WLCG Update

Ian Bird

LHCC Referee's meeting

CERN, 26th February 2019

Final update - Data - 2018

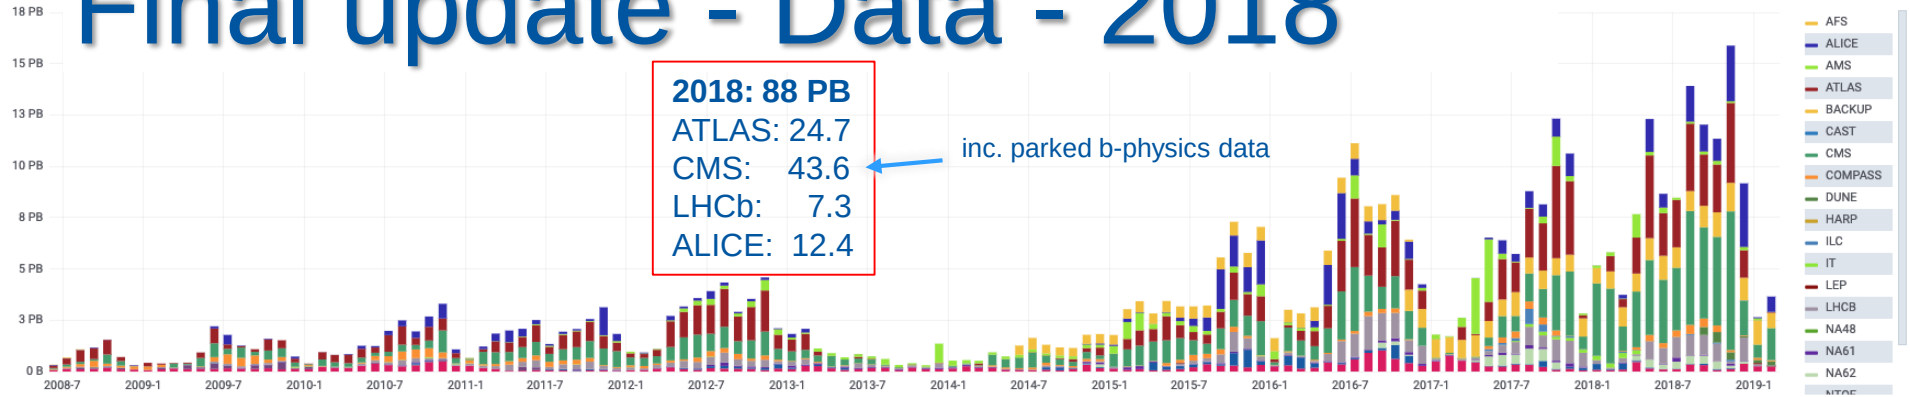

2018: 88 PB
ATLAS: 24.7
CMS: 43.6
LHCb: 7.3
ALICE: 12.4

inc. parked b-physics data

Data transfers

HI Run

