HEP Computing Users Meeting November 2024

David Hutchcroft, Rob Fay and Mark Wong

Purpose of this meeting

- The meeting is to collect information for future requirements and to see what issues people have currently
  - I circulated a request for input last week, results on the next slides
- Please feel free to raise additional issues in this meeting
  - Note this meeting does not replace helpdesk for specific issues
  - We are highly constrained in both the staffing and funding so you may not be able to get everything you want
- There is a reminder of resources available later in the slides

Items requested

- KLOE and MUonE asked for:
    - 1PB of disk space for large scale data storage (Frascati was not able to manage their requirements), this would be the only copy of ~100TB of raw experiment data so needs backups.
    - Estimated 100TB tuples “processed in a day” would be a goal
    - A request for MC processing facilities, with “all 400 cores” needed for months at the start of next year
    - Expect MUonE to also have fairly aggressive requirements, but not finalised

Items requested

- TCAD system for sensors modelling
  - Five people doing this work, very limited CPUs, making the simulation slow
  - Request for more a new dedicated machine but details not specified yet.
- Drawing office
  - Need two new laptops to run Windows 11 for CAD/Ansys (current ones too old to switch), plus third during next CG
  - Request for faster network links to a disk server to allow simultaneous editing of large models

Items requested

- ATLAS
  - Disk space: ~80-100 TB (2025), 100-120 TB (2026-2029)
  - CPU: 500 cores (all experiments)
  - GPU: current users covered but expected requirements to grow
  - Desktops: migrate current 4 to Alma9 asap, one more requested + GPU
- FASER
  - Predominately a CERN resource experiment
  - Disk: ~10 TB (2025-2029)
  - CPU: as above
  - Desktops: none presently

Items requested

- LHCb
  - Mixed use, a lot at CERN + grid
  - Disk: 50TB
  - CPU: sufficient at present but usage will grow
  - Desktops: needed for specific tasks at CERN

Resources now

- Interactive nodes:
  - Alpha (96 threads, 512GB RAM): GPU node with 6 Nvidia A100 GPUs
  - Phi (40 threads, 128GB RAM): Older node with 2 Nvidia RTX A4000 GPUs
  - Hepcuda1 (8 threads, 16GB RAM): Old low-end node with one Nvidia GTX 980 Ti
  - Gamma (64 threads, 394GB RAM)
  - Kappa (16 threads, 72GB RAM)
  - Theta (64 threads, 128GB RAM)
- Batch nodes:
  - Short queue with 4 nodes (8 threads, 16GB per node)
  - Main queue with 9 nodes
    - 5x 24 thread, 128GB RAM
    - 4x 80 thread, 384GB RAM, Nvidia RTX 2080 Ti GPU
  - Accessed by the slurm scheduling software (runs a fair share allocation method)

Resources now

- Storage:
  - Hepuser: Single server for home directories, backed up disk, limited quota (default 80GB soft, 100GB hard)
  - Hepstore: Single server, for individual and small experiment bulk data, default quota of 500GB (can be increased by request), not backed up.
  - Scratch: 11TB on dedicated but old server, deprecated
  - Bundle: Distributed file system across five servers intended for large experiment files (current system has performance issues with large numbers of small files creating potential maintenance and reliability issues, alternatives being explored)
  - Bundle/scratch: Distributed 22TB scratch area using batch node disks
  - Wisp: Single server cloud storage, 100GB quota by default, old hardware but backed up

Resources now

- Tier-2 Grid
  - 4,248 job slots across 68 compute nodes
  - Current ~1.2PB dCache storage across 8 servers
- HEP & Grid Computing Staff
  - Currently 1.5 FTE, down from 2 FTE (2020-2022), 3+FTE (<2020)
- Building networking and laptops are IT services responsibility

Money possibly incoming

CG requests for October 2025 to September 2029

These are requests and have not been granted…

- Tier 3 Computing (staffing): £129,564
- Computing (hardware): £141,235
- PP Computing cluster upgrades: £ 99,282
- PP Networks switches: £ 55,999

- Ongoing grid support for staff (0.5 FTE) and computing for GridPP Tier-2 use

Changes known to be required

- Centos7 to AlmaLinux9
  - Currently there is a working Alma9 image for desktops but few explicit requests came in to switch, now have more requests from this call for information
  - Would like to retire all Centos7 systems as out of support (old code can be run on new OS with wrappers like singularity)