HEP computer users meeting

David Hutchcroft on behalf of the HEP computing team

Agenda

- What we offer locally, CPU + disks
  - Note I will not be discussing Grid resources specifically, these are available to everyone in the VOs so not a “local” resource. You can ask about them though.
- What our staff can help with, what needs to go to IT services
- Replacements and upgrades, during next CG period (2 years)
- Feedback from experiments (very short as there was none)
- AOB

Available resources: interactive nodes

- CentOS 7:
  - Note support for the OS ended in June 2024…

Node	CPU Cores	GPUs	Memory	CPU date
phi	40 Xeon E5-2680	2 RTX A4000	132GB	2012
kappa	15 Xeon E5640	N/A	74GB	2010
hepcuda1	8 Core I7-4790	GeForce GTX 980 Ti	15GB	2014

- Alma 9:

Node	CPU Cores	GPUs	Memory	CPU date
gamma	64 Xeon Gold 5218	N/A	385GB	2019
theta	64 Xeon Gold 5218	N/A	128GB	2019
eta	64 Xeon Gold 6526Y	N/A	385GB	2023
alpha	96 EPYC 7413	6 A100 80GB	515GB	2021

Available resources: batch nodes

```
> sinfo --long --Node
Fri Feb 20 09:11:35 2026
NODELIST   NODES    PARTITION       STATE CPUS    S:C:T MEMORY TMP_DISK WEIGHT AVAIL_FE REASON
alpha          1        alpha        down 24     24:1:1  63200     50000      1   (null) Node unexpectedly re
eta            1         test       down* 2       2:1:1   2048     80000      1   (null) Not responding
hep73          1         test       inval 2       2:1:1   2048     80000      1   (null) maintenance
livdat1        1        short       down* 8       8:1:1  64000     50000      1   (null) Not responding
pc178134       1         test       down* 2       2:1:1   2048     80000      1   (null) Not responding
r27-n06        1        short        idle 8       8:1:1  15800    800000      1   (null) none
r27-n07        1        short        idle 8       8:1:1  15800    800000      1   (null) none
r27-n08        1        short        idle 8       8:1:1  15800    800000      1   (null) none
r27-n09        1        short        idle 8       8:1:1  15800    800000      1   (null) none
r27-n13        1     compute*        idle 24     24:1:1 128000    800000      1   (null) none
r27-n14        1     compute*        idle 24     24:1:1 128000    800000      1   (null) none
r27-n15        1     compute*        idle 24     24:1:1 128000    800000      1   (null) none
r27-n16        1     compute*        idle 24     24:1:1 128000    800000      1   (null) none
r27-n17        1     compute*        idle 24     24:1:1 128000    800000      1   (null) none
r27-n18        1  compute-al9        idle 72     72:1:1 102400    650000      1   (null) none
r27-n19        1  compute-al9        idle 72     72:1:1 384000    650000      1   (null) none
r27-n20        1  compute-al9        idle 72     72:1:1 384000    650000      1   (null) none
r27-n21        1  compute-al9        idle 72     72:1:1 384000    650000      1   (null) none
ronald         1    vidarrcpu        idle 24     24:1:1  63200     50000      1   (null) none
```

Available resources: local disk

```
> df -h
Filesystem                                Size  Used Avail Use% Mounted on
…
ihepslab05.ph.liv.ac.uk:/data             513T  175T  339T  35% /bundle/data
ihepstore2.ph.liv.ac.uk:/scratch           11T  3.0T  8.0T  27% /scratch
ihepstore.ph.liv.ac.uk:/hepstore          688G  8.9G  679G   2% /hepstore
ihepuser.ph.liv.ac.uk:/user                16T   12T  3.7T  77% /user
r27-n13.ph.liv.ac.uk:/scratch              22T   19T  3.2T  86% /bundle/scratch
imuonstore2.ph.liv.ac.uk:/muonstore2      262T   34T  228T  13% /muon/store2
imuonstore1.ph.liv.ac.uk:/muonstore1      373T  296T   78T  80% /muon/store1
```

Bundle data users:

Admin	g-2	MuEDM
ALICE	JUNO	MUonE
ATLAS	KLOE	particle_therapy
CMOS	LHCb	SHIP
Containers	LIVDAT	T2K
DarkSide	LZ	Users
FASER	Mu2e	
FCC	Mu3e	

Helpdesk + IT services

- Rob and Mark are here to help
  - However, they can only manage our systems, the building network/wifi/IT services laptops are all IT services issues
  - We can install missing software if it is easy to do, however most packages are on CVMFS which is created by CERN but publicly available to anyone
- David is responsible for the local GridPP registrations, but I can only grant you a grid certificate. Everything else is a virtual organisation responsibility.
- Reminder you can get 2TB of disk for any project from IT services by request (useful for backups if not live analysis).