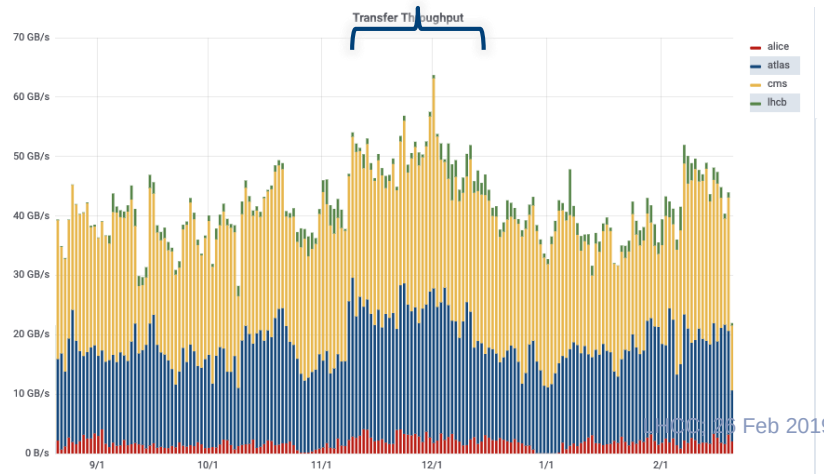

2018: 19.8 PB
ATLAS: 5.2
CMS: 7.7
LHCb: 1.2
ALICE: 5.7

HI Run

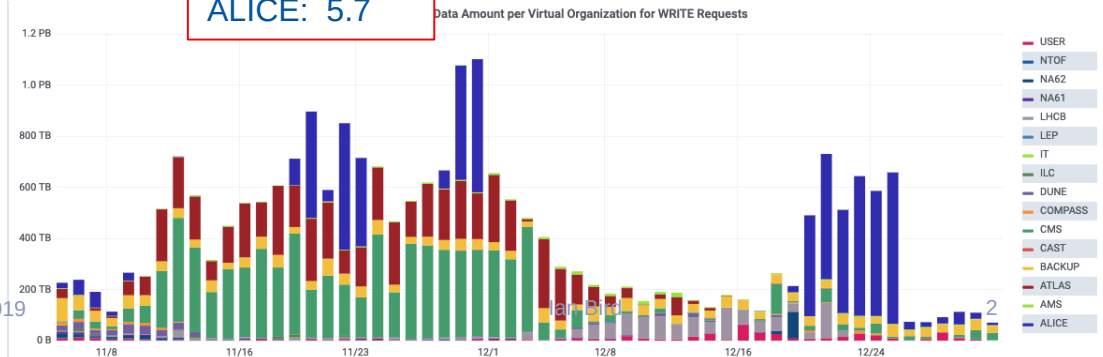

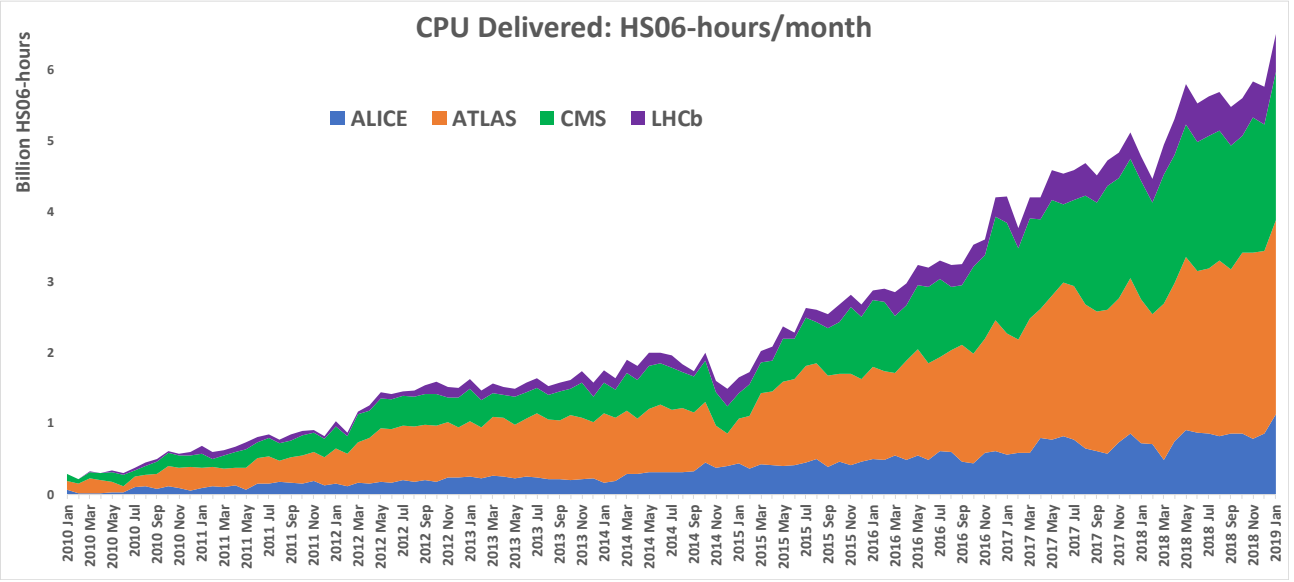

CPU Delivered

New peak: ~271 M HS06-days/month

~ 875 k cores continuous

(From sites that pledge)

