI/CD, virtual directories and other updates

XRootD integration tests with K8s

ROOT JS library to show graphs and transition to RNTuple

Spack for container building (also at Fermilab and Key4hep) and image snapshotters

Quantum computing: Tuesday plenary and track reconstruction (Laura Cappelli)

VEGA integration with Tier-1 and INFN Cloud (Claudio Grandi)

Norwegian/Swedish grid

🧑‍💻 ☁️ 💾

Track 7: Computing Infrastructure

Opportunistic resources, orchestration of virtual machines and containers; cloud; HPC and exascale; networking; computing centre infrastructure; energy efficiency; cost of computing; management and monitoring; quantum networks.

Conveners:

- Henryk Giemza (National Centre for Nuclear Research – PL)
- Flavio Pisani (CERN)
- Christoph Wissing (Deutsches Elektronen-Synchrotron – DE)
- Bruno Heinrich Hoeft (KIT - Karlsruhe Institute of Technology – DE)

Presentation materials:

A Lightweight Analysis & Grid Facility for the DARWIN Experiment

Speaker - Robin Hofsaess (KIT - Karlsruhe Institute of Technology (DE))
Primary author Sebastian Brommer (KIT - Karlsruhe Institute of Technology (DE))

This discuss about the architecture of the facility, its provided services, first experiences of the DARWIN collaboration, and how it can serve as a sustainable blueprint for other collaborations.

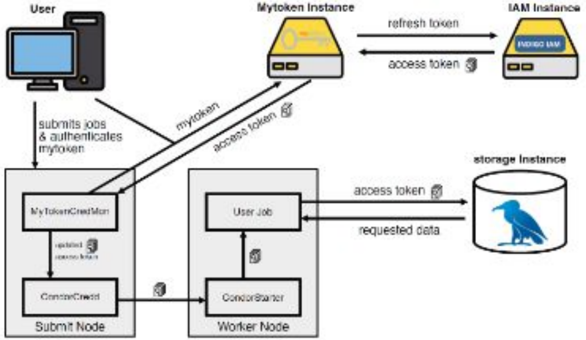

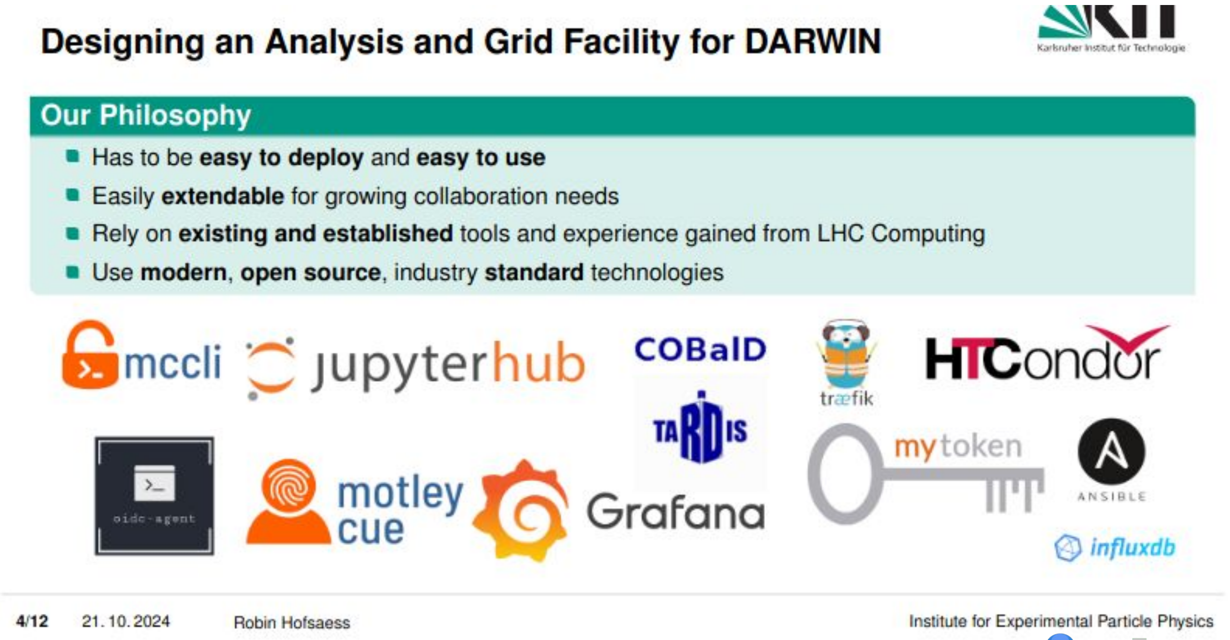

Presentation materials:

The BaBar Long Term Data Preservation and Computing Infrastructure

Speaker & Primary author Dr Marcus Ebert (University of Victoria)

The computing system was put in production in 2022 and this talk described its infrastructure, based on cloud compute in Victoria, Canada, data storage at GridKa, Germany, streaming data access, as well as the possibility to analyze any data from anywhere.

Analysis Environment - Overview

- user accounts/management
- batch system
- shared file system
- data transfer machine
- XRootD system
- redundancy needs to be built in

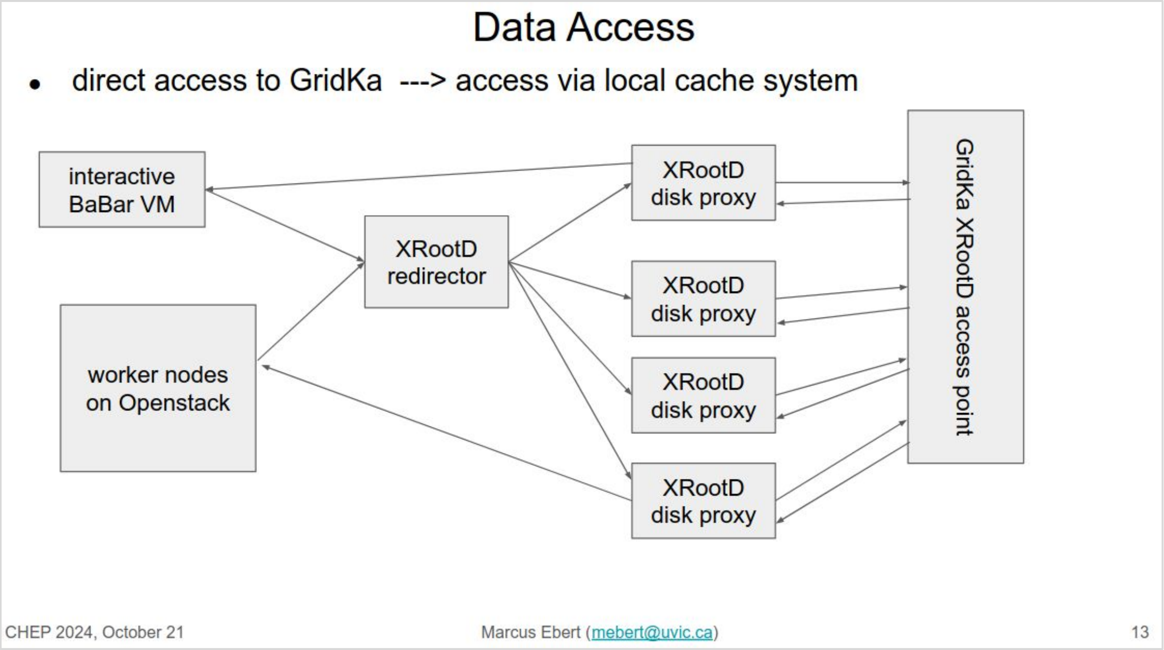

Presentation materials:

The Big Data Processing Infrastructure for monitoring and analysing the ATLAS experiment processing activities at INFN-CNAF Tier-1

Speaker - Aksieniia Shtimmerman (INFN-CNAF) **Co author** Giacomo Levrini (Universita e INFN, Bologna (IT))

Within this framework, the first data pipeline was established for the ATLAS experiment, using input from the ATLAS Distributed Computing system PanDa. This is involve examining the performance metrics of the machines and identifying the log errors that lead to job failures.

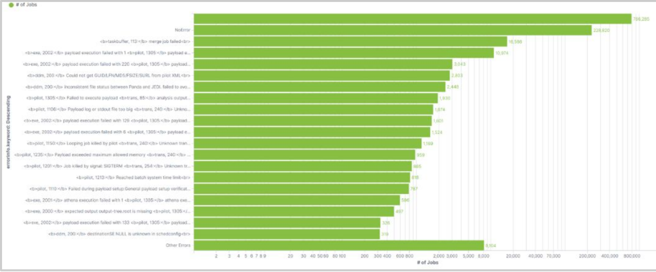

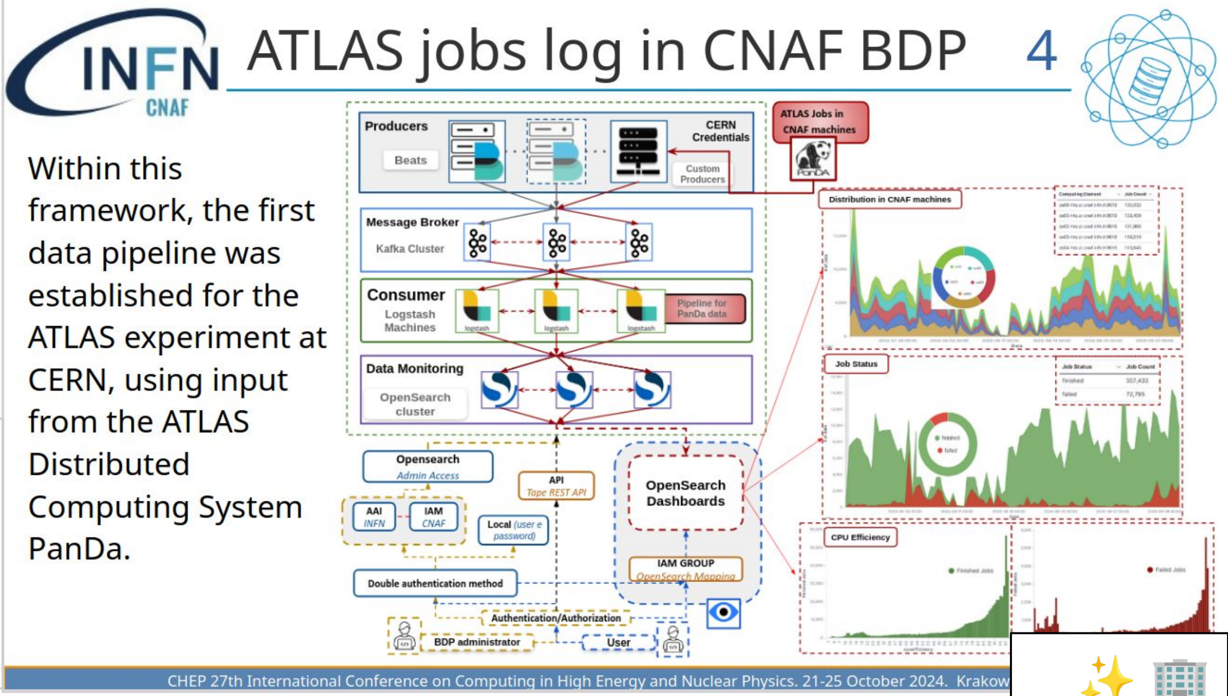

Presentation materials:

Monitoring large-scale dCache installations with storage events using Kafka streams

Speaker - Christian Voss
Primary author Mr Tigran Mkrtchyan (DESY) , Marina Sahakyan, Christian Voss

In this talk, present aggregation and analyses pipelines an on demand Apache Spark cluster on top the National Analysis Facility at DESY and workflows and how they are enabling DESY IT to scale out dCache storages for

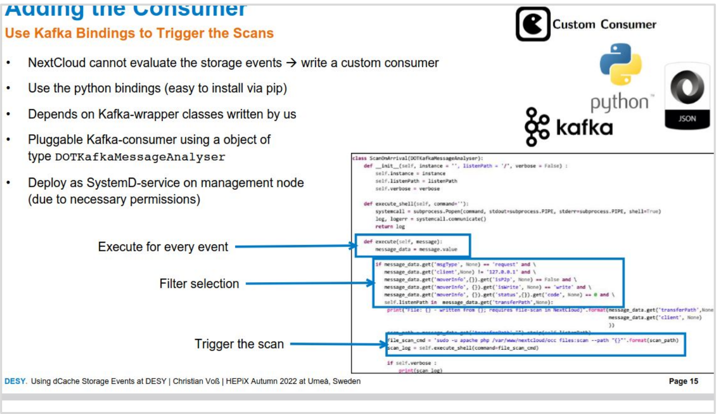

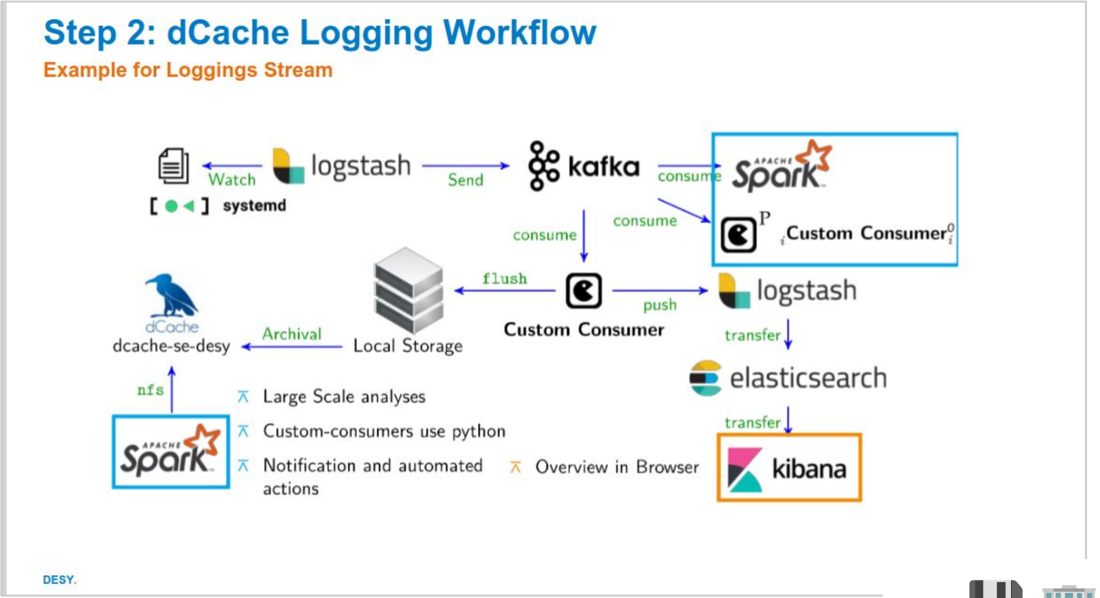

Presentation materials:

A Successful Data Centre Refurbishment Project

Speaker & **Primary author** Daniel Peter Traynor

In this talk, present how Queen Mary University of London (QMUL) has recently finished refurbishing its data centre computing cluster supporting the WLCG project.