IT services HPC

https://www.liverpool.ac.uk/research-it/high-performance-computing/ has the full details

The Barkla HPC cluster

- Liverpool's Barkla platform is a Linux HPC cluster consisting of 67 specialist computers (called nodes). Recent significant upgrades have boosted cluster performance and capabilities several-fold, opening doors for new research via:
- 58 *compute* nodes,* each with 168 cores (two AMD EPYC 9634 CPUs), 1.5 TB RAM (9 GB/core), and 3.84 TB local NVMe storage.
- 2 *visualisation* nodes - as compute nodes but each with two NVIDIA Ada Lovelace L4 GPUs. For remote desktop pre- or post-data GUI work, or debugging lightweight GPU apps.
- 4 *general purpose* GPU nodes - as compute nodes but each with two NVIDIA L40S GPUs.
- 3 *deep-learning focused* GPU nodes, each with 96 cores (two Intel Xeon Platinum 8468 CPUs), 2048 GB RAM (21 GB/core), four NVIDIA H100 SXM GPUs, and 7.68 TB local NVMe storage.
- network storage includes 2 PB for short- and medium-term work (NFS with backup), and 2 PB of Lustre parallel storage for all tasks, including I/O-intensive work. Nodes connect via a fast 200 Gb/s dual Intel Omni-Path interconnect.

(*) while available to all, a growing number of compute- and GPU-nodes are funded by certain research groups who have job priority.

Replacements required

- Need to replace the Centos7 nodes as they are very old + unsupported (note people still using software locked to this OS)
  - Docker/Apptainer/Singularity containers and other similar services allow “old” software to run on newer OSes.
- Main network switches are old, should be replaced but these are expensive
- VM hosts will need replacing, as will several UPSes

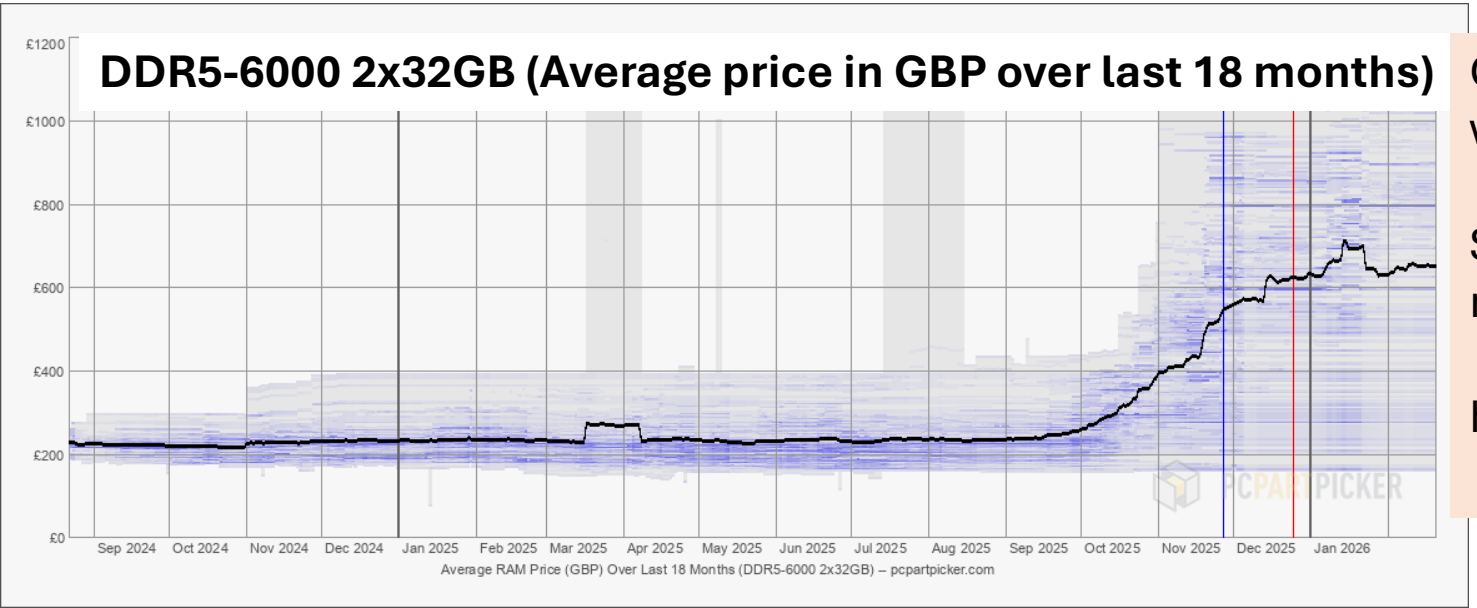

Currently memory prices are more than three times what they were in September last year

SSDs and hard disks are also much more expensive now

https://uk.pcpartpicker.com/trends/price/memory/

ATLAS requirements

thanks Carl, apologies for the heavy summarisation below

- Approx 2-3 more disk for analysis tuples 60 to 180TB and faster access through a high throughput cache disk
- Approx 10TB for pixel endcap production data (short/medium term) as the only copy, so should be backed up
- More local batch nodes on AL9 (moving C7 nodes sufficient)
- Desktop machines with specialised resources: memory for Ansays, large screens for MS Project.
- LCG lags in deploying new software so many ~10GB per user maintaining local python environments
- Maintain licence servers and software licences for several CAD and clean room services
- Several places where single system failures would cause significant problems, but none defined as catastrophic

LHCb requirements

- Work on LHCb SWAN based analysis cluster lead by Liverpool
- Datasets for Run 3 and 4 expected to be larger but manageable
- No specific requests for data storage locally

Any Other Business?