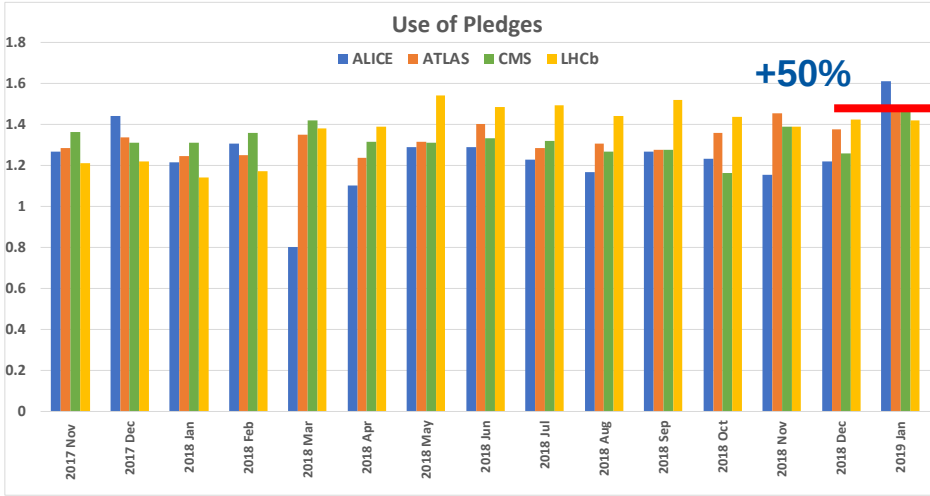

2019 Pledge situation – no change

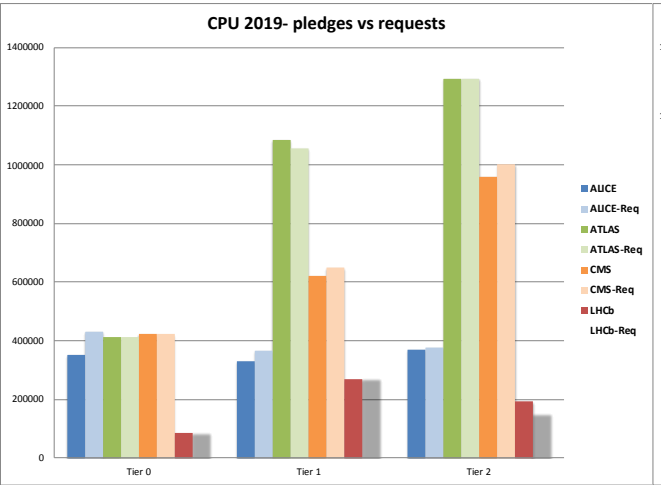

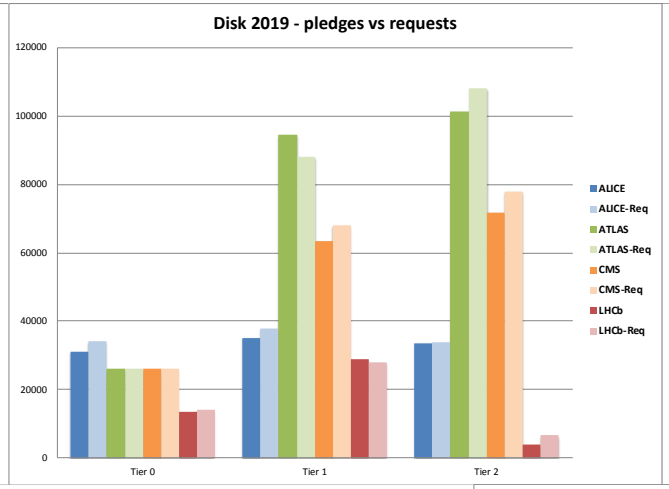

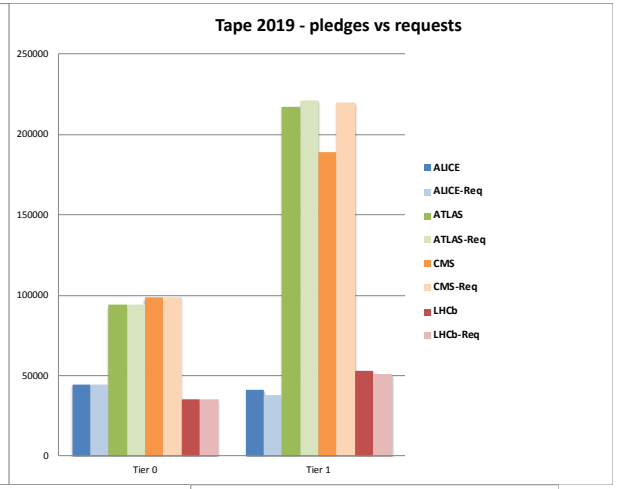

2019 pledges wrt requests:

As given in REBUS

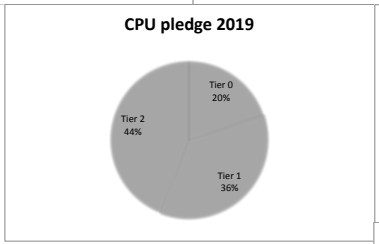

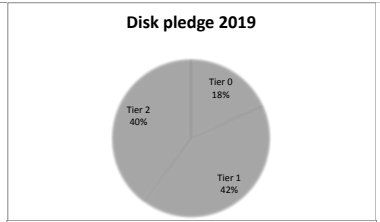

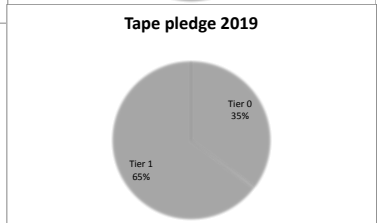

R&D projects - DOMA

- Data access – work is ongoing – next main status update at JLAB workshop
  - 3 groups – Access, 3rd party copy, QoS
- Access – propose a strawman model for how a data lake works for analysis use case
  - To aid thinking about workflows and how it will be used

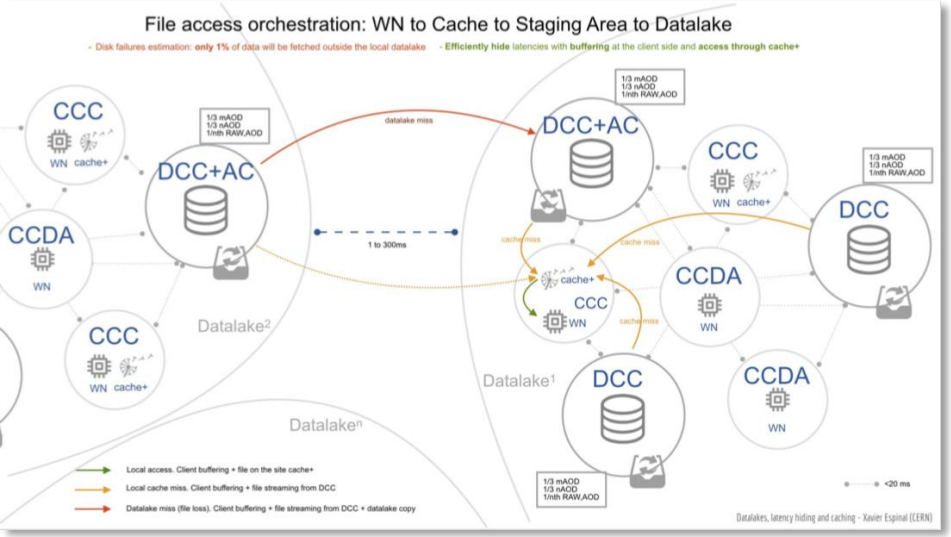

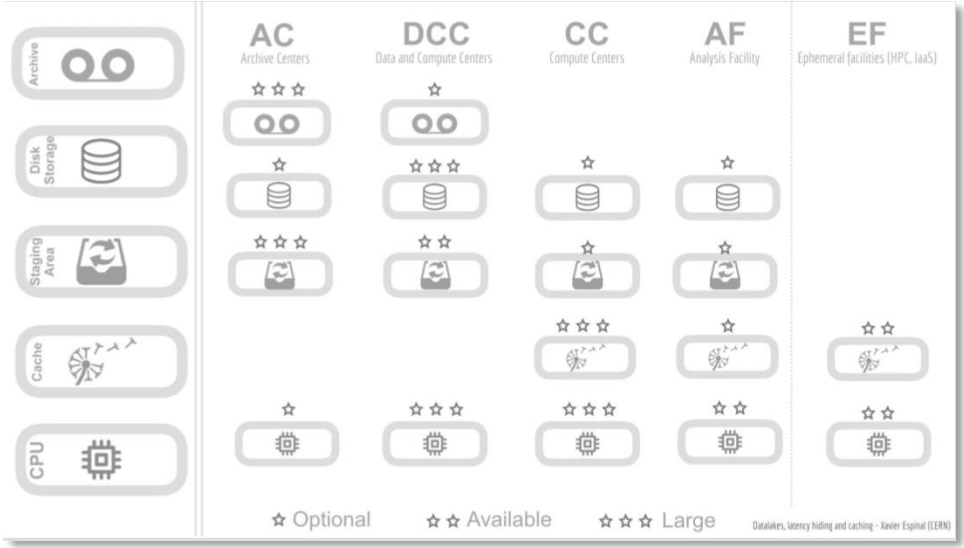

DOMA – 2

- Rucio now seems like a real potential common high-level DM tool
  - Rucio workshop this week
  - ATLAS+CMS, DUNE, SKA are interested, will be part of ESCAPE prototype
- Want to initiate discussions with national efforts to understand how they see storage evolving over the next 5 years
  - How will archive storage look, how much consolidation and/or federation will there be, etc?
  - Storage costs stagnating – potential to lost “cheap” tapes, future models may be very different
  - All goes in the direction of needing a storage lake that serves data (streaming, caching, etc), rather than data distribution and management to expensive storage systems
- Networking activity
  - At protocol level, SDN, bandwidth on demand, etc.
  - Together with SKA w.g.
- WG on storage archiving
  - Ongoing work
