High Performance Computing in ATLAS

HEP/HPC Strategy Meeting
2025-01-30

Mario Lassnig (CERN)
on behalf of the ATLAS Experiment

Introduction :: Opportunistic resources in ATLAS

- ATLAS has a long history using opportunistic resources
- Harnessed HPCs for over a decade now
  - First generation HPC integration
    - For example, Cori & Titan (US) or SuperMUC (DE)
    - Majority requiring dedicated task submission
  - More recently, we have had enormous success
    - EuroHPCs Vega (SI) and Karolina (CZ)
    - Perlmutter and TACC (both US)
  - Several sites deploy HPCs as part of their pledge
    - For example, MareNostrum (ES) and Alps (CH)
    - Not opportunistic in this case
- ATLAS has also been able to use cloud resources
  - Cloud projects at ATLAS institutes such as Victoria & Oslo
  - Commercial cloud integration with Amazon and Google
  - Very recently, and following this work, a US Tier-2 has deployed cloud resources as pledge

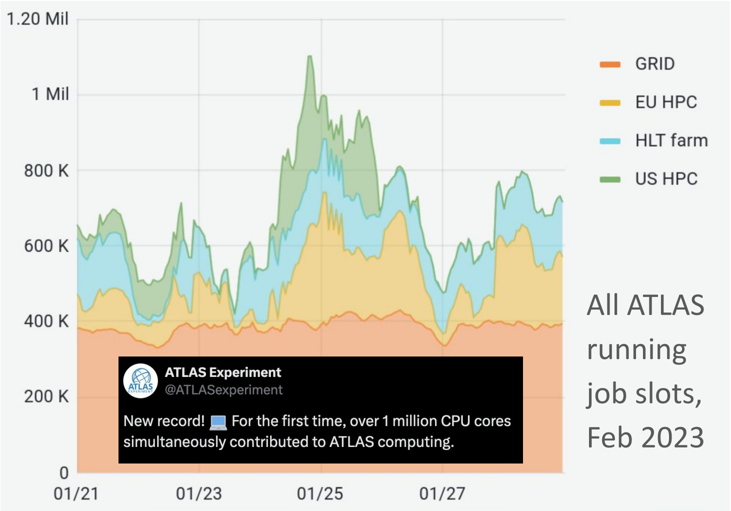

All ATLAS running job slots, Feb 2023

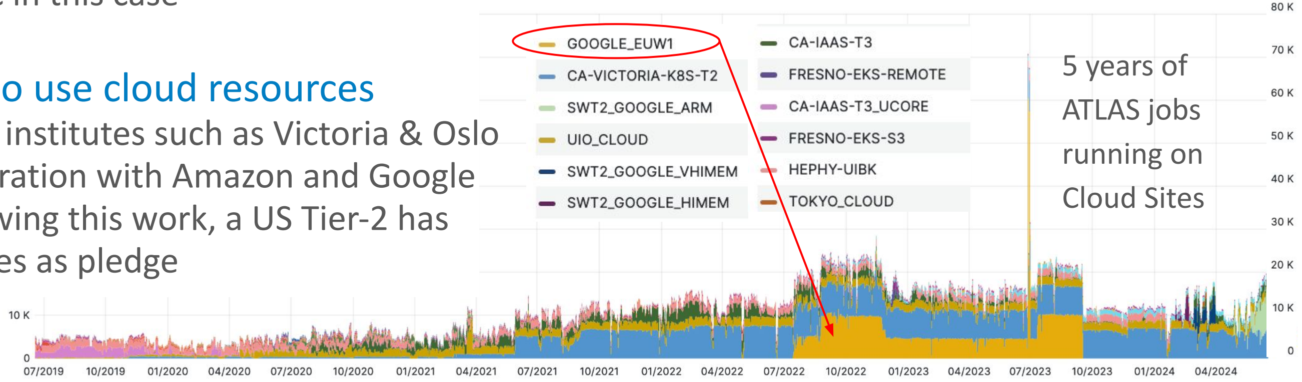

5 years of ATLAS jobs running on Cloud Sites

ATLAS resource usage :: Last 8 years

- HPCs give us access to large-scale compute resources beyond WLCG commitments
  - The WLCG is already providing these resources significantly above the pledge
  - On top of that, HPCs allow us to increase compute capacity even more by 30-40%
  - Looking at the last 90 days, we were running ATLAS jobs on 15+ HPC sites with various levels of integration