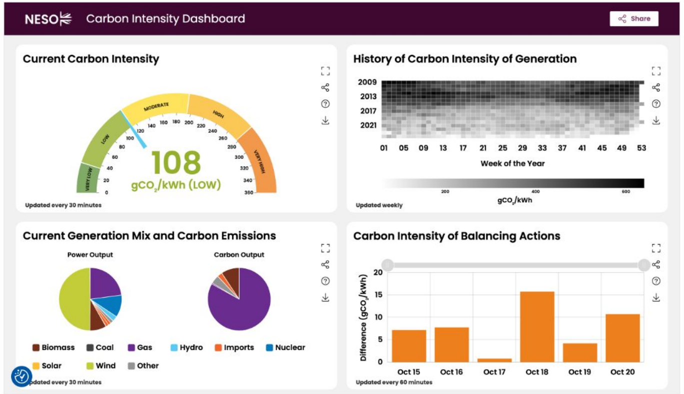

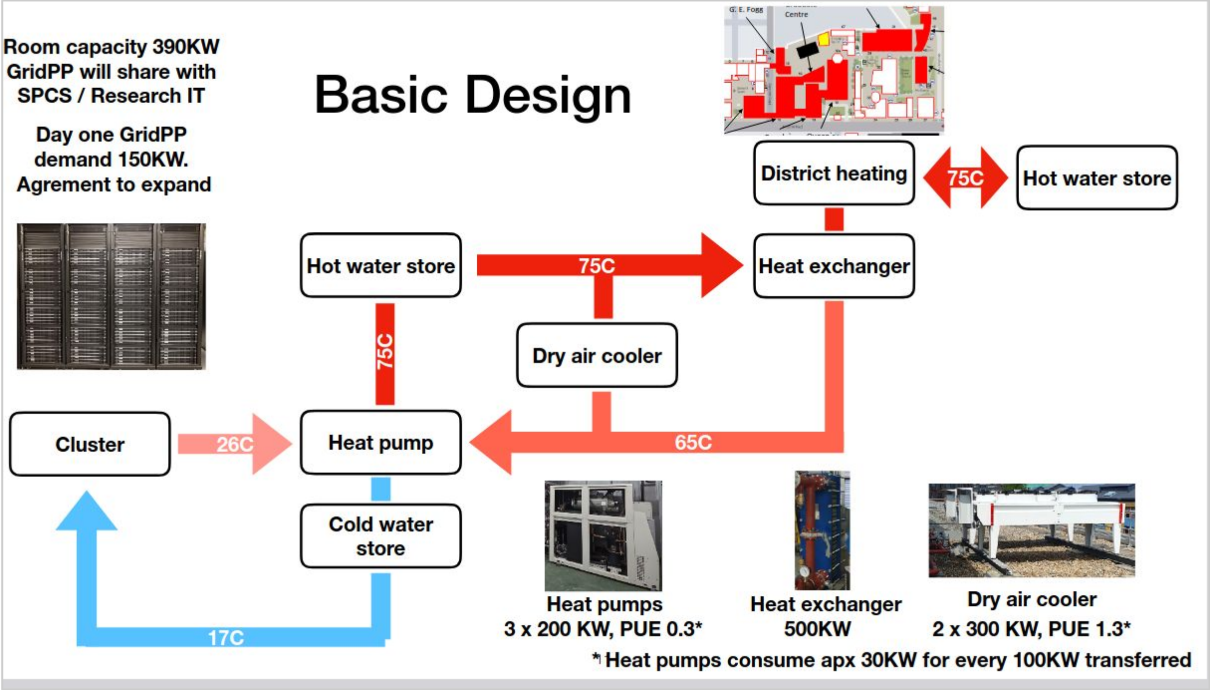

Presentation materials:

A Successful Data Centre Refurbishment Project

Final Design

• Chilled water circuit to cool computing (17C in 23C out).

• Total cooling capacity 390KW provided by three 200KW water source heat pumps (N+1), take 23C water back to 17C. High quality hot water circuit operates at 65C in 75C out.

• Hot water (75C) circuit feeds into either the district heating system via a 500KW heat exchanger or two 300KW dry air coolers. Note dry air coolers operate at 75C! Can be used all year round

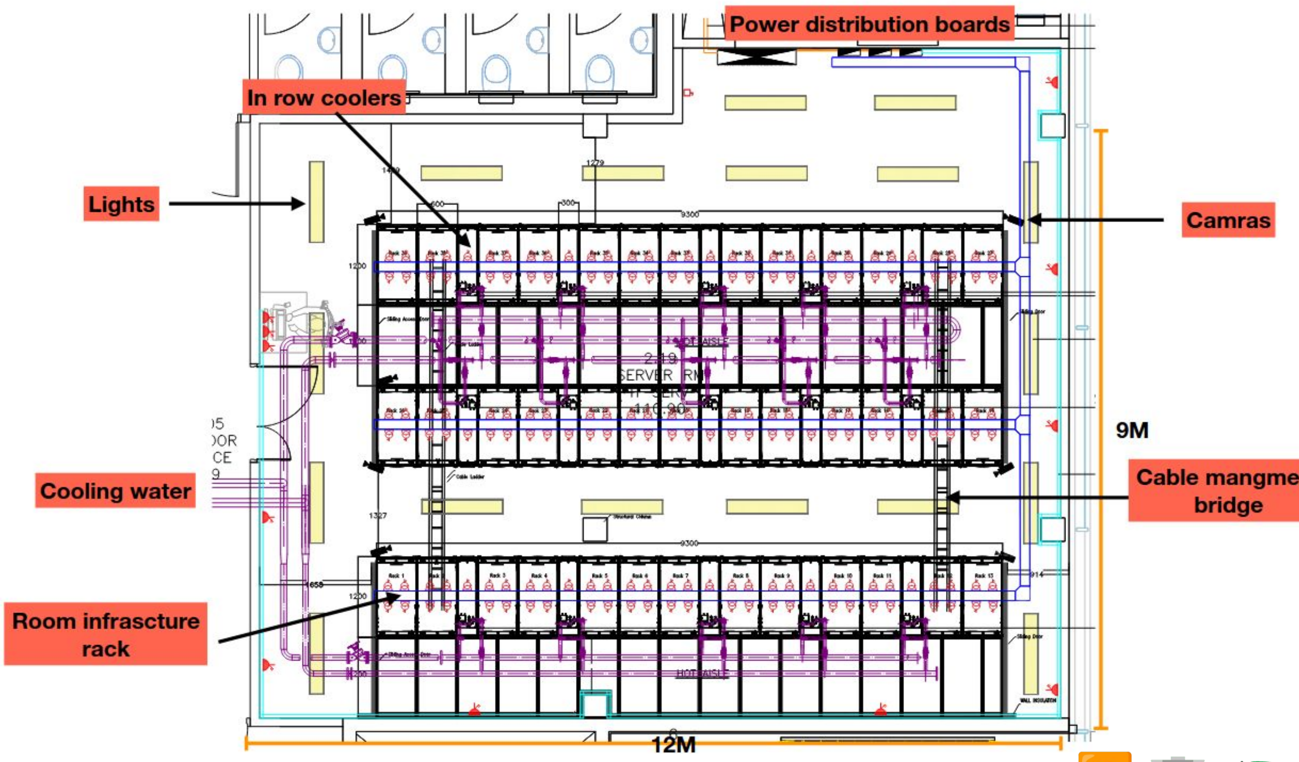

Presentation materials:

A Successful Data Centre Refurbishment Project

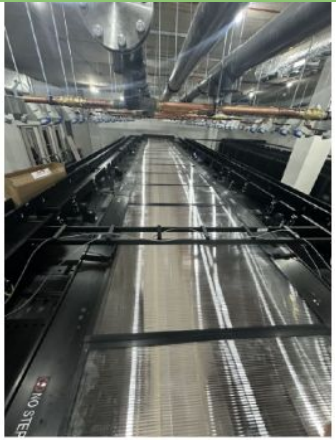

Increase Capacity

- Increase number of usable racks from 15 to 38.
- Increase power supply to DC 150KW to 390KW.
- Bonus - flexible rack power up to 20KW.
- Bonus - Use deep racks (1200mm vs 1070) flexible equipment
- limits.

Improve Efficiency

Move from air conditioning (CRAC-refrigerant) to air handler units (CRAH-water).

Use hot aisle containment (expected to be more efficient than no containment / cold aisle containment).

Temperature deltas chosen to me most energy efficient: cold aisle
26C, Hot aisle 45C*; In water 17C, out water 23C

Presentation materials:

An implementation of cloud-based grid CE and SE for ATLAS and Belle II

Speaker & Primary author: Dr Jonathan Mark Woithe (University of Adelaide (AU))

In this talk describe an implementation of this approach which has been deployed for
Australia's ATLAS and Belle II grid sites.

The sites are built entirely with VM orchestrated by an OpenStack instance. The Storage Element utilises an xrootd-s3 gateway with back-end storage provided through an S3-compatible object store from a commercial provider. The provisioning arrangements required the deployment of some site-specific helper modules to ensure all SE interfacing requirements could be met. OpenStack hosts the xrootd redirector and proxy servers in separate VMs.

The Compute Element comprises VM within the Openstack instance. Jobs are submitted and managed by HTCondor . A Cloud Scheduler instance is used to coordinate the number of active OpenStack VMs and ensure that VMs run only when there are jobs to run.

Benchmarks

	Within cloud	In Australia
davs:// read	108 MB/s	40 MB/s
davs:// write	123 MB/s	74 MB/s
Checksum calc	3.2 s	3.4 s
Checksum fetch	0.72 s	0.98 s
s3 read	213 MB/s	n/a
s3 write	165 MB/s	n/a
root:// read	6.6 MB/s	5.9 MB/s
root:// write	132 MB/s	70 MB/s

Read/write tests used gfal-copy, checksum tests used
Results are the average of 5 tests, each using a 1 GB

Presentation materials:

An implementation of cloud-based grid CE and SE for ATLAS and Belle II

Speaker & Primary author: Dr Jonathan Mark Woithe (University of Adelaide (AU))

A VM in MRC OpenStack hosts HTCondor and HTCondor-CE instances

- Host is running AlmaLinux9
- Token authentication is supported

▸ Jobs submitted to HTcondor-CE on the HTCondor host

▸ After authorisation, jobs passed onto HTCondor by HTCondor-CE on same host

▸ CSv2 processes proceed as for Belle

- Executed on xrootd proxy servers
- Uses bash shell script to interface with xrootd

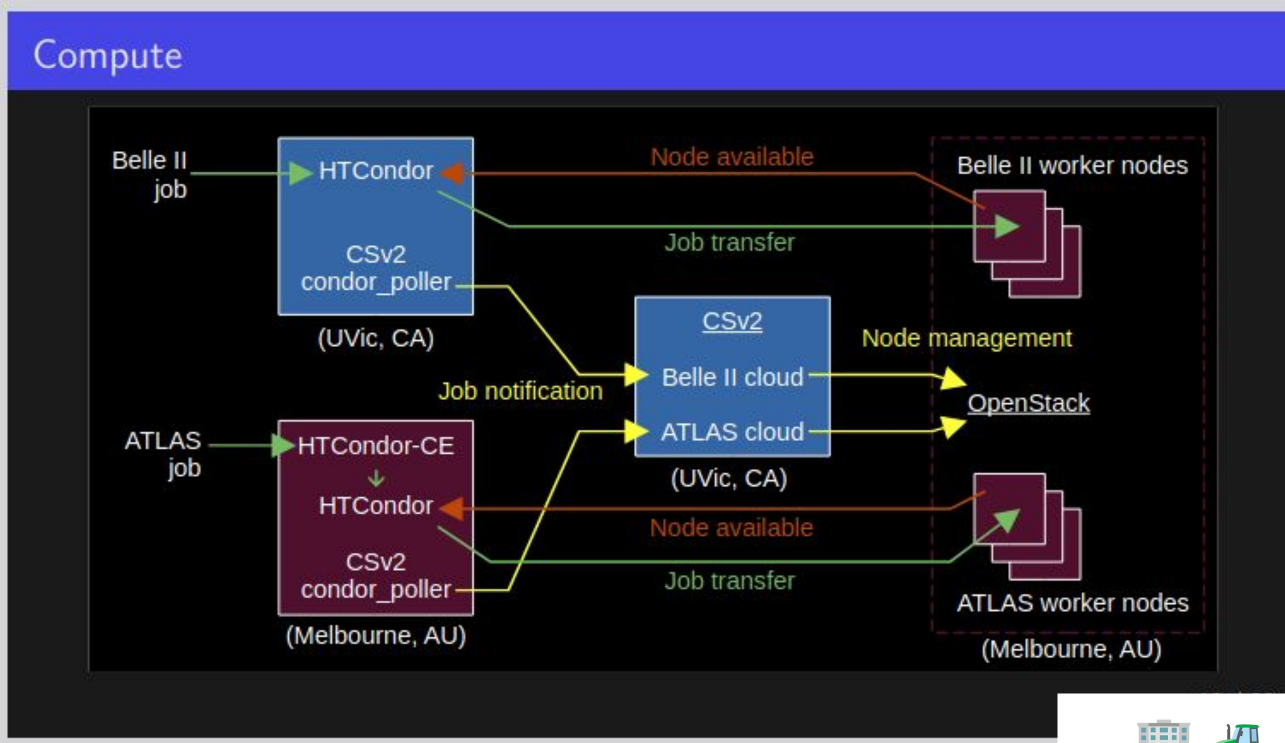

root:// transfers: xrdcp streams content from source, s3cmd sends content into S3 object

Presentation materials:

Provisioning of Grid computing resources in the Norwegian Research and Education Cloud

Speaker & Primary author: Matthias Richter (University of Bergen (NO))

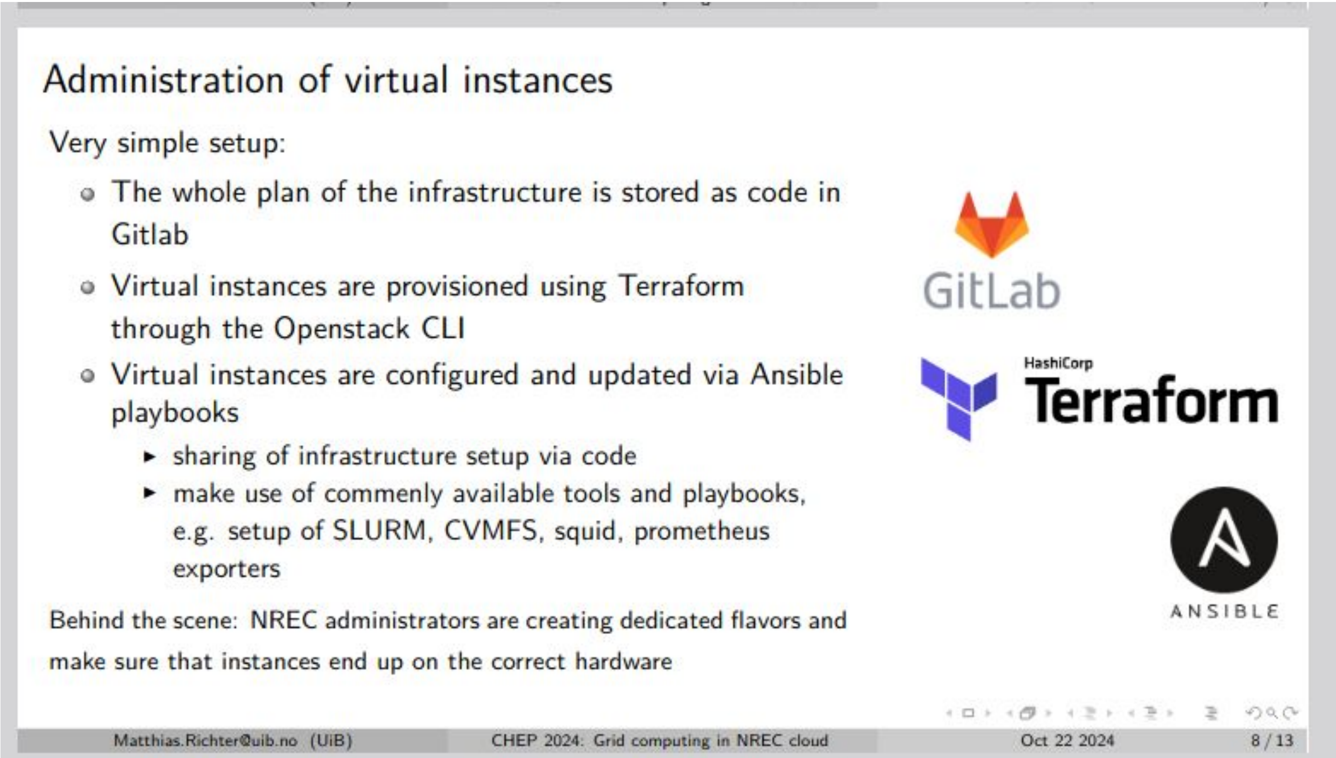

Presentation materials:

Provisioning of Grid computing resources in the Norwegian Research and Education Cloud

Computing resources

- Computing resources are operated in both BGO and OSL regions
  - ▷ OSL 68 virtual instances providing 2170 cores
  - ▷ BGO 325 virtual instances providing 5200 cores
- Both sites provide SLURM clusters to provide compute nodes to Grid middlewar
- ALICE setup in Bergen is running JAliEn which is managing whole nodes
- ATLAS setup in Oslo is running ARC which also handles data staging more flexible

