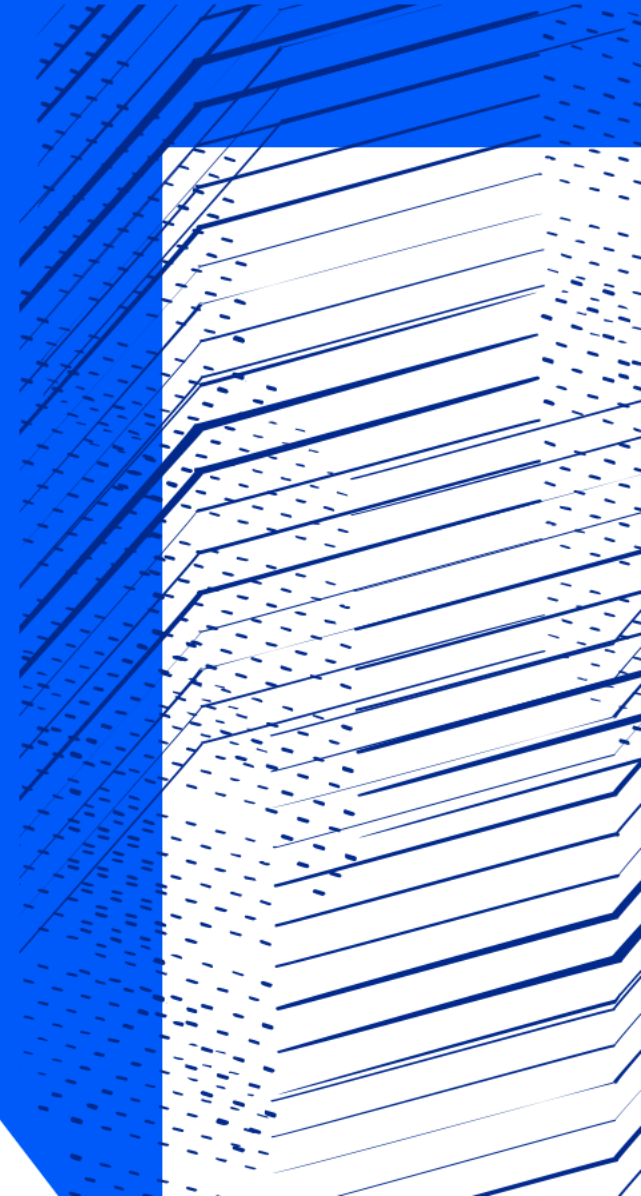




Science and
Technology
Facilities Council

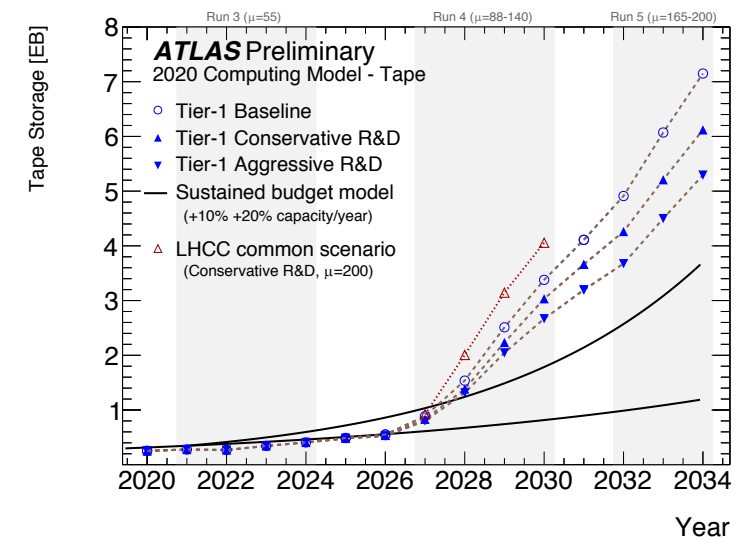
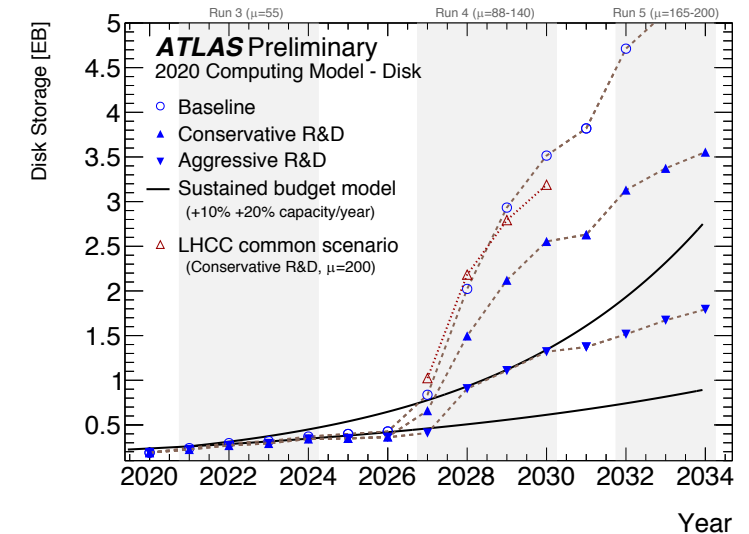
Data Storage