- HPCs play a key role in ATLAS Computing but come with significant challenges
  - They operate outside our standard WLCG policies, especially related to resource planning
  - Lack straightforward WLCG interconnectivity, especially related to grid storage and CVMFS
  - Enforce strict security and access policies with custom onboarding

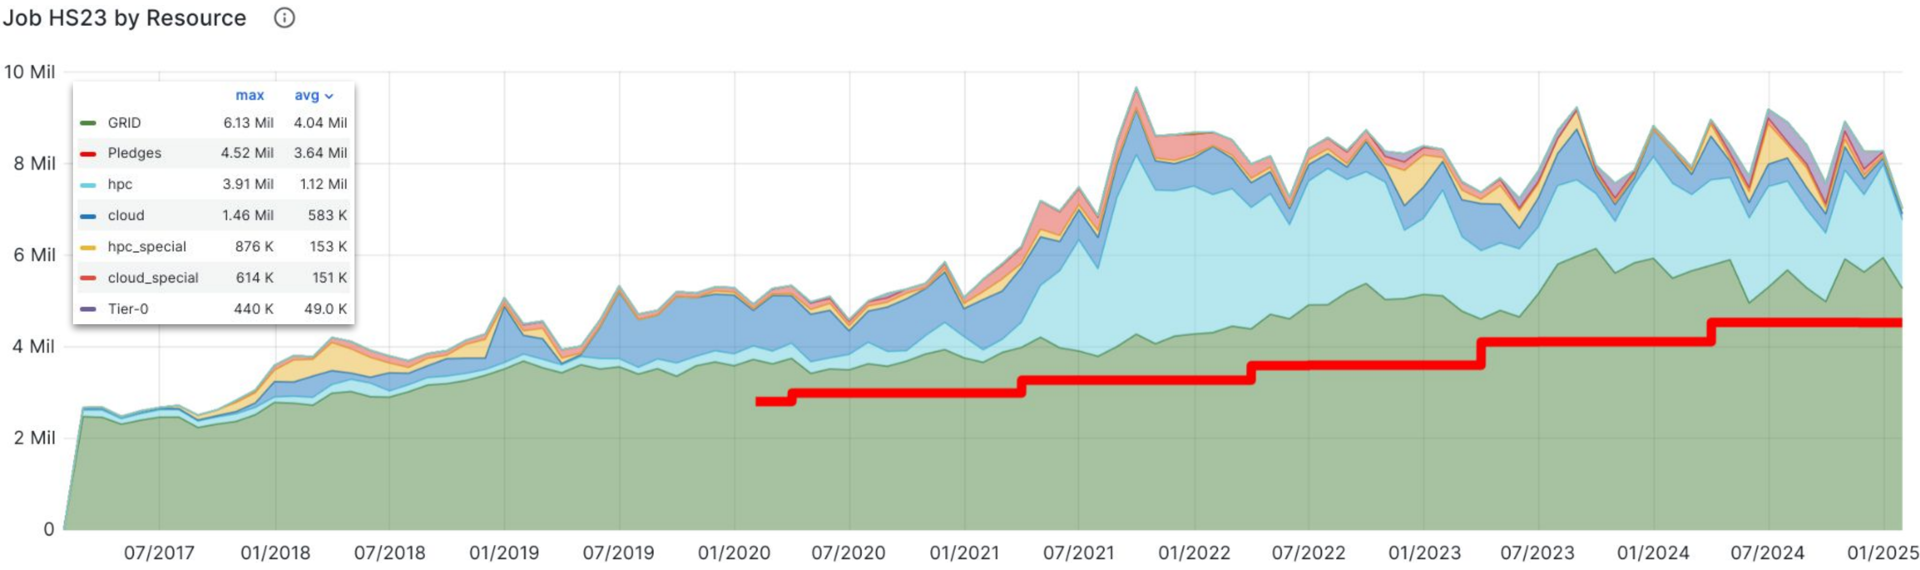

HPC resource usage :: Last 8 years

- Operational in ATLAS for many years, but vary in levels of integration
  - Some act like grid sites, allowing ATLAS jobs to run with minimal adaptation
  - Some need edge services: Harvester and/or ACT with combinations of pull/push jobs
  - Some HPC sites need to use dedicated containers for specific workloads