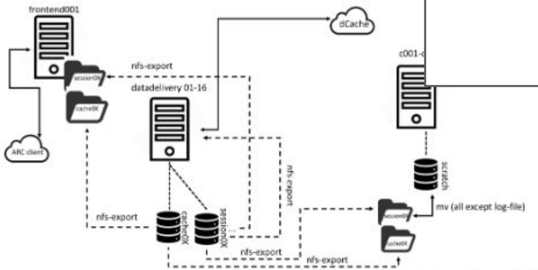

dCache Disk

- The Nordic Tier 1 is operating a distributed dCache cluster
- Only dCache disk storage in Bergen is integrated into NREC
- ATLAS storage is outside NREC, more network traffic in and out NREC
- installed storage servers are managed as CEPH cluster
- NREC provides CEPH storage as block storage volumes
- Dedicated disk pool nodes for integration to dCache mount the CEPH volumes

dCache Tape storage
- Integrated into NREC in BGO region for the ALICE tape storage
- Dedicated 10Gb/s link, write at 800GB/s with two parallel drives
4 drives allow simultaneous mount
4 PB, upgraded by additional 6 PB by end of 2024
Specific hypervisor servers to host dCache pool instances with 10 TB SSD cache
2 write pools, one active, one failover
hot spare slot for new instance
Management is interplay of site admin, dCache admin and IBM SpectrumProtect
operator

💾 🏢

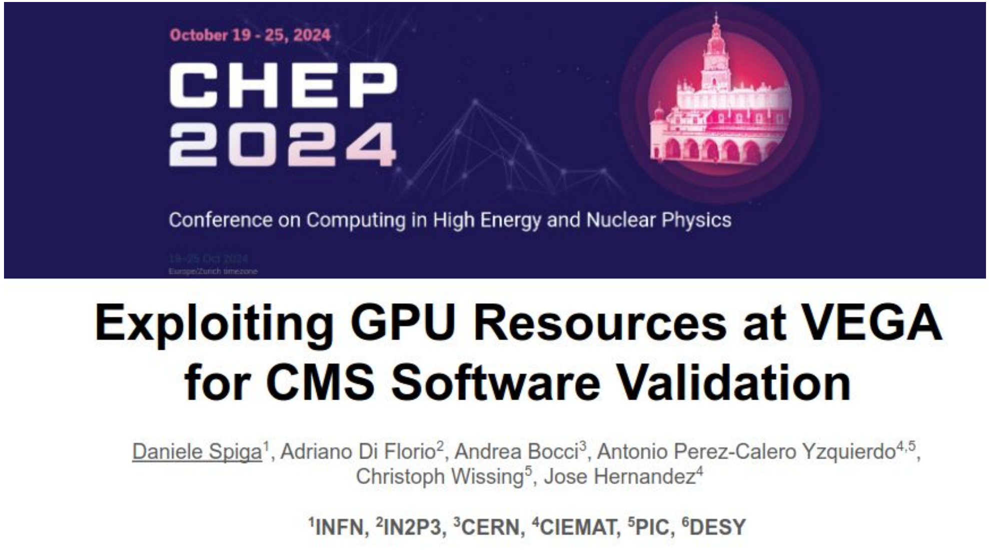

This Talk

HPC resources integration at CMS: brief intro
CMS Grant at VEGA EuroHPC in Slovenia
- Motivation
- Strategy
- Results
Summary and Lessons

The INFN has recently developed a national cloud platform to enhance access to distributed computing and storage resources for scientific researchers. A critical aspect of this initiative is the INFN Cloud Dashboard.
The platform is based on INDIGO-PaaS middleware, which integrates a TOSCA-based orchestration system.

Integrating the Perlmutter HPC system in the ALICE Grid

The SuperFacility API

- HTTP based API to manage jobs
- Uses JWT to authenticate
- Initially we were promised an easy security review process for getting an auth token
    - Policy is to receive 30 days job submission privileges
- SFAPI receives SLURM job specifications and sends them to the HPC system

https://docs.nersc.gov/services/sfapi/

Let's use Kubernetes

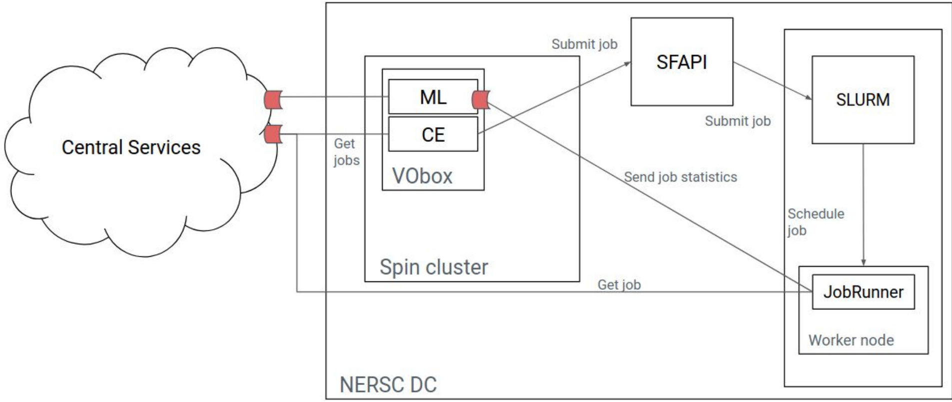

Latest developments of the PUNCH4NFDI compute and storage infrastructures

Compute4PUNCH key ingredients

- Heterogenous compute resources with different schedulers, architectures, operating systems
- Aggregate resources in a single **Overlay Batch System** based on **HTCondor**
- Provision resources to the OBS using **COBalD/TARDIS meta-scheduler**
  → Single federated pool with dynamic extension on a user-demand basis
- **TARDIS** integrates available resources into OBS
- **COBalD** does balancing, job to resource matching, ensures effective usage

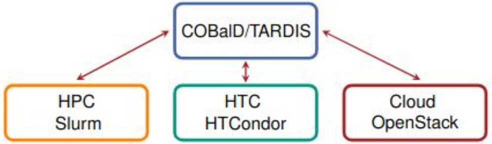

Job submission

```
condor_submit analysis.jdl
Submitting job(s)
Hello benoit_roland!  You are going to submit your HTCondor jobs.
No credential has been found!
Please visit the following url in order to generate your credential:
https://mytoken.data.kit.edu/c/XWSdsgNC
Starting polling and waiting for your approval...
Your credential has been successfully created!
Its remaining life time is 6 d 23 h 58 m 46 s.
1 job(s) submitted to cluster 1142
```

→ **Visit specified URL and approve Mytoken**

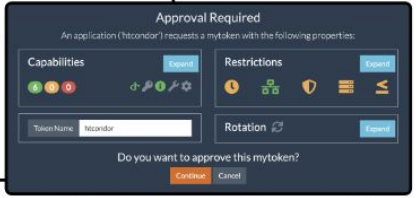

Presentation materials:

Moving a data center keeping availability at the top

Speaker - D.Lattanzio
Authors - A.Pascolini, A.Chierici, D.Michelotto, D.Cesini, G.Sergi

- Describing how we were able to move all the resources without interrupting the service provided to users
- **Redundancy in first place**
- What we moved
- How we moved it
- Lesson learned

Switch on of the first WN rack in production at Technopole 19/03/24

Presentation materials:

Prévessin Data Centre Powers Up

Speaker - Max Dupuis (CERN)
Authors - Joel Murray Davies (CERN), Max Dupuis (CERN)

- CERN's state-of-the-art Prévessin Data Centre (PDC) is now operational, complementing CERN's Meyrin Data Centre Tier-0 facility to provide additional and sustainable computing power to meet the needs of High-Luminosity LHC in 2029
- In 2019, it was decided to tender the design and construction of a new, modern, energy-efficient (PUE of ≤ 1.15) Data Centre with a total of 12 MW IT capacity spread across six IT rooms

PDC: Design

Highly efficient data centre with a Power Usage Effectiveness (PUE) [2] ≤ 1,15

Construction in 1 phase

Installation will be in 3 phases over time for a total of 12MW

- **1st Phase**: 4MW - 2nd floor
- **2nd Phase**: +4MW - 1st floor
- **3rd Phase**: +4MW - ground floor

4x 2MW Jaegi dry coolers in 1st Phase (2x2MW per additional phase) - N+1 redundancy plus 1 for technical areas

Adiabatic cooling from 20 degrees outside temperature with water recycling to improve Water Usage Effectiveness(WUE)[3]

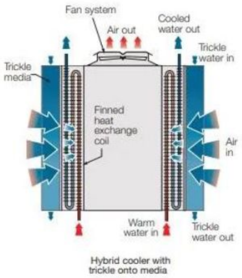

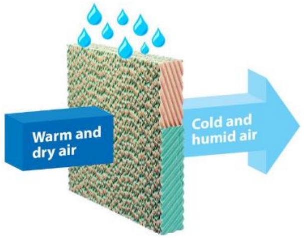

Computational Fluid Dynamics (CFD) has been used to determine the optimal IT room layout and to evaluate various failure scenarios.

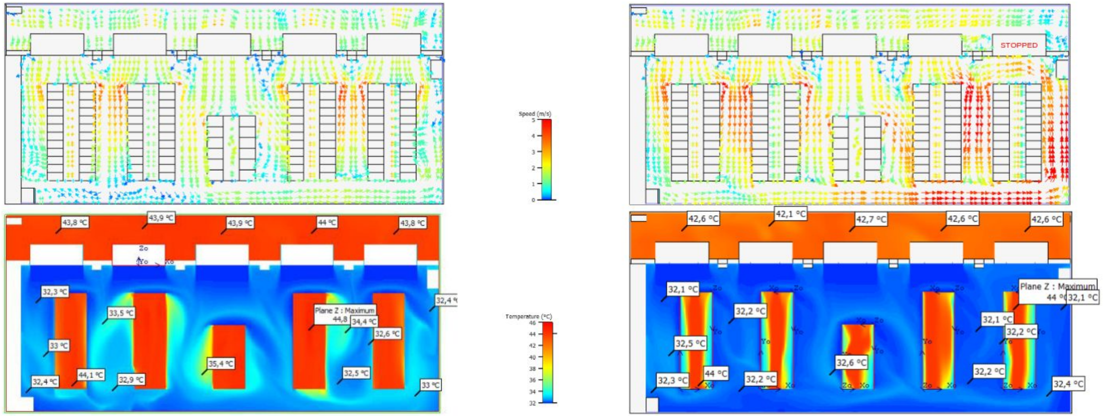

PDC: 1st year of operation

Since the beginning of 2024, the two rooms of the 1st Phase have been operational. In 6 months, we successfully commissioned the PDC at 50% capacity

NEXT: In preparation to Run4, we will need to initiate the installation of the 2nd Phase of the PDC and begin procuring additional servers

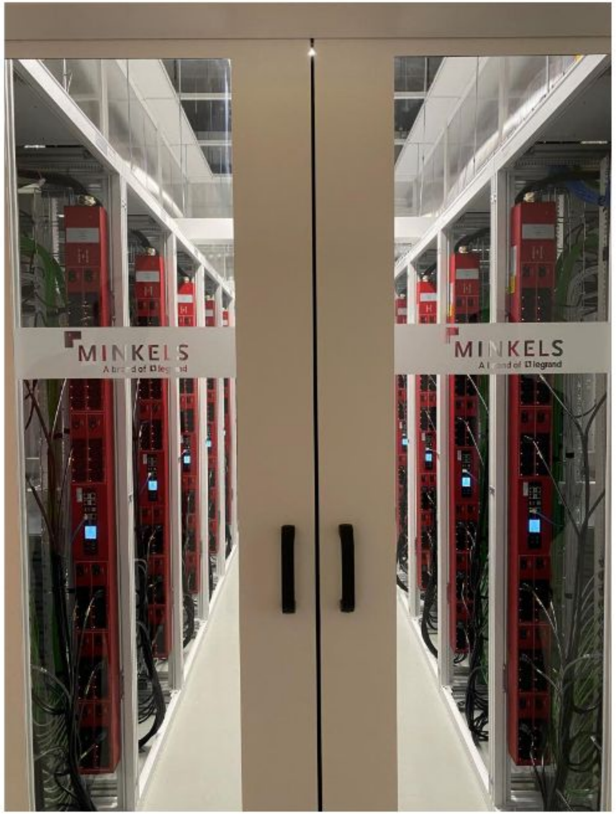

Presentation materials:

Using the ATLAS experiment software on heterogeneous resources

Speaker & Primary author - Johannes Elmsheuser (Brookhaven National Laboratory (US))

- With the large dataset expected from 2029 onwards by the HL-LHC at CERN, the ATLAS experiment is reaching the limits of the current data processing model in terms of traditional CPU resources based on x86_64 architectures and an extensive program for software upgrades towards the HL-LHC has been set up.
- The ARM CPU architecture is becoming a competitive and energy efficient alternative. Accelerators like GPUs are available in any recent HPC
- In the past years ATLAS has successfully ported its full data processing and simulation software framework Athena to ARM and has invested significant effort in porting parts of the reconstruction and simulation algorithms to GPUs.

Using the ATLAS experiment software on heterogeneous resources

ATLAS Grid setup: PanDA and Rucio

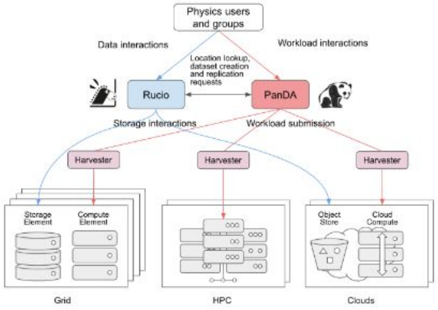

- More details in arxiv:2403.15873
- Dedicated ARM and GPU PanDA queues configured in Computing Resource Information Catalogue (CRIC)
- Full ARM/aarch64 grid setup available with OS container, middleware, Kubernetes etc.
- For NVIDIA GPUs need matching CUDA linux kernel module version and redistributable CUDA libraries

- Pilot Job configuration:
  - Dedicated ARM queue on the CE or using ”WantARM=True” in PanDA pilot job jdl
  - Dedicated GPU queue on the CE or using ”+RequireGPUs = True” and ”+RequestGPUs = 1” in PanDA pilot job jdl
- ATLAS user job submission:
  - Use following options for PanDA job submission tools (prun/pathena): *–architecture* ”&nvidia-*” or *–architecture* ”@el9#aarch64”

Usage of ARM in ATLAS

Last CHEP23 presented ”The ATLAS experiment software on ARM” (link) - much has happened since then

- 7 nightly builds with gcc13 for development and production branch and different projects of ATLAS software stack Athena, which are built on 4 build machines provided by CERN IT (Ampere Altra/Neoverse-N1) - nightly and stable releases automatically installed on CVMFS
- Running MC simulation and reconstruction on 5 PanDA queues with up to 15k concurrent job slots
- Configured PanDA queue as extension of US Tier2 in UT Arlington in *Google Cloud* with up to 9.5k job slots

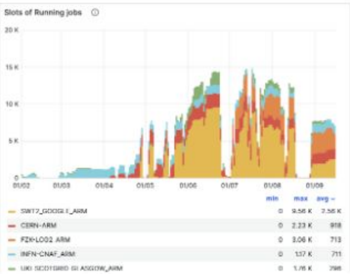

ATLAS is the first WLCG experiment which will accept ARM resources as pledge in 2025/26

Using the ATLAS experiment software on heterogeneous resources

Using GPUs in ATLAS - introduction and caveats

- ATLAS is not using GPUs in production for Run3
- Major R&D effort on-going to port parts of simulation, reconstruction and high level trigger code to use GPUs for HL-LHC (see several talks in other sessions at this conference)
- Since March 2024, Athena main branch uses CUDA 12.4 (and later 12.4.1) - this version of CUDA supports gcc13 (current production compiler version in ATLAS)
- Parts of CUDA SDK are redistributable (see file list at link) and usage works fine also outside of CERN via CVMFS and PanDA
- N.B. ATLAS HLT group successfully ran fully automated offline reprocessing of Calorimeter topo clustering algorithm on GPUs via PanDA
- **But:** every PanDA GPU queue/site has to have at least the kernel driver version 550.54.15 (or newer) from CUDA 12.4.1 installed
  → Tedious process in reaching out to sites and asking for CUDA kernel driver version update
  → Same process repeats potentially when moving to a new major CUDA version
  → A (automated) procedure should be discussed within WLCG when GPUs are more commonly used on the Grid

7/11

Using Athena on NVIDIA Grace Hopper

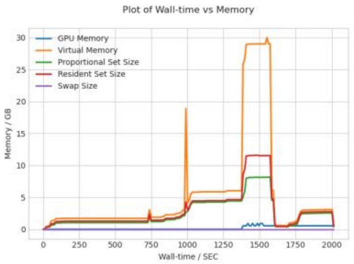

prmon plot of memory usage over time for ATLAS HLT reconstruction workflow on Grace Hopper (72 core Arm Neoverse v2 + GH200 GPU)

- Athena HLT reconstruction code ported to GPUs benchmarked on NIVIDA Grace Hopper testbed provided through LBNL
- GPU workflows run out-of-the box on this ARM CPU+GPU testbed (but slower due to missing frontier/squid in this testbed)
- Reliable GPU benchmark needed in future HepScore version - so far only CPUs:

Name	nCPU	HepScore23	HepScore23 per nCPU
HepScore23 reference	64	1018	15.9
Grace Hopper	72	2319	32.2
Apple M2 Air	8	141.4	17.7
Ampere Neoverse-N1	20	349.4	17.5
Intel Xeon E5-2683 v4	16	258.5	16.2

Presentation materials:

Heterogeneous Computing and Power Efficiency in HEP

Speaker - Emanuele Simili,
Authors - Albert Gyorgy Borbely, David Britton, Gordon Stewart, Samuel Cadellin Skipsey

- In 2021 started investigating alternative architectures for Grid computing, starting with ARM chips

Methodology:

- As benchmark, we rely on the HEP-Score & the HEP-Benchmarking Suite:
  - **HEP-Suite**: https://gitlab.cern.ch/hep-benchmarks/hep-benchmark-suite
  - **HEP-Score**: https://gitlab.cern.ch/hep-benchmarks/hep-score

- While the benchmark executes, a script collects and exports CPU, RAM, Frequency and Power usage (via IPMI tools) into a CSV file. Performed a few tests to validate **IPMI** power readings against a metered **PDU**

Presentation materials:

Heterogeneous Computing and Power Efficiency in HEP

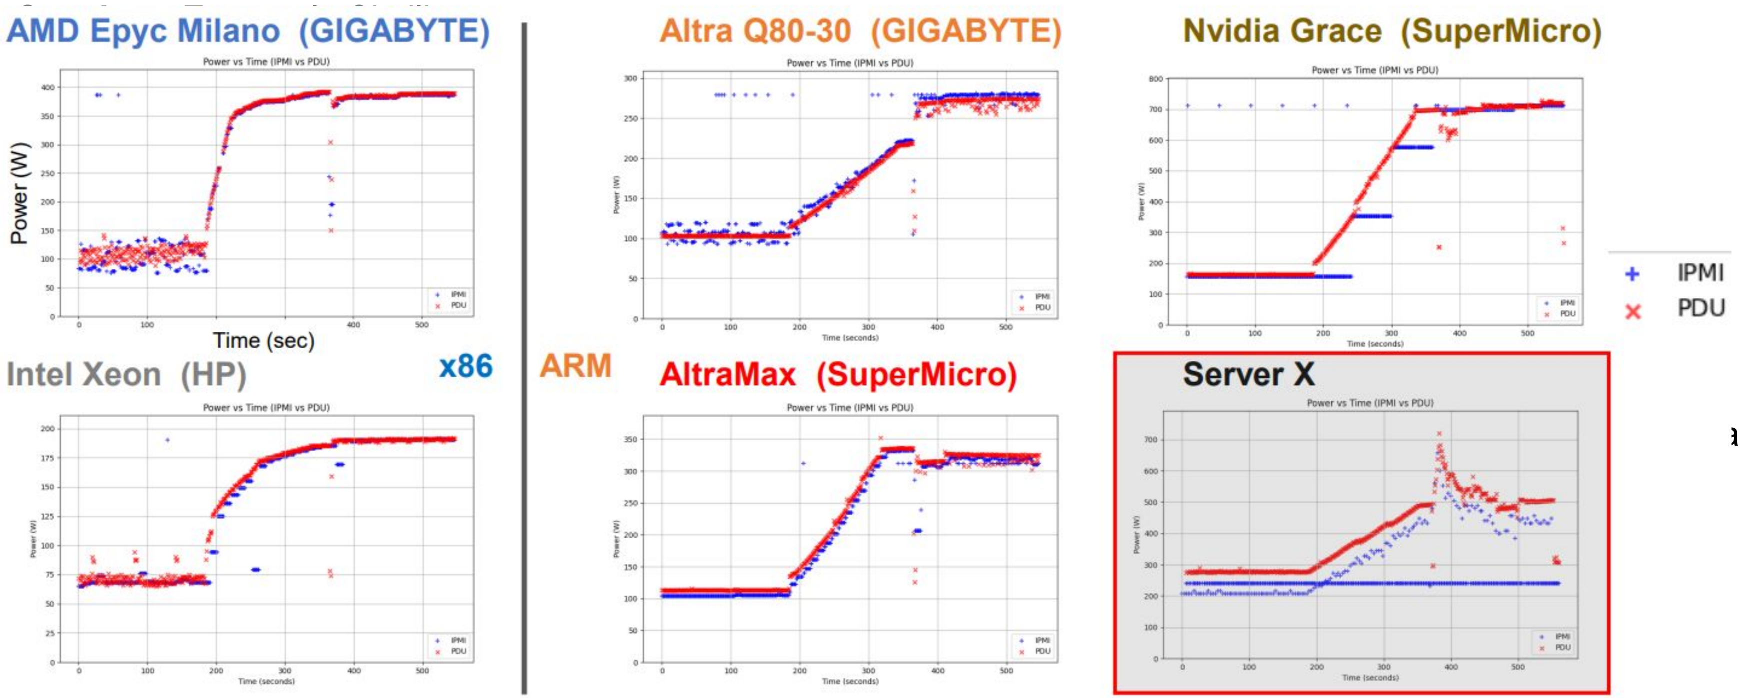

Presentation materials:

Heterogeneous Computing and Power Efficiency in HEP

- Since **HEP-Score** cannot yet run on every hardware (e.g., **RISC-V**, **GPU**), we may need other benchmarks to assess the performance of new architectures. We have tried a few other standard HEP benchmarks: **ROOT bench**, **Geant4** with **CMS** geometry, and **DB12** (single core and whole node).

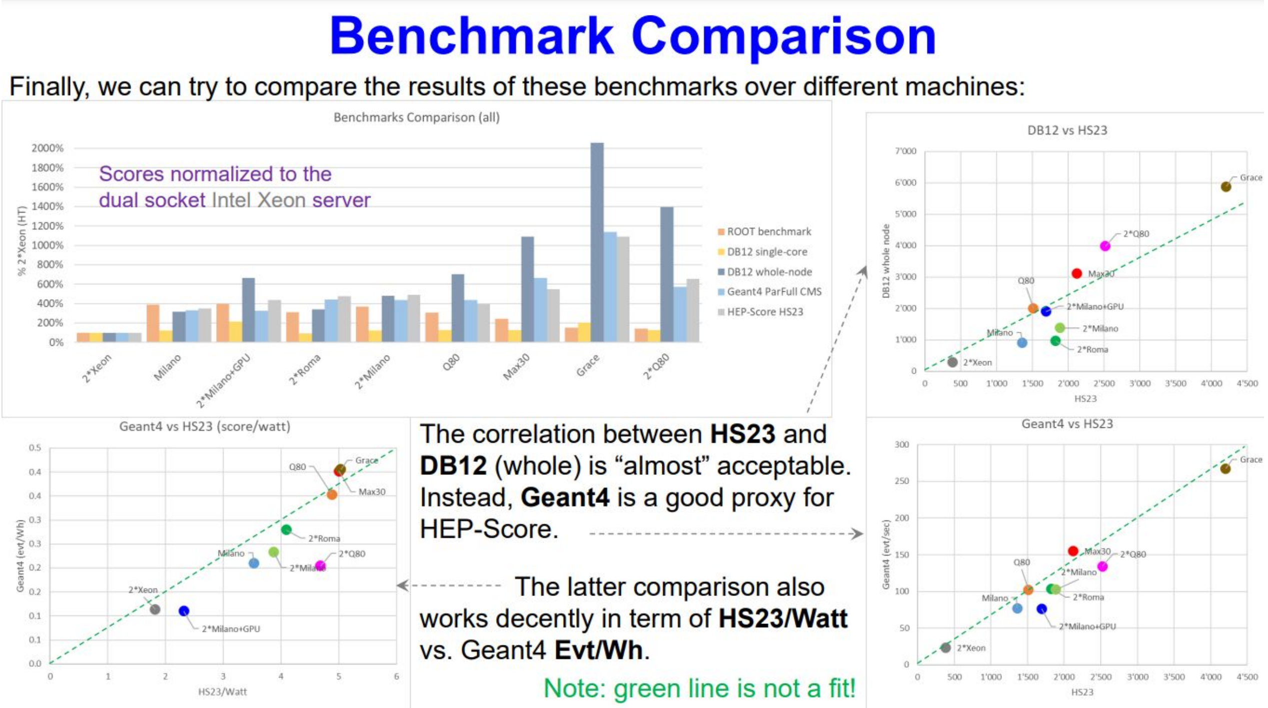

Heterogeneous Compute Cluster

We started providing **ARM** resources at the our WLCG Tier2 cluster by creating a separate queue for ARM (former **ce-test**). After upgrading, we joined both queues within our standard ARC-CE endpoints.

This is a simplified view of our heterogeneous computing cluster (we still keep **ce-test** alive for testing):

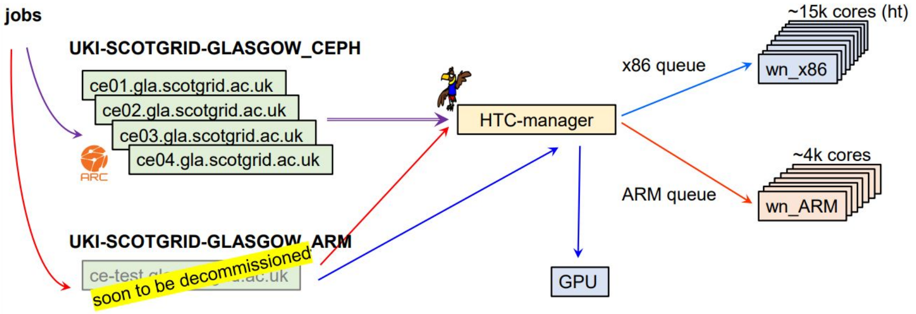

The **condor_requirements** setting in the ARC-CE configuration modifies the **ClassAd** for the jobs that **ARC** submits to **HTCondor** by inserting an architecture request (x86, ARM, GPU) …

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 **VO**s to start sending ARM jobs our way … we have over 4k ARM cores !

Presentation materials:

Taking on RISC for Energy-Efficient Computing in HEP

Speaker - Emanuele Simili
Authors - Shahzad Muzaffar, Tommaso Boccali, Albert Gyorgy Borbely, David Britton, Gordon Stewart, Samuel Cadellin Skipsey

- In this work, we have taken a first look at the **RISC-V** architecture for HEP workloads
- We introduce the **Pioneer Milk-V**, a 64-core RISC-V machine running Fedora Linux, as our new testbed, available at **ScotGrid Glasgow** (UK) and **INFN Bologna** (Italy)
- HEP & WLCG software: **ROOT, Geant4, CVMFS, XRootD**
- Selected benchmarks: **ROOT bench, DB12, Geant4 (2x)**

But in practice, what do we have today?

- Up to 6 months ago, “Raspberry PI” like devices: **Visionfive 2**
  - 4 GB RAM, 4 cores, Fedora 37, Ubuntu 20.04,
  - Useful for initial tests, bootstrap of some HEP codes (Geant4, ROOT, some benchmarks),
  - 12 hours to (fail to) compile gcc.

Visionfive 2
~100 € on Amazon

- More recently, fully integrated desktop PC: **Milk-V Pioneer***
  - 128 GB RAM, 64 cores, Fedora 38, Ubuntu 24.04, Debian,
  - Out-of-the box compatibility with HEP codes (Geant4, ROOT, more benchmarks),
  - Node Performances almost comparable to a standard x86 mid range desktop.

Milk-V Pioneer
~2.5k € from China

- Upcoming: **MIPS P8700**, up to 512 cores !
  - detailed specs yet to be be defined.

<u>Cores ramp per chip</u>: **4** → **64** → **512** in just 1 year!

* Now available at CERN, Glasgow, Pisa, Bologna, and CNAF (as far as we know).

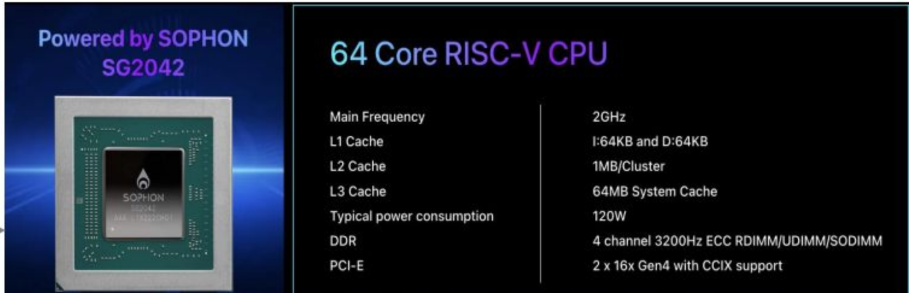

Benchmarks Comparison

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.

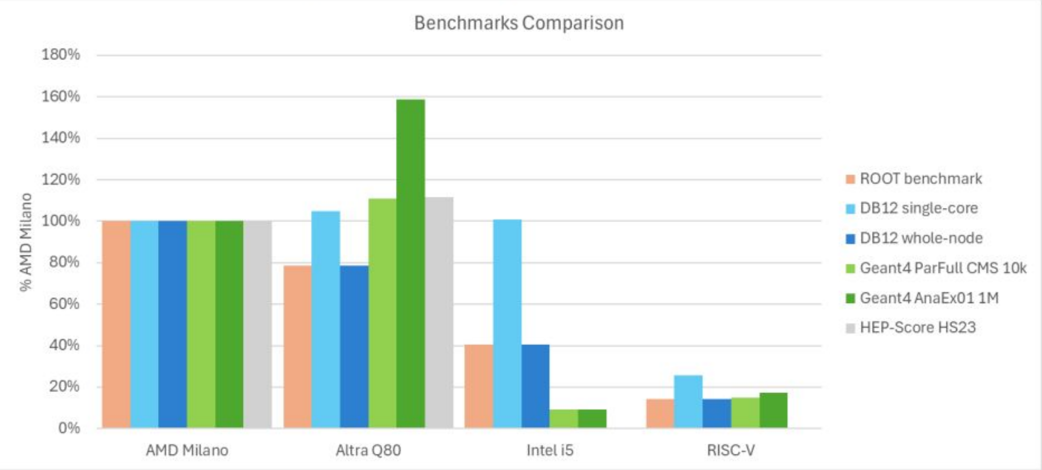

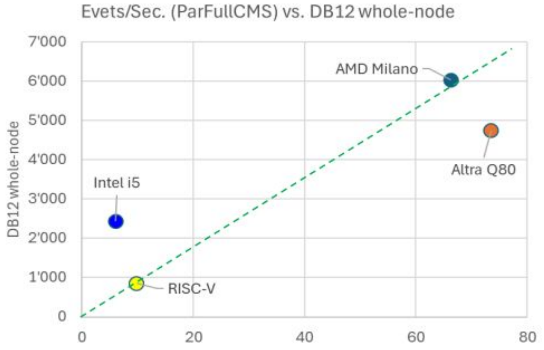

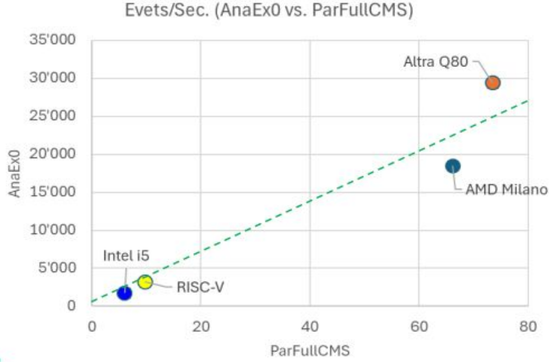

As mentioned, the first two benchmark (ROOT & DB12 single) are single-threaded, the others are multi-threaded (DB12 whole, both Geant4 simulations, and the HEP-Score).

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

Note: the green line is not an actual fit.

Power Usage

To get a better feeling for the power usage of the **Milk-V Pioneer**, we have measured the idle baseline while running **Geant4** jobs with increasing number of threads.

