Vega Overview and Plans

Andrej Filipcic

Vega EuroHPC

- Established as a first EuroHPC in Q1 2021
- In production until Q1 2027

Specs:

- 130k cores (260k HT slots) AMD 7H12
- 240 A100 40GB GPUs
- 25PB Ceph (HDD/NVMe)
- 1PB Lustre (NVMe)
- 300 Gb/s WAN
- 3M HS06 on CPU

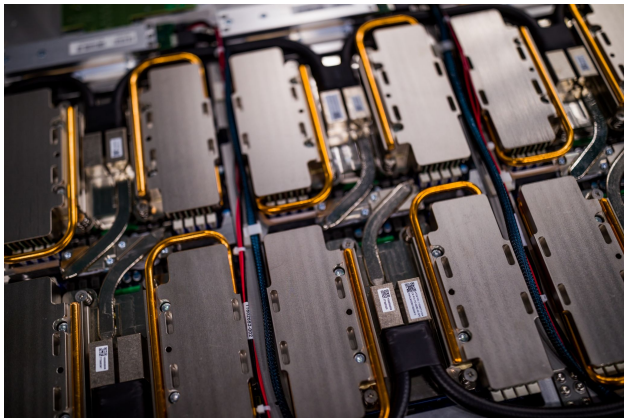

Running ATLAS

- All workflows including analysis
- Continuous slots since 2021 (up to 150k)
- Most slots used for simulation - not easy to fill 100k cores with data processing

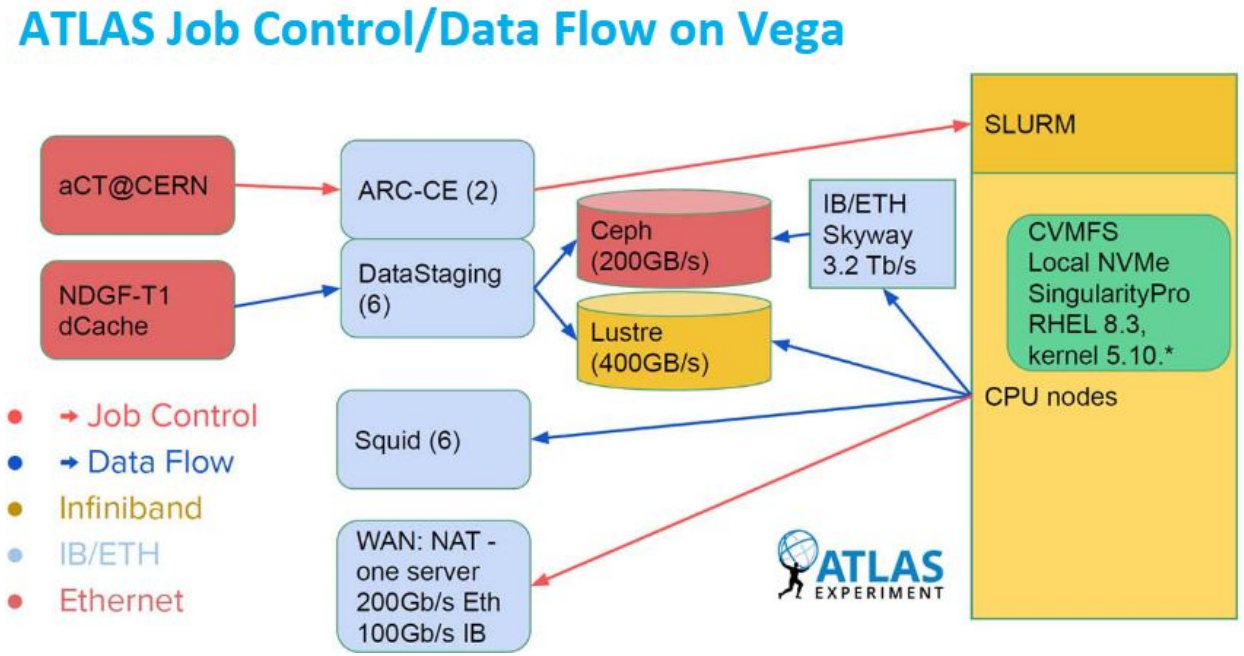

dCache pools

- Last week traffic
- 12 pool VMs on proxmox
- CephFS as a storage
- 6PB net space
- Very stable and little maintenance
- Some issues due to networking/ceph hiccups

