Facility Priorities During LS3

Fred Luehring (Indiana University)

Kaushik De (UT-Arlington)

9 June, 2026

US ATLAS @ HTC26

Discussion Topics

- Meeting the pledge
- Purchasing priorities
- Infrastructure upgrades
- Dell common pricing
- Storage and compute server purchas
- Common tools

Meeting the Pledge

- No hardware funding during I&C
  - Meet 2027 and 2028 pledge with end of CA funding in 2026
  - Purchases must be “completed” by Jan 31, 2027
  - Possibility of NCE
- Each site needs to take account of retirements
- Assume flat CPU pledge
- Slight increase in storage pledge
- US T2 pledge now equal for each of 4 sites
- Continue x2.2 for CPU, x1.2 for storage
  - Should we change x2.2 to x2.0?

Purchasing Priorities

- Highest priority is infrastructure.
  - Must cover infrastructure needs for the foreseeable future using the lump sum funding from the end of CA funds.
    - Network (ideally with 400G to the WAN
      - Too early to do 400G at some sites.
    - Power, Cooling
    - Servers that run the sites
  - Infrastructure can disable the part or all of a site if it fails.
- Second priority is storage servers
  - Failing storage can disrupt may jobs because data is spread over many servers.
    - Causes us to retire storage at no more than 7 years.
    - Also adding storage to an existing site takes time and effort.
- Last priority is compute servers
  - Failing CPU is far less disruptive than failing storage
    - We often run cmpute more than 10 years
    - Easy to add more compute to an existing site.

Infrastructure Upgrades

- Can we pool together network purchases for T2's?
  - Will we get better Dell pricing if we add infrastructure to CPU and storage purchases
- Sharing of ideas about equipment to buy
- In general try to get the universities to pay for as much infrastructure as possible.
  - The willingness of the universities to pay varies depending on site.

Dell Common Pricing

- Leverage large end of CA purchases by T2 sites
  - Need to push Dell to move quickly
    - We need to tell Dell a the purchase size.
      - And stick to that amount of spending.
      - I am leery of doing this until NSF releases the funding.
    - Need some dedicated meetings?
- Can we use T1 also in negotiations?
  - This is tricky because the T1 has DOE funding and different (stricter) purchasing rules

Storage Server Age

- We are currently in reasonable shape but things will be tough the start of HL-LHC without purchases now.