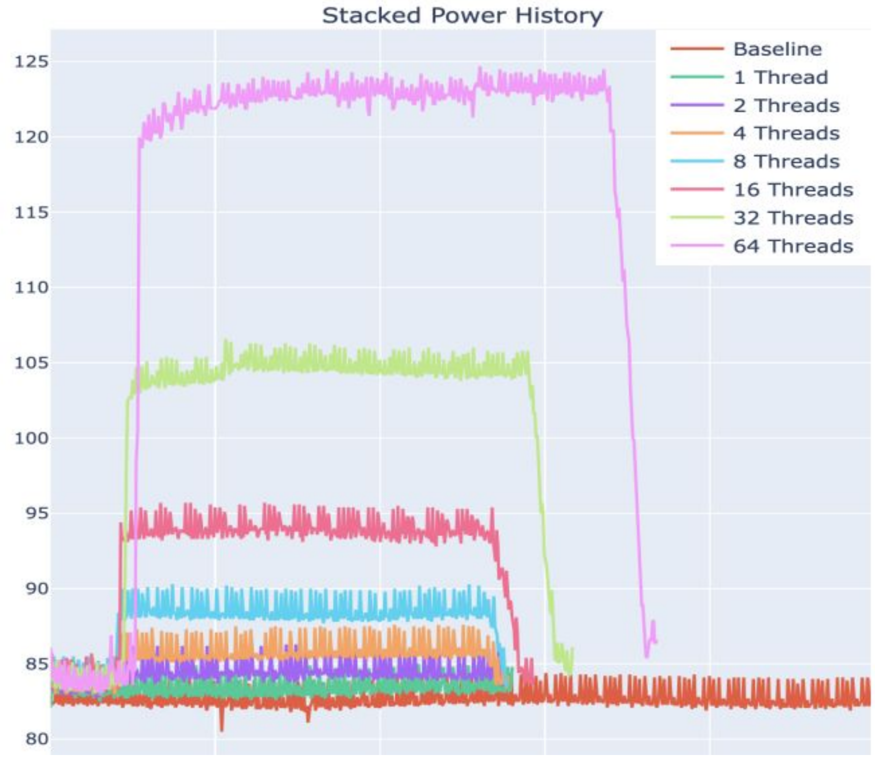

We notice that the **Milk-V Pioneer** has a large power usage in idle mode, comparable to the idle power of a single socket servers, but definitely too high if compared to the mid-range **x86** desktop!

idle power

AMD Milan	~	110W
Altra Q80	~	100W
Intel i5	~	13W
RISC-V	~	80W

Presentation materials:

Heterogeneous computing at INFN-T1

Speaker - D.Lattanzio
Authors - A.Pascolini, A.Chierici, D.Michelotto, G.Sergi

- ARM resources
- RISC-V resources
  - **An interesting architecture for the future**
- LHC experiments point of view
- **Power consumption**

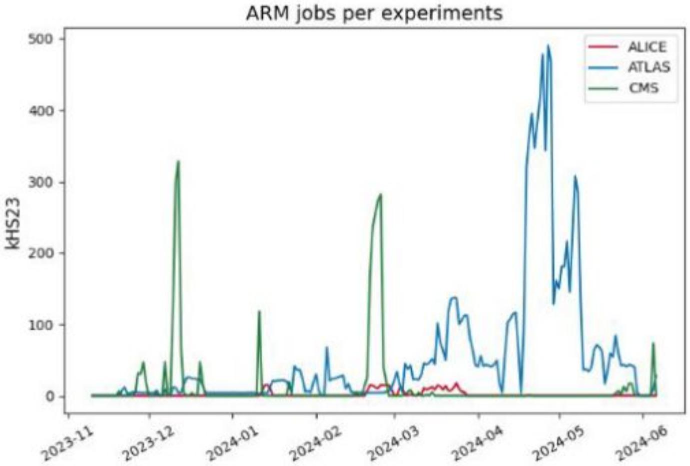

Track 9: Analysis facilities and interactive computing

- Conveners:
  - Enric Tejedor Saavedra (CERN)
  - Marta Czurylo (Heidelberg University)
  - Nick Smith (Fermilab)
  - Nicole Skidmore (University of Bristol)
- 18 talks, 4 posters

Analysis Facilities: the perspective from *Data Centers*

- CERN Analysis Facility Pilot
  - instance currently under study for **scale out of interactive analyses**
  - [contribution link][summary slide]
- CERN Virtual Research Environment
  - **PoC of a middleware interface** allowing users to access services and computing resources
  - [contribution link][summary slide]
- Interdisciplinary Analysis Facility at DESY
  - computing system driven by **generic concepts** and **interdisciplinary applications** (beyond HEP)
  - [contribution link][summary slide]
- The HEPS scientific computing system
  - Chinese perspective of data centers and analysis facilities, aiming at **multidisciplinary applications**
  - [contribution link][summary slide]

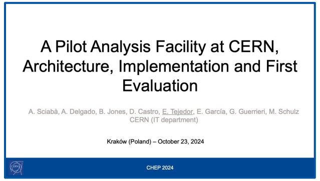

A Pilot Analysis Facility at CERN, Architecture, Implementation and First Evaluation

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

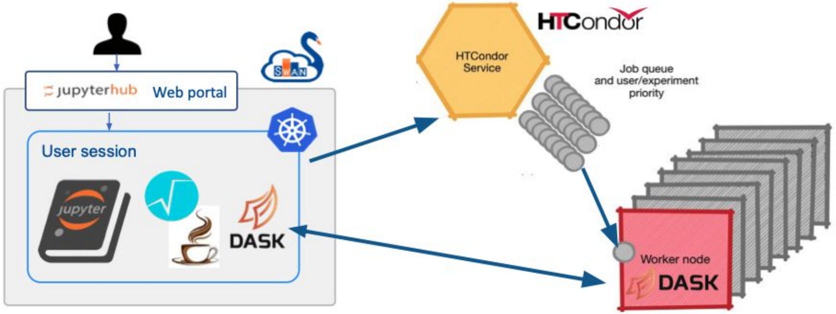

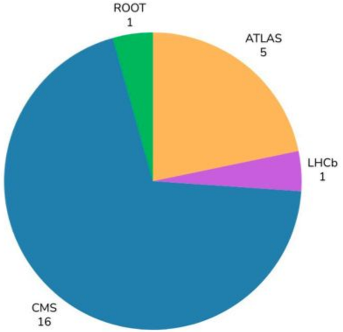

Testers by experiment

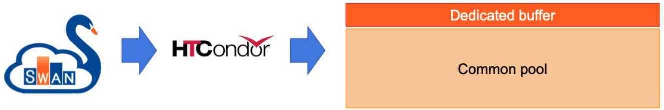

Dedicated buffer of ~2k cores for quick allocation of resources

Feedback: Software provisioning

- **Pre-configured stacks** are appreciated (experiment stacks, LCG releases), but more **flexibility** is also requested
  - Custom software environments (e.g. conda)
  - Custom container images
- As an alternative, the pilot allows to install **additional packages on CERNBox/EOS**, to be used on top of an LCG release
  - Some errors accessing the mount on the batch side have been reported
  - Approach seen as a stopgap solution

Feedback: Data access

- Both access to data on EOS and to **external datasets** (i.e. not at CERN) has been reported
  - External access usually done via remote URLs (e.g. XRootD), sometimes via Rucio copy
  - Added support for X.509 proxy generation (tokens could follow)
- Suggests that setting up a **cache** could be beneficial
  - Work in progress with storage group at CERN IT to have an XCache PoC

Feedback: User Interface

- The pilot offers the JupyterLab user interface, meaning:
  - Notebooks, terminal
  - Dask extension to **graphically reserve resources** (i.e. create a Dask cluster) and dashboard to **monitor** the progress of the distributed application
- Users mostly satisfied with the interface, but reported some issues:
  - Creating / destroying a cluster can be slow (minutes) if many (a few 100s) Dask workers are requested, since their jobs are submitted / removed sequentially
  - The Dask monitoring dashboard can be sometimes laggy and freezing
- **Programmatic reservation of resources** and execution (from the terminal) is appreciated

source: https://indico.cern.ch/event/1338689/contributions/6010680

Data discovery, analysis and reproducibility in Virtual Research Environments

Enrique García[1], Giovanni Guerrieri[1], Rubén Pérez[1], Michael Zengel[1], Georgy Skorobogatov[2], Andrés Tanasijczuk[3], Hugo Gonzalez[1] and Xavier Espinal[1]
[1] CERN, [2] ICCUB Barcelona, [3] UCLouvain

CHEP 2024 | Track 9 | 24 October

Data discovery, analysis and reproducibility in Virtual Research Environments

CERN Virtual Research Environment

VRE@CERN

- ESCAPE outcome
  - The place where most of the data-related R&D activities happened.
  - PoC platform provisioning different solutions for users:
    - Built on top of a common scientific tool: Jupyter
    - Software environment & repositories
    - Connection with resources: EOS, CVMFS
    - Ease interacting with underlying services

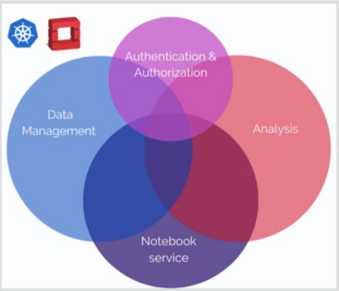

- To date CERN VRE acts as a **middleware interface** allowing scientific users to access different services and computing resources.

- CERN VRE presents an Analysis Facility concept
  - Allows performing most of the stages of a complete scientific analysis
  - All from the localhost of a Jupyter session - compatible with a remote / cloud-based architecture
    - Accessing and discovering data via the **Rucio** Plugin
    - (re)Triggering a workflow on a **Reana** instance
    - Publishing results on **Zenodo**
- Fully modular
  - Can help establishing the basic tools for certain communities

RUCIO

reana

zenodo

source: https://indico.cern.ch/event/1338689/contributions/6010696

Evolution and Broadening of the National Analysis Facility at DESY

CHEP 2024

Christoph Beyer, Stefan Dietrich, Martin Flemming, Sandro Grizzo, Thomas Hartmann, Jürgen Hannappel, Yves Kemp, Joja Meyn, Johannes Reppin, Krunoslav Sever, Christian Sperl, Alexander Trautsch, Christian Voß

https://naf.desy.de https://docs.desy.de/naf

Evolution and Broadening of the National Analysis Facility at DESY

Lessons for the NAF

- Avoiding being technology/implementation driven
- Aiming for generic concept driven approach wrt. user needs

NAF is a *whitelabel* Analysis Facility

- Users today from a broad spectrum of communities end experiences
  - User support crucial
    - Any technology can only be auxiliary
  - Evolved beyond HEP
    - **Interdisciplinarity** central

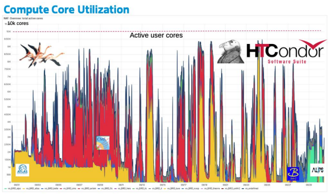

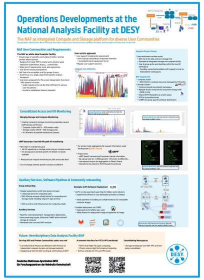

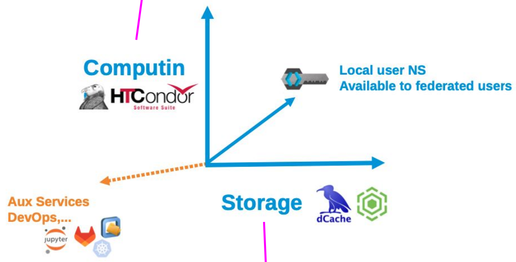

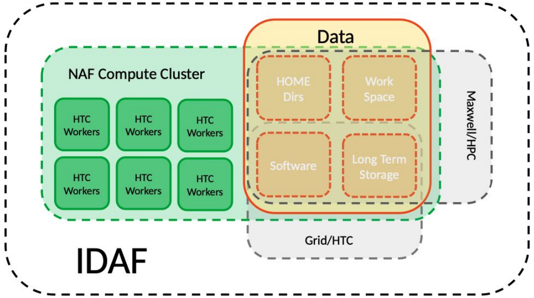

eBPF based per User/Job file path I/O monitoring

- Zoo of storage systems specific solutions prohibitive

eBPF

- Project aims:
  - eBPF based: harnessing all necessary kernel information
  - By cgroup task list → LRMS agnostic: HTCondor, SLURM, K8s,...
  - Job requests as json for aggregation in Elastic Search
  - Intended to be extend to TCP/IP based I/O protocols

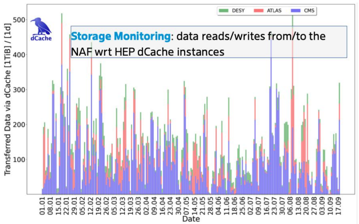

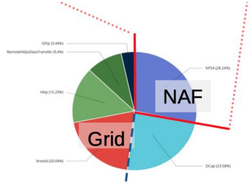

Ratio of protocol uses on HEP dCache instances

Future: Interdisciplinary Data Analysis Facility IDAF

Serving HEP and Photon Communities under one roof

- Currently Particle Physics and Research with Photons on independent compute clusters and storage backends
- Integrating all into the IDAF serving all communities at DESY

A common interface for HTC & HPC workloads

- NAF & Grid Hight Throuput Computing
- Photon science High Performance Computing
- IDAF as umbrella for HTC and HPC needs

Consolidating Namespaces

- Storage namespaces over NAF, HPC and Grid being consolidated

sources: https://indico.cern.ch/event/1338689/contributions/6010713, https://indico.cern.ch/event/1338689/contributions/6010687

HEPS scientific computing system design for interactive data analysis scenarios

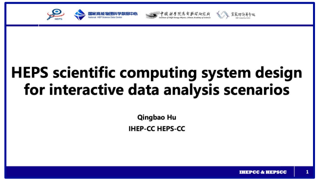

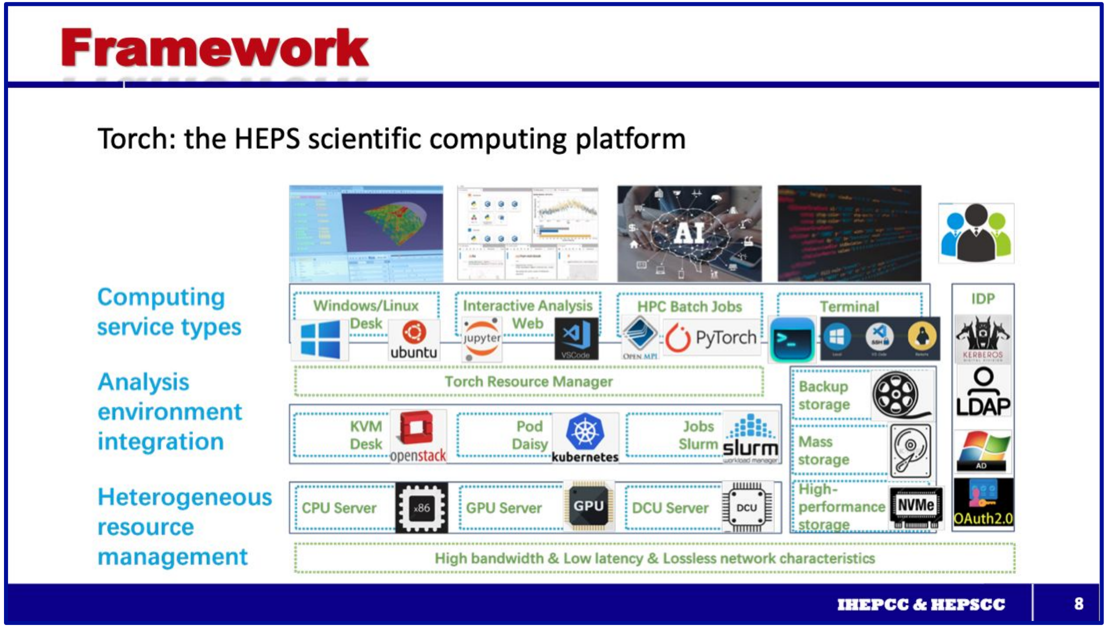

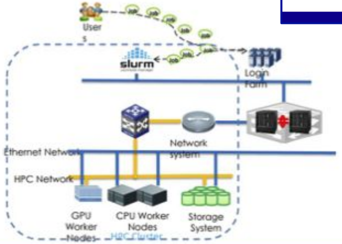

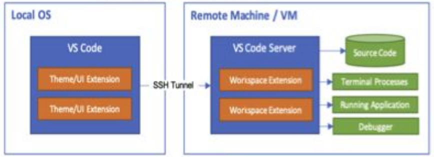

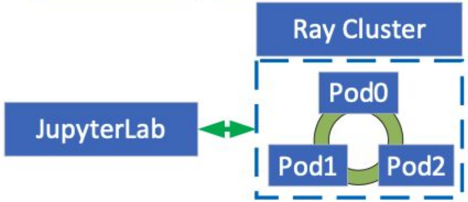

The **HEPS scientific computing system** realizes centralized management and unified scheduling of heterogeneous computing resources, and provides diversified computing service solutions according to different usage patterns.