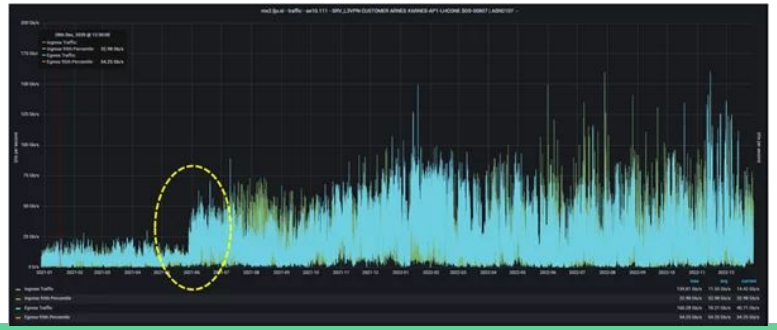

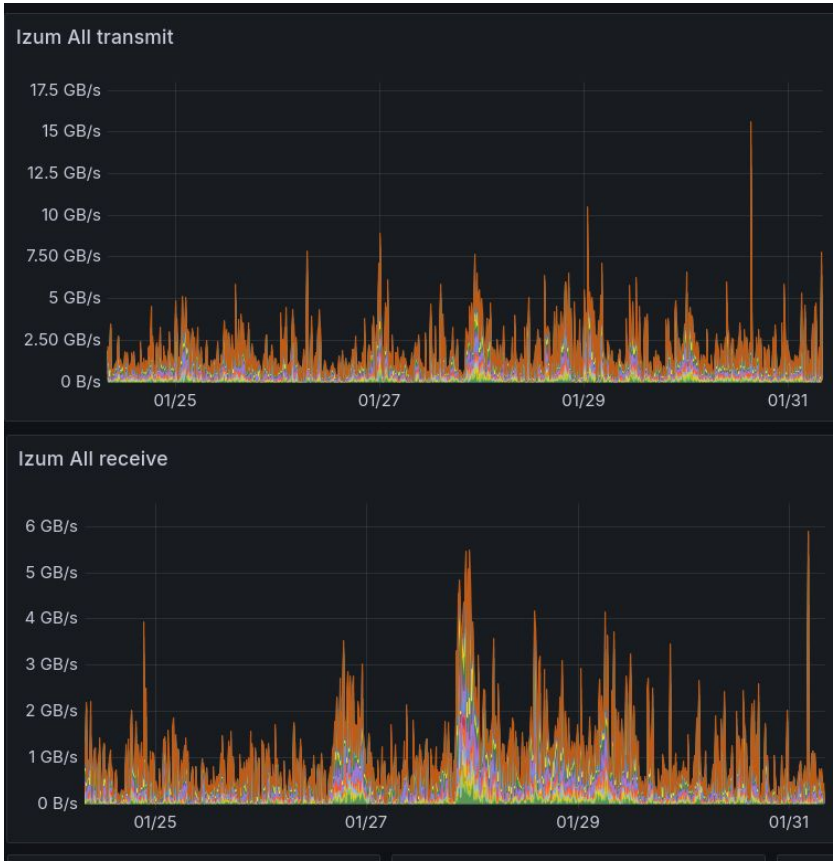

Access & Usage

- EuroHPC ⅓
- National ⅔ including support (pledge) for
  - ATLAS
  - Vera Rubin Observatory
  - Belle II
- Strategic projects:
  - Open-science data
  - Weather Forecasting
- Other projects:
  - InterTwin, DaFab, EuroHPC CoE, ...
- Development & Benchmarking access for LHC Experiments
  - Limited production
  - Scalability tests

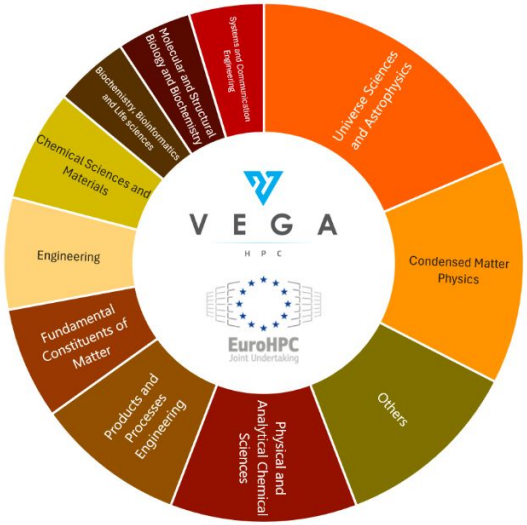

Vega-AI - New HPC after Vega (proposal stage)

- 150M budget (HPC, AI Factory, collaboration with Cineca)
- To start in 2027, 7-year lifetime
- CPU and GPU partitions, most of the funding for GPU
  - GPU likely in two phases
- Try to keep Vega functionality and usefulness for all communities
- Pressure for AI resources increasing in Slovenia as well
- 1.2 Tbps WAN