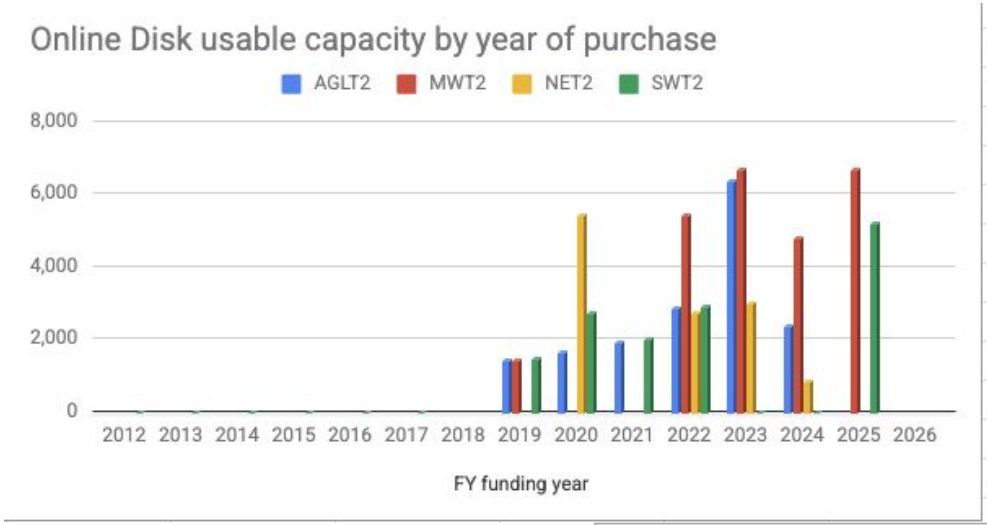

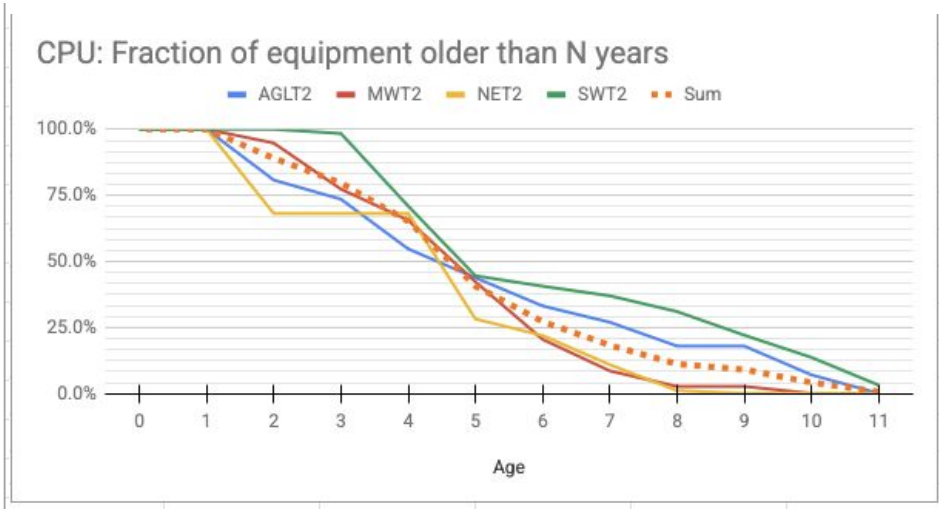

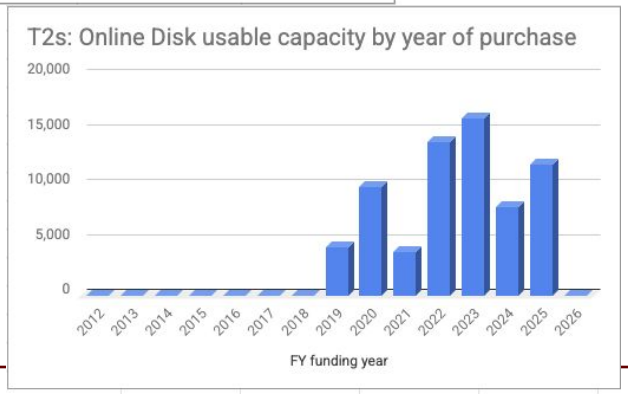

Compute Server Age

- We are currently in reasonable shape but things will be tough the start of HL-LHC without purchases now.