Analysis Facilities: the perspective from *Experiments*

- Benchmarks for ATLAS and CMS use-cases within the ICSC context
  - interactive analyses via Jupyter notebooks using **ICSC resources** for HEP experiment use-cases
  - [contribution link][summary slide]
- Spanish perspective for CMS and ATLAS on Analysis Facilities
  - CIEMAT AF architecture for CMS and Tier-1 & Tier-2 facilities for ATLAS
  - [CMS contribution link][summary slide] [ATLAS contribution link][summary slide]
- Analysis Facilities federation for US ATLAS
  - prototype under development to **federate** BNL, SLAC, and UChicago facilities
  - [contribution link][summary slide]
- Analysis Grand Challenge with Snakemake and REANA
  - **evolution REANA** to face AGC in combination with Snakemake as workflow manager
  - [contribution link][summary slide]
- Stressing ATLAS and CMS Analysis Facilities with the 200 Gbps challenge
  - ATLAS and CMS experiences facing the **200 Gbps challenge** for Analysis Facilities
  - [CMS contribution link][summary slide] [ATLAS contribution link][summary slide]

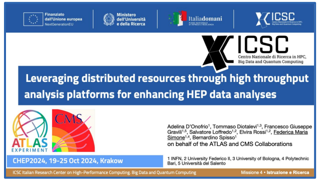

Leveraging distributed resources through high throughput analysis platforms for enhancing HEP data analyses

Need to:

- Optimize the usage of CPU and storage
- Promote the usage of better data formats
- **Develop new analysis paradigms!**
- New software based on declarative programming and interactive workflows
- Distribute on geographically separated resources

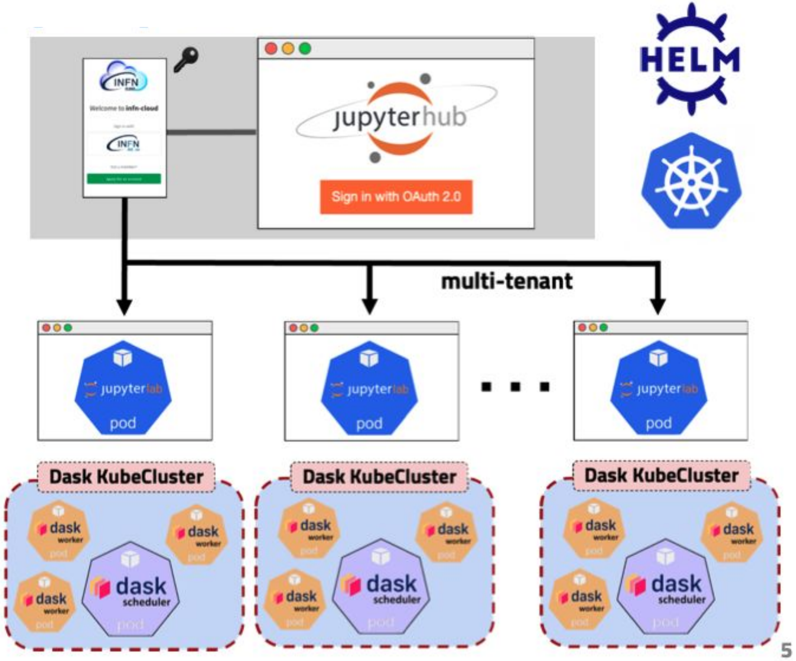

Benchmarks for interactive analyses

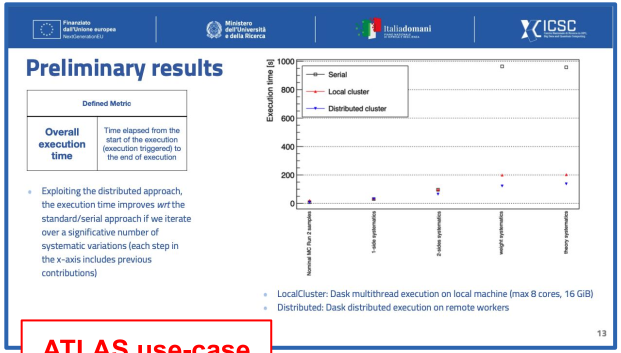

ATLAS use-case

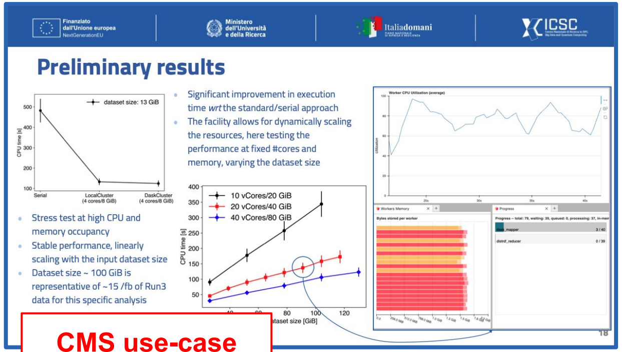

CMS use-case

source: https://indico.cern.ch/event/1338689/contributions/6010647

Operational experience from the Spanish CMS Analysis Facility at CIEMAT

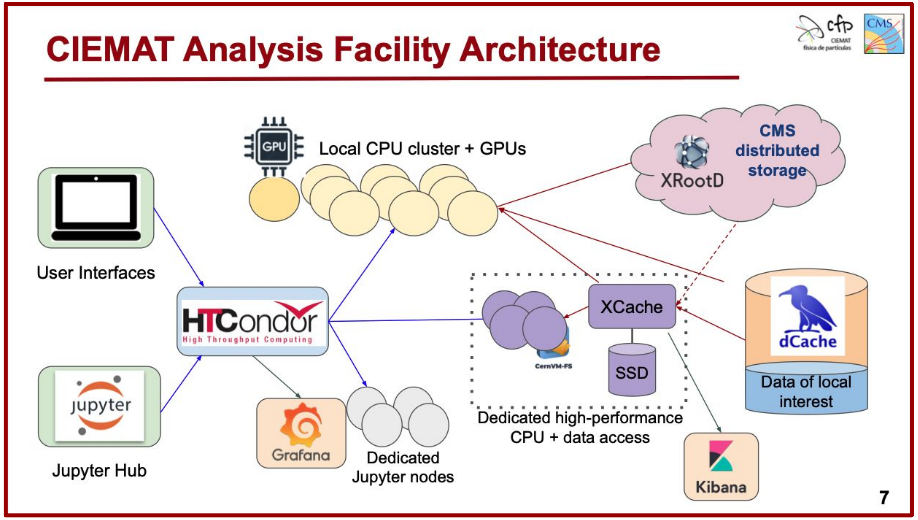

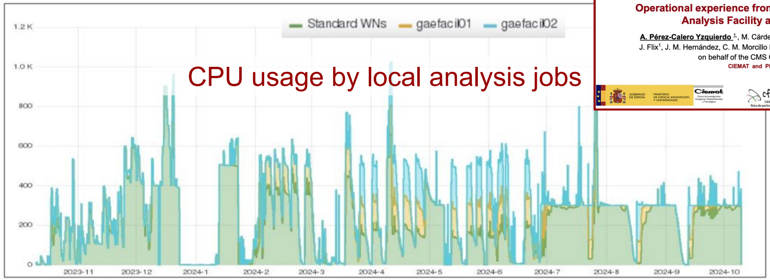

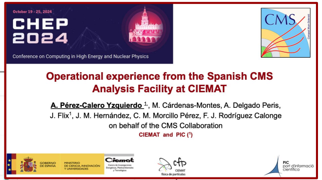

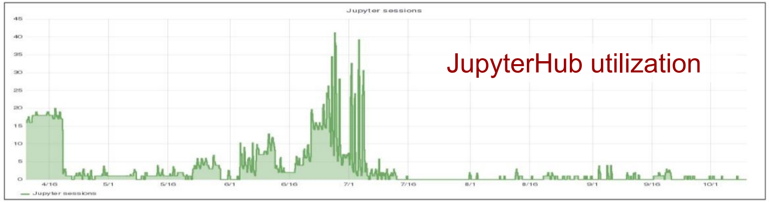

- Our objective from the start was to **empower** users…
  - Reduce "*time to insight*"
  - Fight the "*I can only really run at lxplus*" feeling
- …and help/**encourage** them **transition to new tools/technologies**
  - Jupyter, NanoAOD, RDataFrame (and others), GPUs

- New infrastructure is a **friendly** environment for non-HEP scientific groups at CIEMAT
  - Jupyter interface found convenient by many groups
  - NFS makes massive storage (dCache) widely accessible to users
  - HTCondor allows to ask for required resources (GPU, high performance SSD nodes, etc.), both from command line or through Jupyter

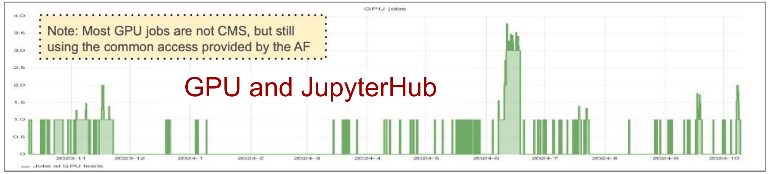

source: https://indico.cern.ch/event/1338689/contributions/6011544

Computing activities at the Spanish TIER-1 & TIER-2s for the ATLAS Experiment in the LHC Run-3 period and towards High Luminosity

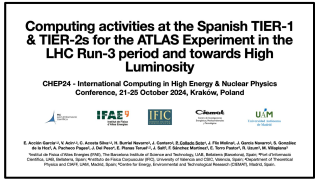

Leveraging HPCs

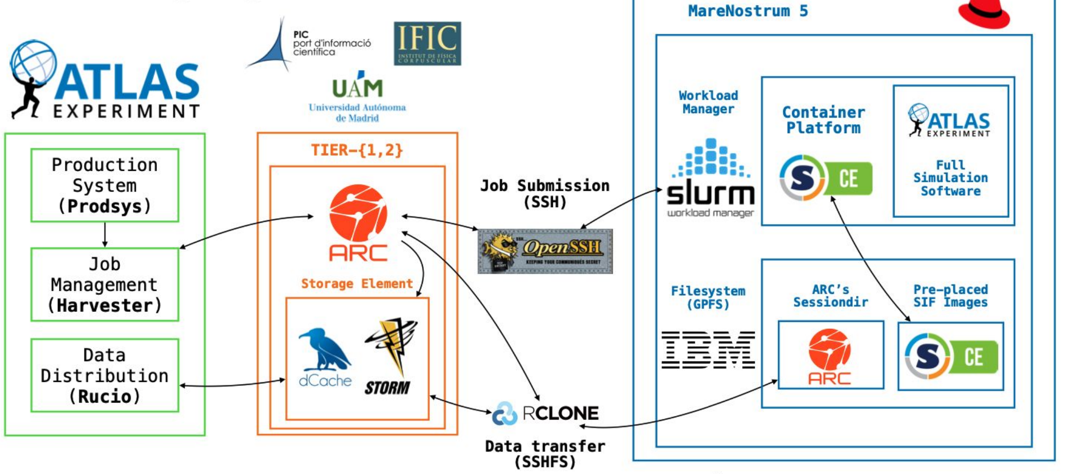

Note: MN5 **disallows** outbound and inbound **Internet** access…

HPCs in numbers

- On **April 2024** the BSC migrated from MN4 to **MN5**.
- Only the **simulation workflow** has been validated.
- **30 million hours** are approved at MN5 every year for ATLAS:
  - Through **Spanish R&D gateways**.
  - This corresponds to **50%** of **simulation** jobs assigned to Spain.
- HS23 / core = **27,1155**

Growing interest in incorporating ***HPC resources*** *into computing budgets available for experiments*

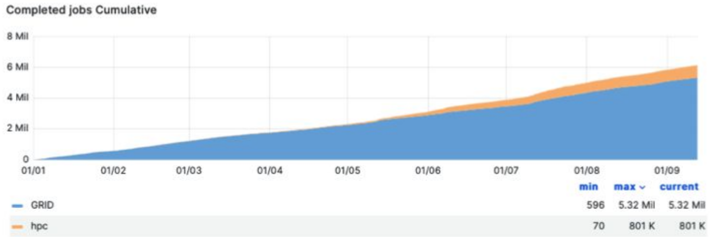

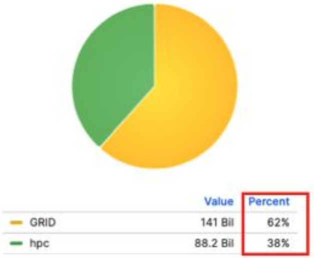

source: https://indico.cern.ch/event/1338689/contributions/6010707

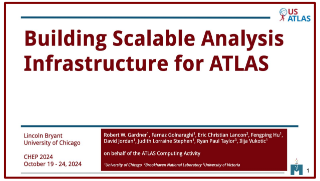

Building Scalable Analysis Infrastructure for ATLAS

US ATLAS is working on a **federated analysis platform** prototype that unifies resources from multiple facilities

- **Identity**
  - Leveraging existing, experiment-specific identity for authentication and authorization
- **Network sites**
  - Utilizing modern network overlay technologies to provide seamless connectivity between
- **Data**
  - Broad deployment of data caching infrastructure
- **Compute**
  - Embracing notebook-centric technologies and Pythonic frameworks, leveraging advances in Identity, Network and Data
- **Policy**
  - Developing policy framework(s) that provide an easy on-ramp for experiment end-users to use resources at all AFs

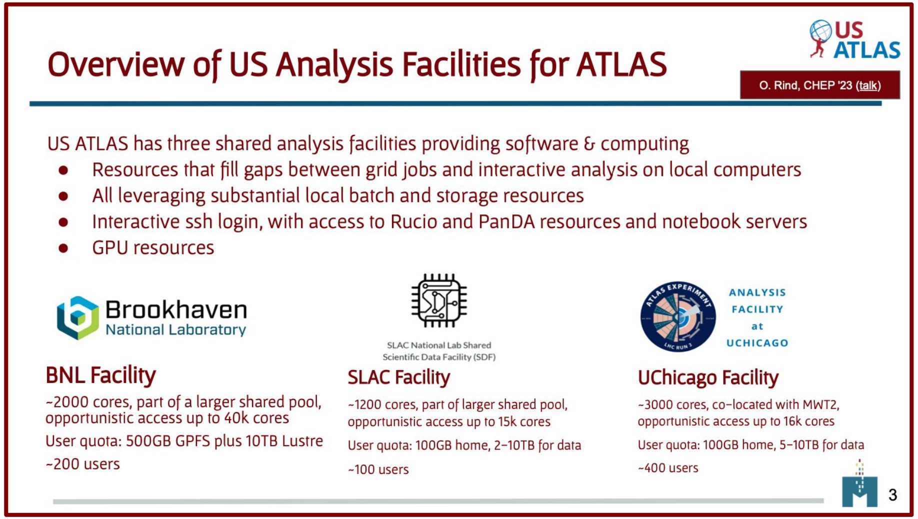

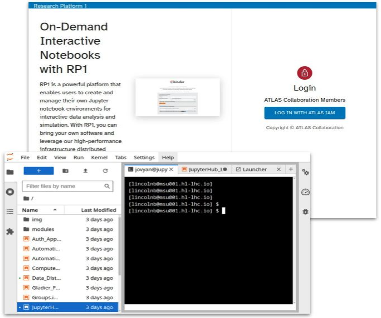

source: https://indico.cern.ch/event/1338689/contributions/6010704

Benchmarking massively-parallel Analysis Grand Challenge workflows using Snakemake and REANA

Marco Donadoni[1] Matthew Feickert[2] Alexander Held[2]
Andrii Povsten[3] Oksana Shadura[4] Tibor Simko[1]
[1]CERN [2]University of Wisconsin Madison (US)
[3]Princeton University (US) [4]University of Nebraska Lincoln (US)

27th Conference on Computing in High Energy and Nuclear Physics
October 21st-25th 2024, Krakow, Poland

Benchmarking massively-parallel Analysis Grand Challenge workflows using Snakemake and REANA