- ATLAS uses HPCs in various modes of operation
  - Can run the full job mix, including analysis and is fully integrated with the workflow systems, e.g. Vega
  - Require special submission tasks set up by the production managers and are usually only Full Sim, e.g. NERSC
    - But… Karolina can run any simulation, or Barbora can run any production MCORE workload…

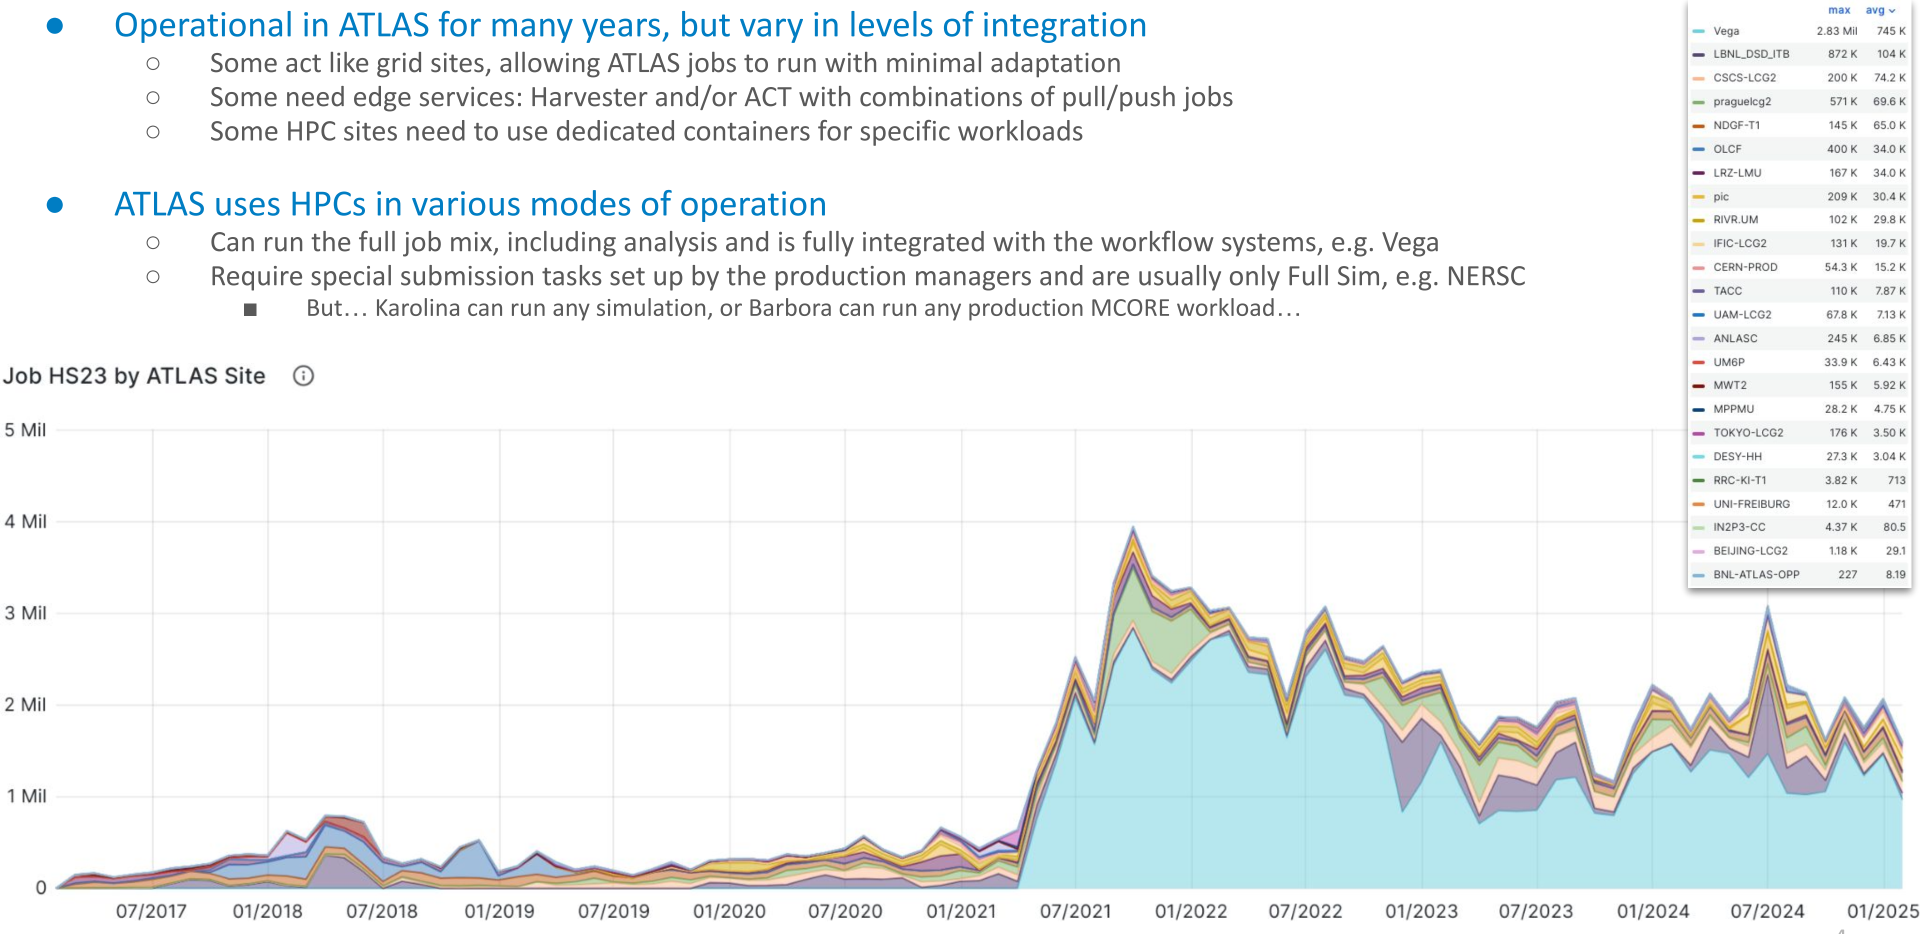

HPC workloads :: Last 8 years

- Operational in ATLAS for many years, but vary in levels of integration
  - Some act like grid sites, allowing ATLAS jobs to run with minimal adaptation
  - Some need edge services: Harvester and/or ACT with combinations of pull/push jobs
  - Some HPC sites need to use dedicated containers for specific workloads

- ATLAS uses HPCs in various modes of operation
  - Can run the full job mix, including analysis and is fully integrated with the workflow systems, e.g. Vega
  - Require special submission tasks set up by the production managers and are usually only Full Sim, e.g. NERSC
    - But… Karolina can run any simulation, or Barbora can run any production MCORE workload…