Alastair Dewhurst



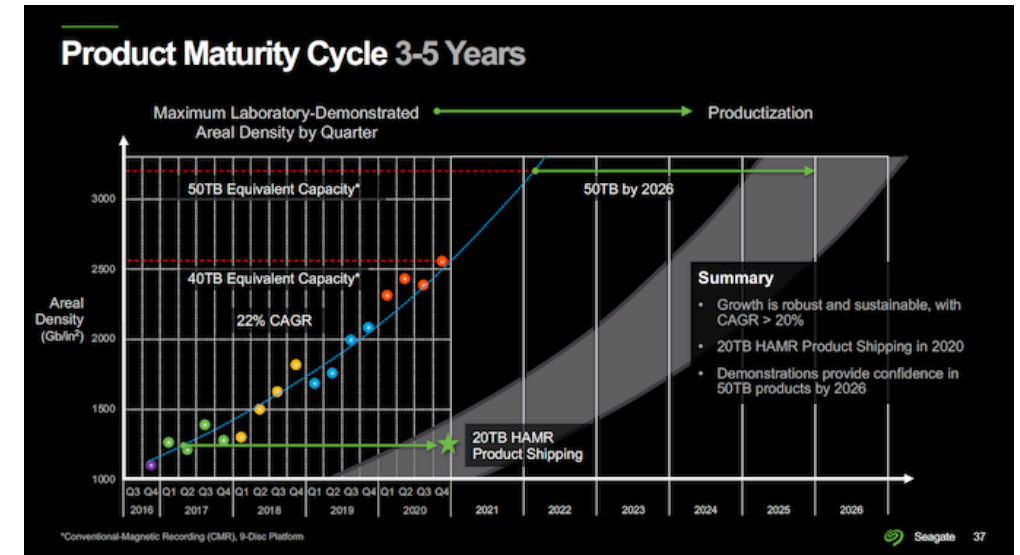
Introduction

- As has almost certainly already been shown, the LHC models predict significant growth in data.
- Even with aggressive R&D this will be at the edge of what Moore's law can provide.
- We can learn a lot from industry.
 - Amazon, Google etc. store many Exabytes already.