    - ATLAS testing on tape carousels
    - Investigations of xcache – proposing test deployments

Cost evolution – new concerns

- Concern raised by several countries – looking at planning for next years
  - The “20%” per year capacity improvement/$ does not seem realistic
- Ongoing investigations – Cost Model working group looking at real costs where possible and trends
  - Until now the outlook has been based on CERN figures tracked over last 20 years
  - We have remarked that “Moore’s law” is no longer a reality
    - But indications are that actual cost trends in several countries are significantly different from that of CERN
- In addition, the dominating factors are market drivers, not technology
- We enter a regime of unpredictable pricing, availability and evolution
  - Forecasting costs over 5-10 years is not possible. We will have to update the outlook ~yearly
  - “Flat budget” does not help understand what capacity may be available!

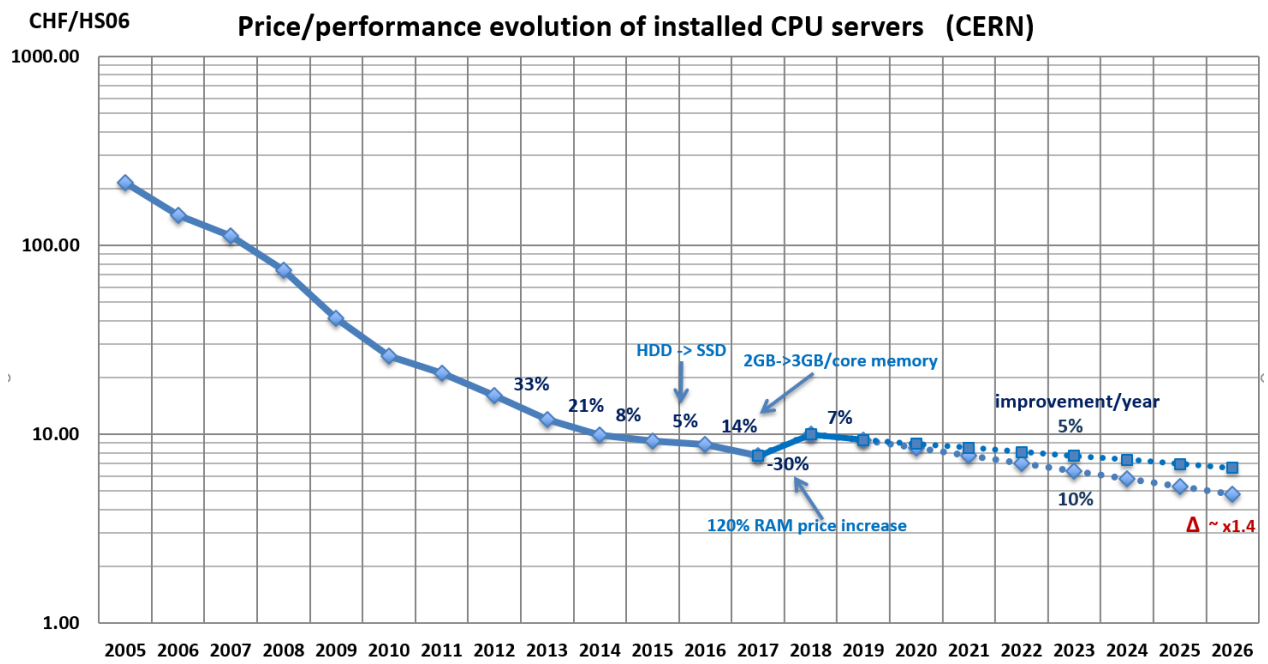

Fluctuating memory price evolution

- Slowdown in feature size reduction in processors (14nm → 10nm → 7nm
  Especially the Intel problems with 10nm in 2018/2019
- Also increase in SSD prices
- Memory and NAND prices are improving in 2019,
  But industry (Samsung, Micron, Hynix,…..) have started to throttle production to stop the price decreases