Slots of Running jobs by ADC activity

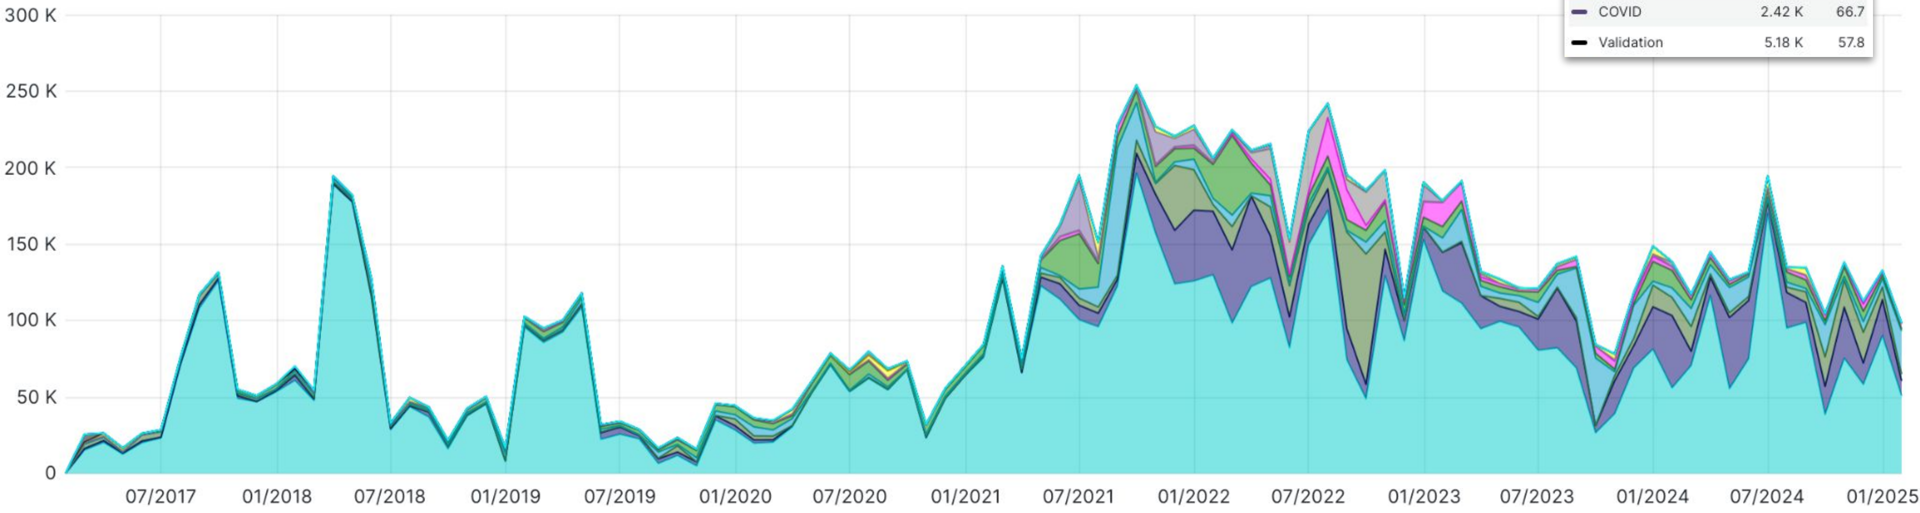

Thoughts on HPC integration

- Operational management
  - This is very decentralised, with individual sites assuming responsibility for their operations
    - Personal connections between people carry the success of the integration and operation
    - Yet still, in many cases, ADC central operations team needs to commit effort to help solve issues
  - Shared usage of HPC resources across multiple sites is observed in some instances, and vice versa
    - For example: PIC, IFIC, UAM for MareNostrum; CZ for Karolina, Barbora, LUMI; ND and CH for LUMI & PUHTI, DE NHRs …

- Commitment of resources
  - The "use-it-or-lose-it" policy for CPU allocations at some HPCs is unfortunate
    - Sometimes we might not have the right amount of jobs, or the right job mix, available
  - National FAs typically co-fund HPCs and some contribute them as pledged resources
    - For ATLAS, only resources that can execute the full job mix can be pledged
    - Must avoid a scenario where large resources can only run limited number of workflows – but there are technical constraints
    - Supplying resources in bursts or only for a limited period of time also doesn't fit our needs – we prefer steady state
  - Commercial interconnect mechanisms, like Globus for data exchange, goes against the open ATLAS spirit

- Commitment in time
  - Grid sites are typically expected to remain operational for the duration of the experiment
  - HPC machines typically operate for several years before being decommissioned and/or replaced
    - This can involve significant architectural or technological changes, requiring extra work from us
    - Especially, as of now, GPU-heavy HPCs are not useful for ATLAS
    - Some partitions, e.g., GPU, can be come outdated after a few years
    - The potential impact of HEP in the architectural/technological decision making process is unclear
    - Need to be compatible with different timescales of HPC sites vs HPC machines

Thoughts on future integration

	Examples
Provisioning	
○ Dynamic allocation based on available resources	*50k HIMEM submissions/day*
○ Plannable long-term allocations	*10M HS23/month*
Interfaces	
○ Ensure compatibility with our central systems	*Harvester, CRIC, …*
○ Integration with HEP AAI mechanisms	*Specifically IAM-compatible tokens*
○ The human interface	*Operator on site to help with issues*
Software	
○ Heterogeneous hardware support in our software	*GPUs, RISC-V, FPGAs, …*
○ Portable AI/ML libraries	*<insert favourite here>*
○ Task splitters	*Fully exploit manycore slots*
Data	
○ Open data management protocols and interconnect	*HTTP/WebDAV*
○ Scalable ingress and egress integrated with experiment data management	*Latency, namespace, access*
○ Exploitation of caching & network capabilities	*DTNs to/from grid sites*
○ Local & remote access to support data and libraries	*Frontier, CVMFS*
Scheduling	
○ Complex workflows	*AI/ML-style jobs*
○ User analysis	*Can HPCs be analysis infrastructure?*