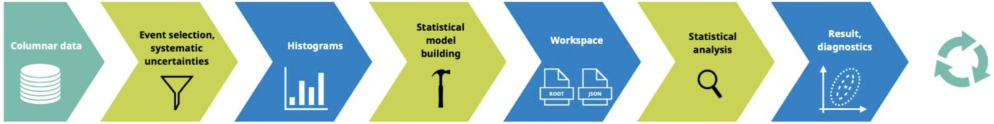

Setting up REANA test cluster

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

8

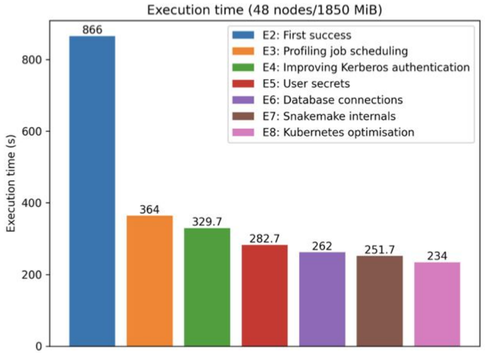

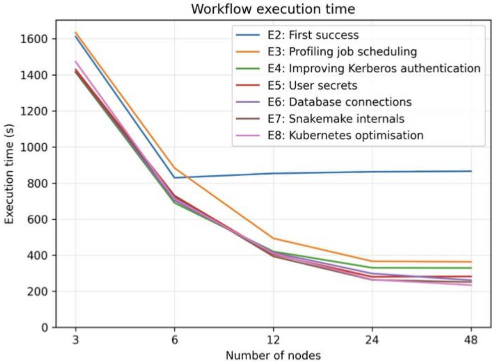

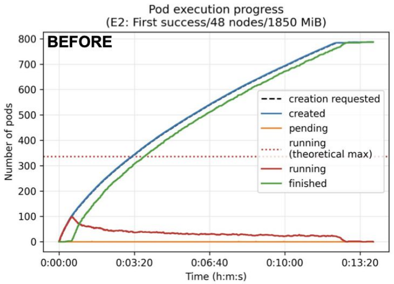

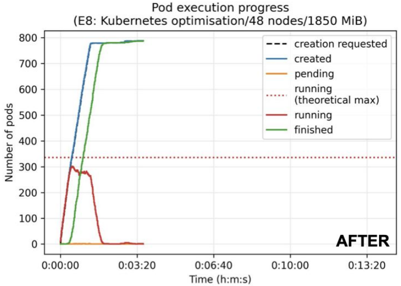

Final results

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

Conclusions

- Snakemake and declarative workflows can efficiently express massively-parallel particle physics computational paradigms
- Analysis Grand Challenge project is very useful to optimise performance of analysis platforms
- Optimisations allowed to improve REANA performance for massively-parallel workflows by a factor of ~3x

source: https://indico.cern.ch/event/1338689/contributions/6010709

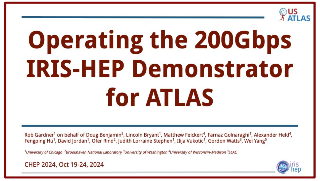

Operating the 200Gbps IRIS-HEP Demonstrator for ATLAS

Rob Gardner[1] on behalf of Doug Benjamin[2], Lincoln Bryant[1], Matthew Feickert[4], Farnaz Golnaraghi[1], Alexander Held[4], Fengping Hu[1], David Jordan[1], Ofer Rind[2], Judith Lorraine Stephen[1], Ilija Vukotic[1], Gordon Watts[3], Wei Yang[5]

[1]University of Chicago [2]Brookhaven National Laboratory [3]University of Washington [4]University of Wisconsin-Madison [5]SLAC

CHEP 2024, Oct 19-24, 2024

Operating the 200Gbps IRIS-HEP Demonstrator for ATLAS

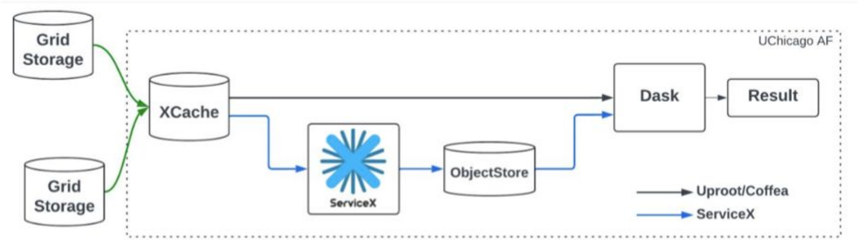

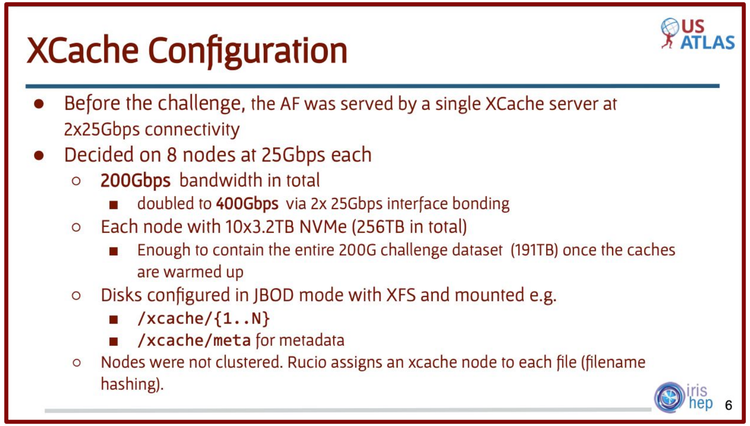

XCache Configuration

- Before the challenge, the AF was served by a single XCache server at 2x25Gbps connectivity
- Decided on 8 nodes at 25Gbps each
  - **200Gbps** bandwidth in total
    - doubled to **400Gbps** via 2x 25Gbps interface bonding
  - Each node with 10x3.2TB NVMe (256TB in total)
    - Enough to contain the entire 200G challenge dataset (191TB) once the caches are warmed up
  - Disks configured in JBOD mode with XFS and mounted e.g.
    - `/xcache/{1..N}`
    - `/xcache/meta` for metadata
  - Nodes were not clustered. Rucio assigns an xcache node to each file (filename hashing).

6

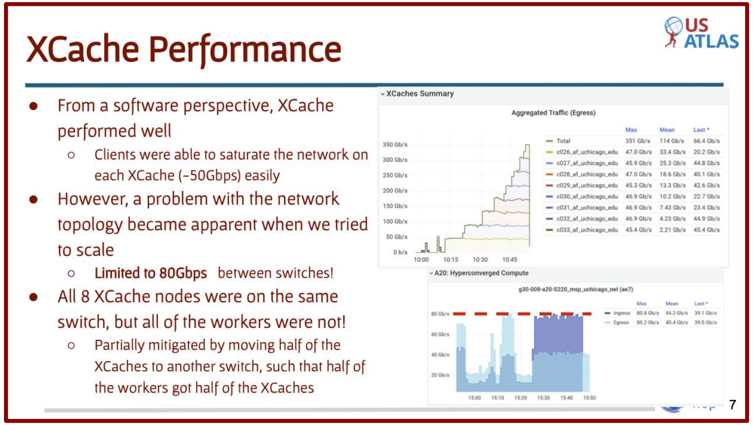

XCache Performance

- From a software perspective, XCache performed well
  - Clients were able to saturate the network on each XCache (~50Gbps) easily
- However, a problem with the network topology became apparent when we tried to scale
  - **Limited to 80Gbps** between switches!
- All 8 XCache nodes were on the same switch, but all of the workers were not!
  - Partially mitigated by moving half of the XCaches to another switch, such that half of the workers got half of the XCaches

7

Results – Success!

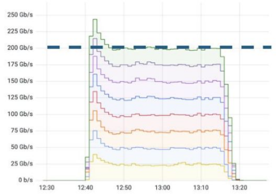

Read Method: uproot

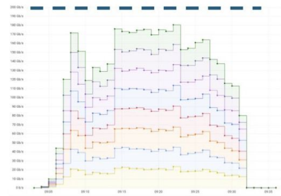

Read Method: ServiceX

source: https://indico.cern.ch/event/1338689/contributions/6010702

Tuning the CMS Coffea-casa facility for 200 Gbps Challenge

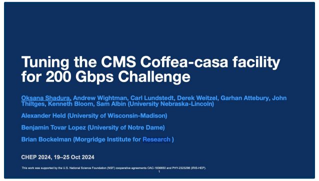

*The goal is to read **25% branches** out of **180 TB dataset** and to process it in **30 minutes***

⇩

In case of NanoAOD (~2 kB event size), the dataset size should be **<u>90 B events</u>** and **<u>to analyze it at 50 MHz</u>**

⇩

With current rate 25 kHz / core,

<u>we will need 2000 cores</u>

(12.5 MB/s per core)

Preparing next generation of Analysis Facilities

Running analysis frameworks and tools at scale

- **Adopt diverse computing executors** to support execution of complex task graphs
  - **Dask, TaskVine**
- Flexible computing **resource provisioning model** optimised for given facility
  - **Kubernetes, Tier-2 resources**, HPC cluster
    - e.g. Dask Gateway, dask-jobqueue, Dask Operator

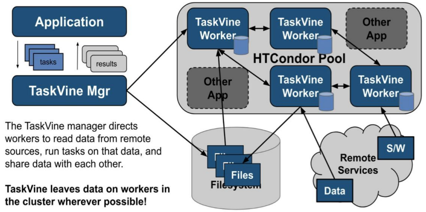

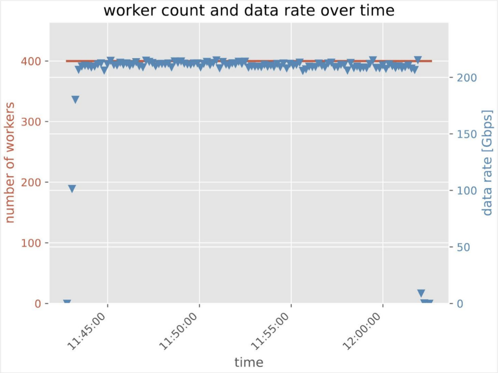

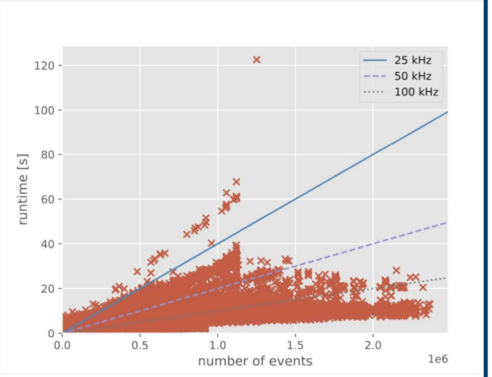

(Some of the) lessons learned

- Very **successful exercise format**: huge amount of progress and activity within 8 weeks
- Faced some challenges with memory use and scaling to all available resources
- **NanoAOD:** very large effect of compression algorithm: switching from LZMA to ZSTD brought 2.5x event processing rate improvement
- **Scaling Dask to 2k+ workers** generally works fine, need more testing combining large numbers of workers and very complex graphs
- Good performance observed also with **TaskVine** as alternative scheduler for graphs
- Scale of challenge allowed to identify **new bottlenecks** (many of which have already been fixed)

source: https://indico.cern.ch/event/1338689/contributions/6010711

Analysis Facilities approach for *ML and GPU access*

- Analysis Facility-like approach to provision GPU for ML/AI developments in INFN Cloud
  - platform providing **interactive mode** for developments and **offloading** (prototype) to scale up
  - [contribution link][summary slide]
- GPUs provisioning via interactive HTCondor jobs
  - **interactive access** and/or **SSH connection** to GPU resources via HTCondor at University of Glasgow
  - [contribution link][summary slide]
- MLFlow-powered pipeline for Machine Learning training
  - packing the entire **lifecycle of a ML application** for remote submission via SLURM
  - [contribution link][summary slide]

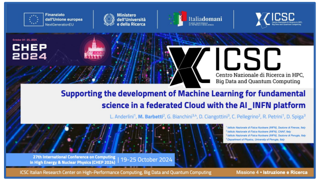

Supporting the development of Machine Learning for fundamental science in a federated Cloud with the AI_INFN platform

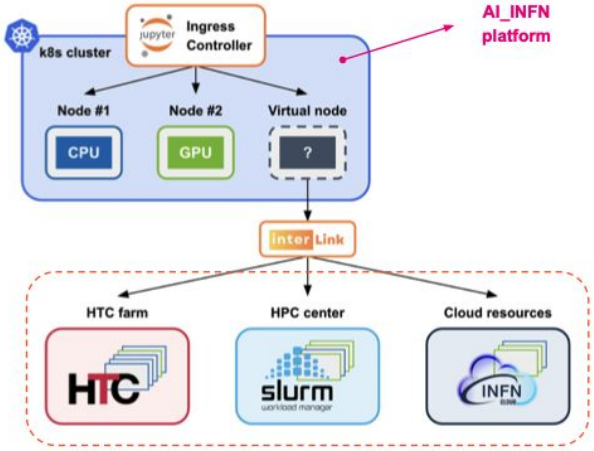

extension of the AI_INFN platform through the VK mechanism

AI_INFN's primary activities include:

- 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

AI_INFN
Artificial Intelligence technologies for INFN research

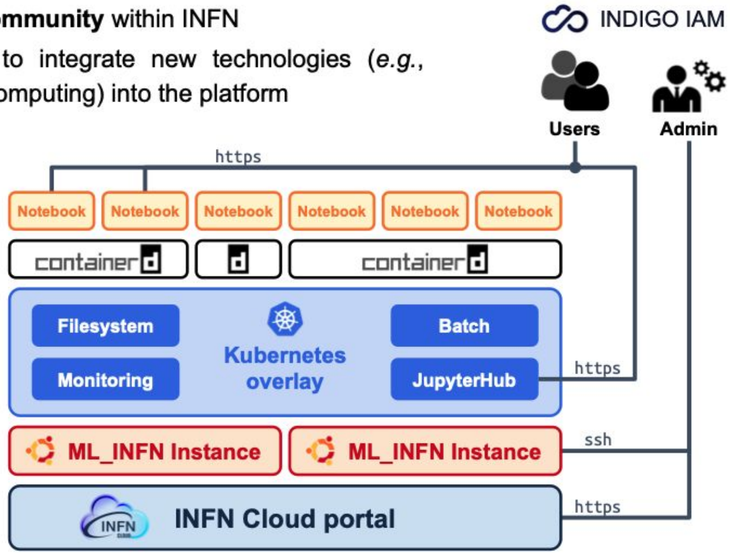

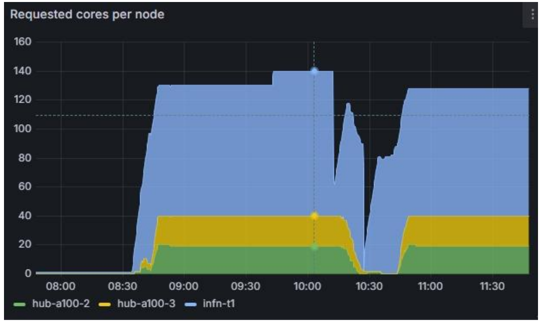

source: https://indico.cern.ch/event/1338689/contributions/6010715

✨ 🚜 ☁️

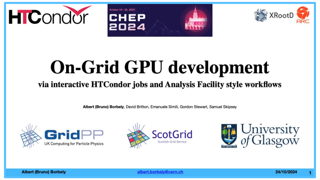

On-Grid GPU development via interactive HTCondor jobs and Analysis Facility style workflows

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

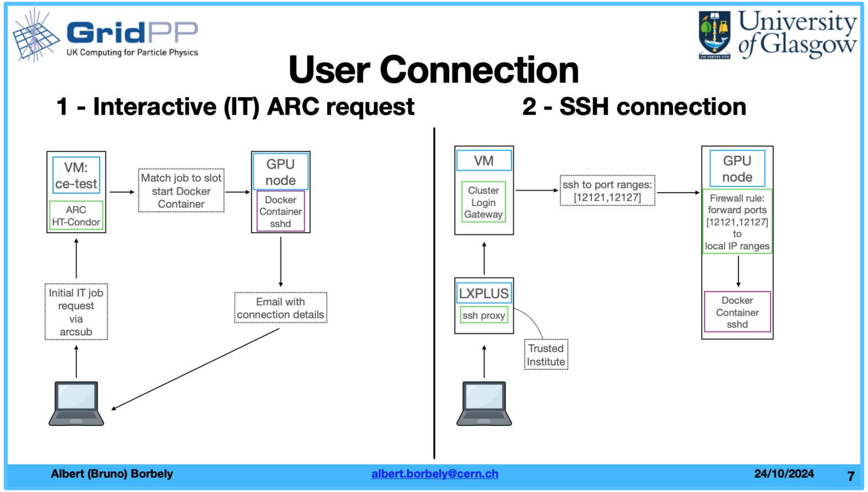

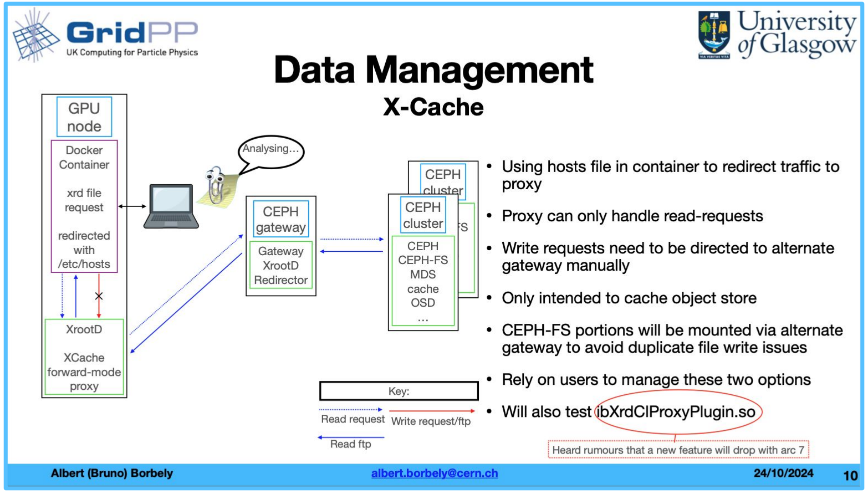

source: https://indico.cern.ch/event/1338689/contributions/6010708

Machine Learning Training Facility at Vanderbilt University

MLTF.VU

Machine Learning Training Facility at Vanderbilt University

Jethro Gaglione, Andrew Melo, Satkar Dhakal, Lindsey Fox, Matt Heller, Prasanna Koirala, Samuel Ninchritz, Brandon Soubasis, Alan Tackett