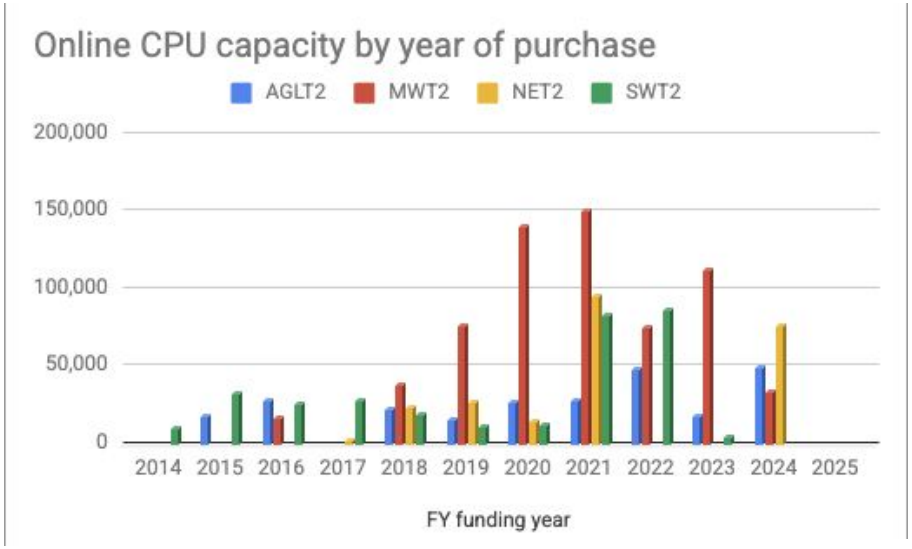

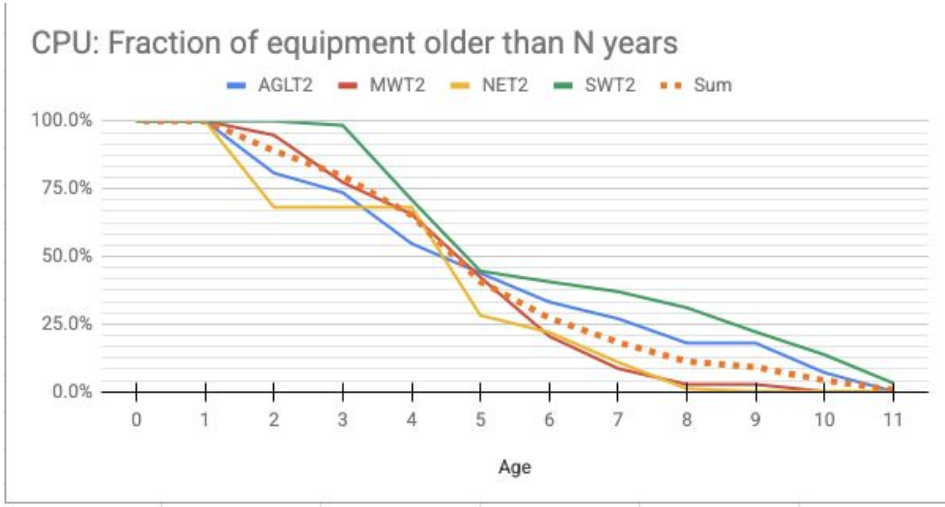

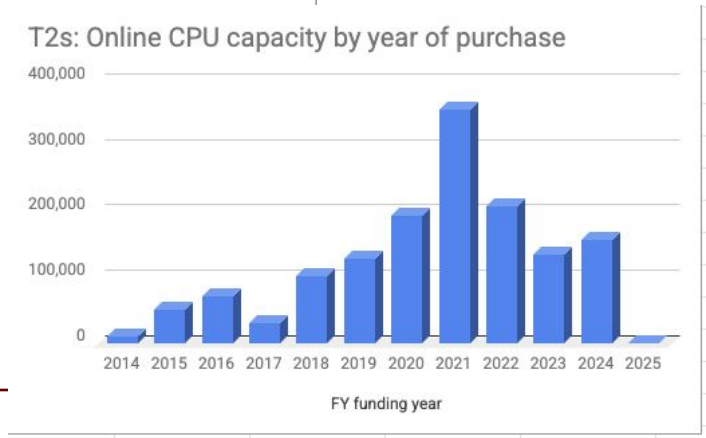

Server Purchasing Considerations

- Andre Sciaba gave an excellent presentation at HEPiX on recent hardware trends:
  - https://indico.cern.ch/event/1471803/contributions/6966459/attachments/3280524/5862278/Hardware%20Technology%20in%20the%20AI%20era%20-%201.0.pdf
- I covered much of the benchmarking in the talk at last Friday's procurement meeting:
  - https://docs.google.com/presentation/d/1BG3f3G3Xaiu8BUUIH5d4DOaiSyR0nX6kWGKxNXio3Nc/edit?slide=id.g3e87592a3c7_0_92#slide=id.g3e87592a3c7_0_92

Storage Considerations

- Basically we're stuck on the storage costs because of the age of our storage.
  - The costs are up for $48/usable TB to about $100-$110/usable TB for SAS HDD.
    - We need to consider buying 28 disks per R760xd2 storage server instead of the usual 24 disks but we need to make sure that this does not cause significant reduction in transfer speeds.
  - When buying from Dell, 20 TB HDDs give the lowest cost per TB ignoring the infrastructure costs of need more servers than if we bought the servers with 24 TB HDDs.

CPU Considerations

- CPU server cost is driven by memory and solid state storage costs.
  - AI demand has driven memory cost up by a factor of 3-4.
  - Andre Sciaba found that CPU prices are flat.
- Kevin Casella at BNL points out that because of these costs, we should turn of hyper-threading and buy servers with half as much memory and storage.
  - Hyper-Threading gains at most 33% more production from a CPU and if we pay 4 times as much for the memory/storage this does not make economic sense.
    - In making this calculation, one must account for the costs of the additional infrastructure needed: racks, PDUs, and network ports.
- Nominally we want 3 GB of memory and 30 GB of storage per job slot. These are soft requirements.

Common Tools

- Do we have fabric commonality to make common tool development worthwhile?