Samsung DIMM 16GB, DDR4-2133, CL15-15-15, reg ECC

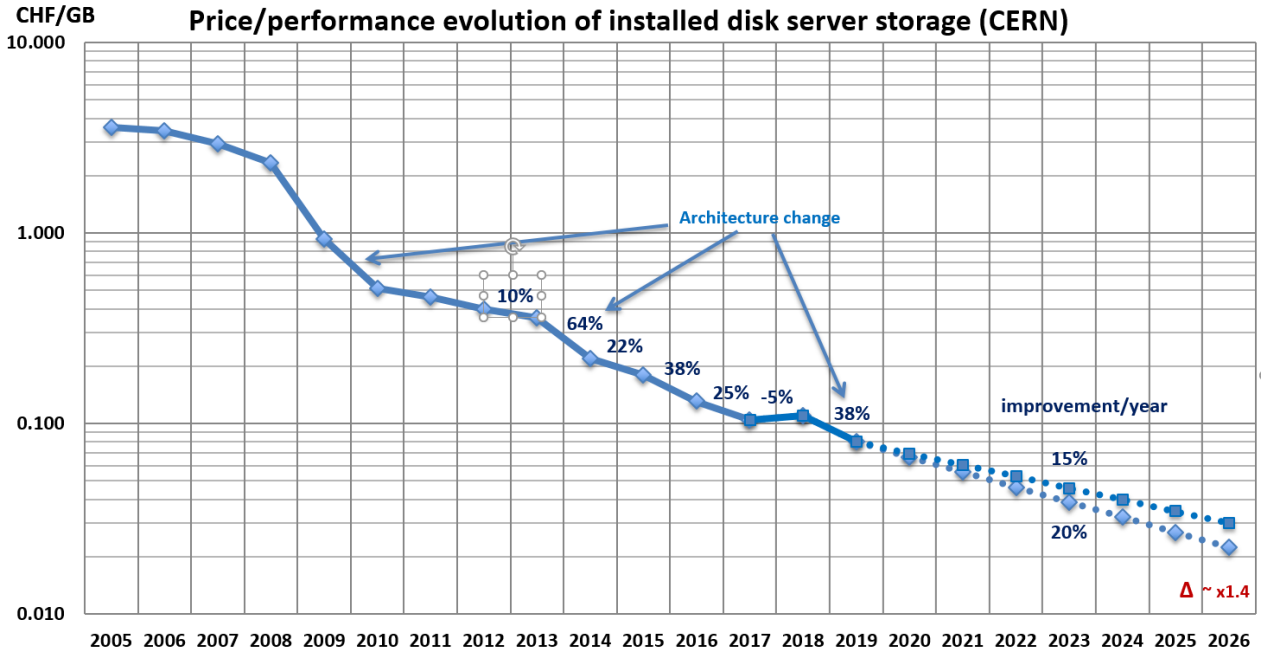

How expensive will the new 16 TB HAMR/MAMR drives be ?

When do we need more spindles then the experiment space requests will require ?
Space ←→ Performance

The total number of disks sold per year is continuously decreasing.

Strongly site depended ! Economy of scale; different server architectures; special one-off pricing; different data management systems

Cost improvements of basic components ← → infrastructure overhead reduction (varying contribution to the yearly improvement rate)
Changes of the disk server architecture at CERN have a made a large impact on the cost decrease:
One-front end node (CPU server) with disk trays attached: went from 1 to 2 to 4 to 8 trays.
Overhead < 10% and much less sensitive to CPU/memory/SSD price changes.
Possible next steps: mirrored → erasure-code; replace only disks and keep trays/front-ends for >8 years.

Summary

- pp run ended with 68 PB on tape,
- HI run ended with ~20 PB on tape
  - Tier 0 performance at exceptional levels
  - No resource contention
- ESCAPE project has started – kick-off meeting was held
  - HL-LHC leads WP on data-infrastructure
- WLCG/HSF/OSG workshop in March
  - Updates and work on R&D projects
  - Planning for the future
  - Discussion on resource/cost expectations