Artifacts

- model
  - data
  - MLmodel
  - conda.yaml
  - python_env.yaml
  - requirements.txt
- tensorboard_logs
- model_summary.txt

Full Path:mlflow-artifacts:/0/a81a0ab31f0447e48b5021637b877945/artifacts/model/MLmodel

Size: 705B


```yaml
artifact_path: model
flavors:
  python_function:
    data: data
    env:
      conda: conda.yaml
      virtualenv: python_env.yaml
    loader_module: mlflow.tensorflow
    python_version: 3.10.4
  tensorflow:
    code: null
    data: data
    keras_version: 2.14.0
    model_type: keras
    save_format: tf
mlflow_version: 2.8.0
model_size_bytes: 1694039
model_uuid: a38eac5c4cb24de691c0a3a0ea5419bc
run_id: a81a0ab31f0447e48b5021637b877945
signature:
  inputs: '[{"type": "tensor", "tensor-spec": {"dtype": "float64", "shape": [-1, 28,
    28]}}]'
  outputs: '[{"type": "tensor", "tensor-spec": {"dtype": "float32", "shape": [-1,
    10]}}]'
  params: null
utc_time_created: '2023-11-22 23:50:51.899498'
```


Motivation and Facility Overview

- Machine learning/AI has been widely adapted for an array of uses in HEP and beyond.
- However, there are considerations about training these models -
  - *Reproducibility - Can someone later consistently produce the same weights?*
  - *Scalability - How does one scale the training from card->node->cluster scale?*
  - *Efficiency - What are the barriers to using high-speed transports and interconnects?*
- MLTF is a prototype facility that enables training w/these considerations in mind -
  - *Hardware - At this stage, capabilities were chosen over large scale*
  - *Software - Easy-to-use and documented hooks/wrappers for clients to enable functionality*
  - *Infrastructure - Services run at the site to connect the client software to hardware resources*

MLflow Gateway

- Remote submission of MLflow projects opens interesting use-cases
  - *No need to get local accounts @ target cluster*
  - *CI/CD applications*
- MLflow has built-in support for remote submission via DataBricks
  - *Extension API allows this to be pluggable*
  - *Support for SLURM submission exists due to MLflow-SLURM*
- We are developing MLflow Gateway, a Client/Server interface to job submission
  - *Client side is an MLflow backend plugin*
  - *Server side is a REST service which interacts with local cluster*
- Since the gateway is an "intermediary" between the client and the facility, this provides a hook to add customization to the jobs before they are submitted to the batch system
  - *Difficult to do in direct SLURM submission*
  - *A task requesting a Rucio dataset as an input can trigger the Rucio rule, and wait for the replication to complete before submitting jobs to batch system*

source: https://indico.cern.ch/event/1338689/contributions/6010710

More from Track 9

- “Distributed Analysis in Production with RDataFrame” (link)
  - ongoing developments to perform distributed computations with RDF
- “Web-based graphics in ROOT” (link)
  - JavaScript ROOT for seamless integration and enhanced functionality across various platforms
- “Reshaping Analysis for Fast Turnaround” (link)
  - columnar data and task graphs offer a promising new paradigm for analysis software optimization
- “Optimizing Resource Provisioning Across Diverse Computing Facilities with Virtual Kubelet Integration” (link)
  - k8s orchestration system to federate scattered resources and optimize workflow via Bayesian model

List of presentations of Track 4 related to tokens

- WLCG transition from X.509 to Tokens: Progress and Outlook
- Evolving INDIGO IAM towards the next challenges (Enrico)
- Early recommendations from the Token Trust and Traceability Working → Security
- Fermilab's Transition to Token Authentication
- CMS Token Transition
- Supporting medium/small-sized experiments in the transition from X.509 to JWTs (Carmelo)
- Addressing tokens dynamic generation, propagation, storage and renewal to secure the GlideinWMS pilot based jobs and system.

Others from Track 4 related to Monitoring and Security:

- Advanced monitoring capabilities of the CMS Experiment for LHC Run3 and beyond → Monitoring
- Unified Experiment Monitoring → Monitoring
- Designing Operational Security Systems → Security

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!
  - Lessons learned: Token implementations of middleware need to improve
- WLCG updates:
  - Version 2.0 of the WLCG Token Profile is under preparation
  - WLCG is engaging with the Grand Unified Token Profile Working Group
  - Migration of WLCG IAM instances from OpenShift to Kubernetes
    - old services will be reused for stress tests → good for IAM developers
  - Token Management Enhancements: stop storing access tokens in the DB to improve the performance. They expect a IAM release before end of 2024.

🧑‍💻 🔒 ☁️

Tokens - Experiments & FTS

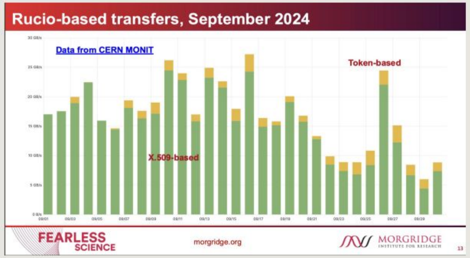

Brian Bockelman - CMS Token Transition

- CMS has made strides in token usage:
  - Every CE but 3 use tokens for pilot submissions, analysis to follow
  - CMS is using tokens in production (via Rucio) since early September (in over 30 sites)
  - 1 token per dataset, IAM can handle just fine
  - By the end of 2025 all services should be able to handle both x509 and tokens!
- ATLAS uses tokens in production since late August
  - no tokens during weekends, typical token rates are 1-2 Hz, spikes of 5 Hz, lifetimes currently are 2 weeks, to be reduced with more experience
- Use of tokens → focus is on FTS workflows: a new model was proposed for testing purposes
  - scopes per individual file and long-ish lifetimes, no exchanges or refreshing, token expires → transfers will fail → ball back to Rucio/DIRAC
  - the process used in DC24 still remains required by communities within and outside of WLCG
- Fermilabs uses a Vault instance to manage tokens since one year (with CILogon OIDC Provider)
  - Vault hides the token complexity from users
  - Vault is paid but there is a very promising open source version (OpenBao)
  - Vault is used to store high-value long-lived refresh tokens → configured with package htvault-config → the combination is called HTVault
    - Adds modified OIDC plugin, oauth secrets plugin, and ssh-agent authentication plugin
  - CMS will do the same, possibly via a CERN instance

🧑‍💻 🔒 ☁️

Tokens - Security

- TTT WG slides → 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
  - The Token Trust and Traceability WG Aims to form best practices, for users, devs, service providers + issuers
    - per workflow → after a process of Risk Analysis
    - “token-aware” incident response procedures
- Designing Operational Security Systems
  - Security operations centre (SOC) fits with an overall cybersecurity plan such as the Trusted CI Framework → Be proactive to prevent cybersecurity incidents: monitor, detect, respond
  - The pDNSSOC package was suggested as a lightweight way for smaller sites to get the benefits of a SOC → focused on correlating DNS logs with Threat Intelligence

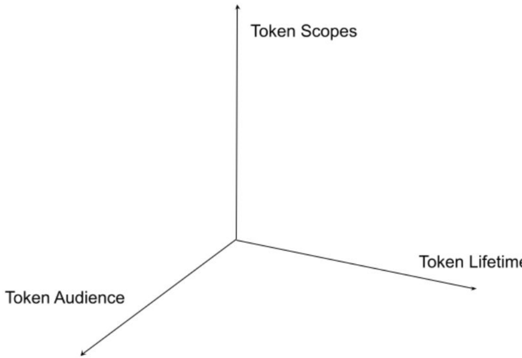

pDNSSOC outline

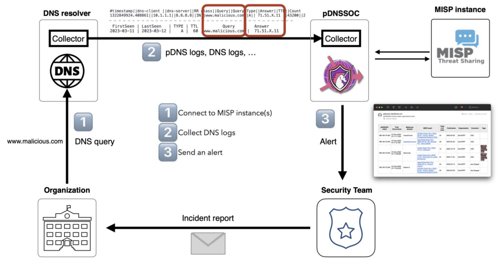

Operations/Monitoring

- “Data-flow: or 'death by a thousand acronyms!'” from WLCG central operations presentation:
    - Focus is on integrating heterogeneous resources, while keeping the grid operating

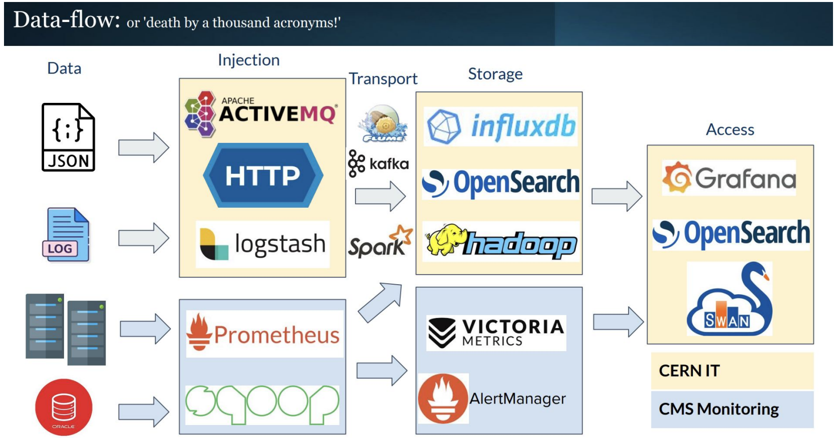

- Dashboard as a code described within Unified Experiment Monitoring presentation
    - Jsonnet and grafana extension grafonnet

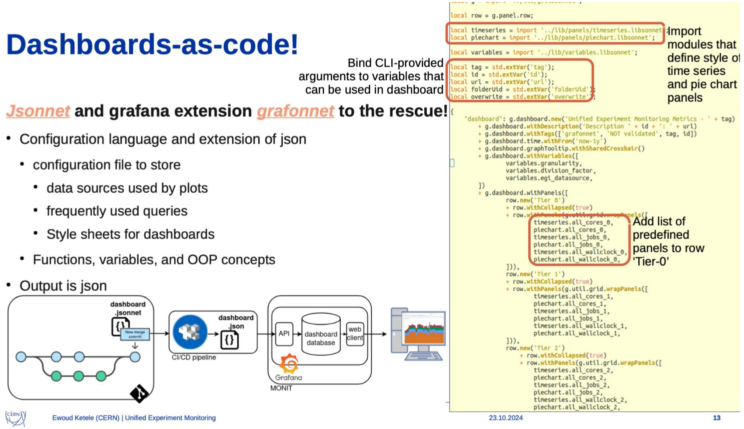

CVMFS: Pushing performance on highly parallel, many-core clients

Typical CPU Specs of current HPC nodes

- AMD EPYC 7h12: 64 cores, 128 threads
- AMD EPYC 7763: 64 cores, 128 threads
- AMD EPYC 7742: 64 cores, 128 threads
- AMD EPYC 7452: 32 cores, 64 threads
- ARM A64FX: 48 cores, 48 threads
- Intel Xeon-SC 8628: 24 cores, 48 threads
- Intel Xeon Platinum 8480: 56 cores, 112 threads
- Intel Xeon 9480: 56 cores, 112 threads
- Intel Xeon 8460Y: 40 cores, 80 threads

Maximum Core Count per Node

out of 34 ATLAS sites

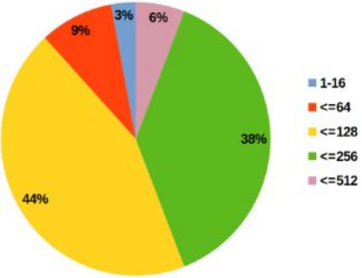

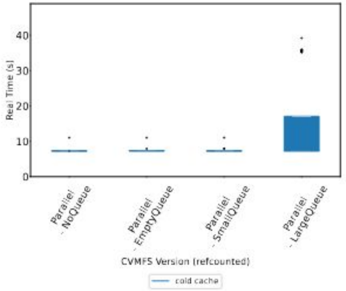

Different ROOT versions - 1 thread

Parallel Decomp can be same speed

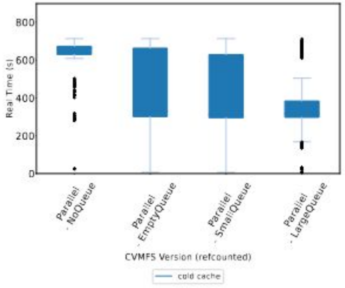

Different ROOT versions - 256 threads

Parallel Decomp 30 - 40% faster

DownloadManager: Old implementation

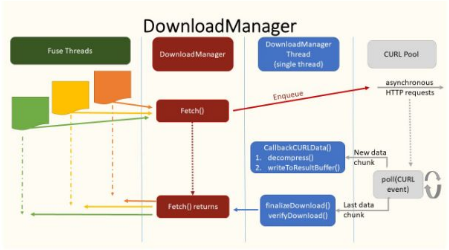

DownloadManager: New implementation

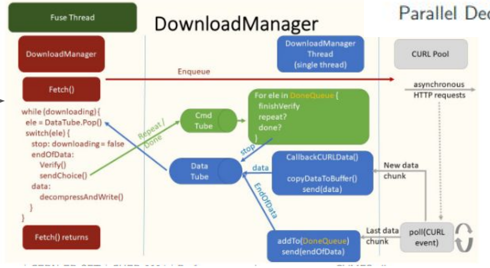

Automazione del supporto utenti

HEP-Help

Can a Large Language Model (LLM) be a first responder, either to answer questions or to send people to a forum where their question can be answered or has already been answered?

PyHEP 2019: I tried to get everybody to *use* StackOverflow, with tags to carve out our space within it.

JLab Future Trends 2022: I acknowledged that it's not going well.

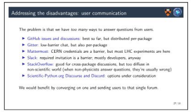

Analysis Ecosystem II 2022 and PyHEP.dev 2023: Brainstorming sessions, never landed on a solution.

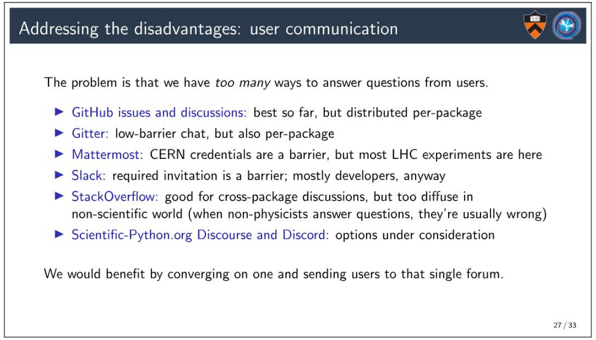

Addressing the disadvantages: user communication

The problem is that we have *too many* ways to answer questions from users.

- GitHub issues and discussions: best so far, but distributed per-package
- Gitter: low-barrier chat, but also per-package
- Mattermost: CERN credentials are a barrier, but most LHC experiments are here
- Slack: required invitation is a barrier; mostly developers, anyway
- StackOverflow: good for cross-package discussions, but too diffuse in non-scientific world (when non-physicists answer questions, they're usually wrong)
- Scientific-Python.org Discourse and Discord: options under consideration

We would benefit by converging on one and sending users to that single forum.

27 / 33