Disk Roadmap

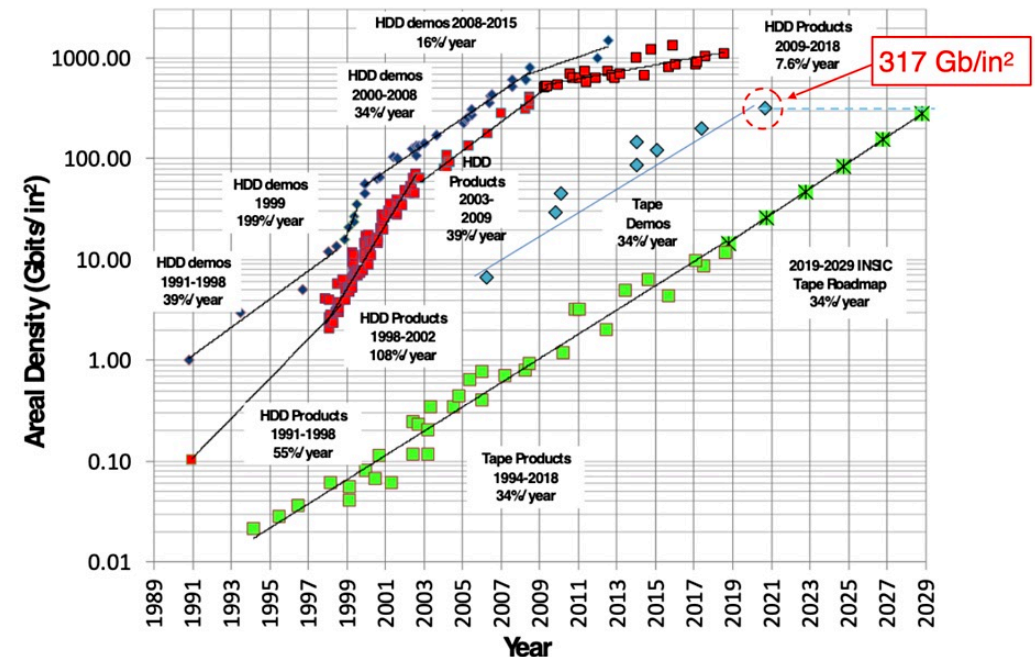
- SSD are taking over in the consumer world and for data intensive workflows.
 - Focus on performance rather than capacity.
- HDD remains critical for data centre use cases.
 - Data does not in general get deleted, so larger fractions are becoming “cold”.
 - Clear roadmap for higher capacity HDD.



Tape Roadmap

- Oracle unexpectedly pulled out of the market in 2017.
- Tape has been declared dead many times, but development continues at a rapid rate.
- Tape has a few strong selling points:
 - No power costs to store data.
 - Tape media lasts a long time (~30 years)
 - New: The air gap means it is immune to ransomware attacks!
- In December 2020, IBM demonstrated a 580TB Tape.

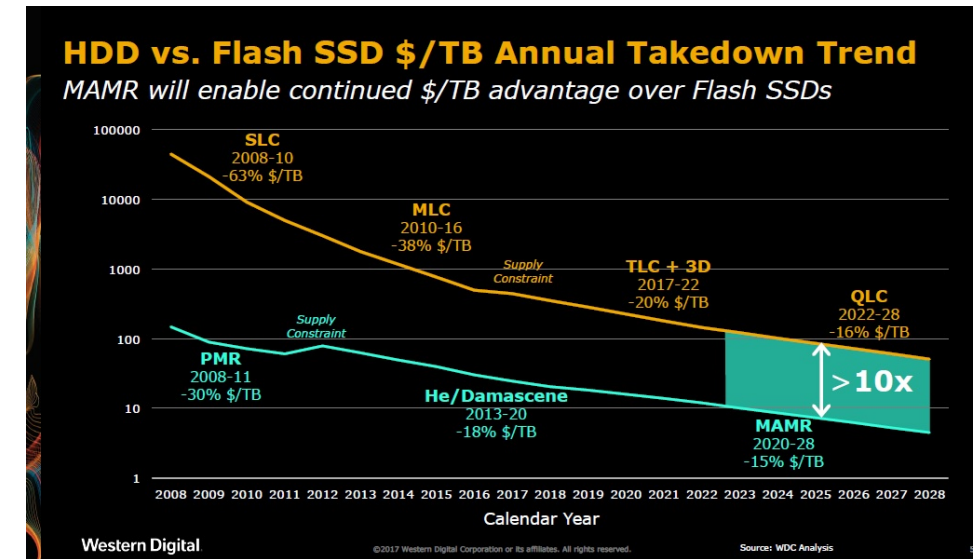
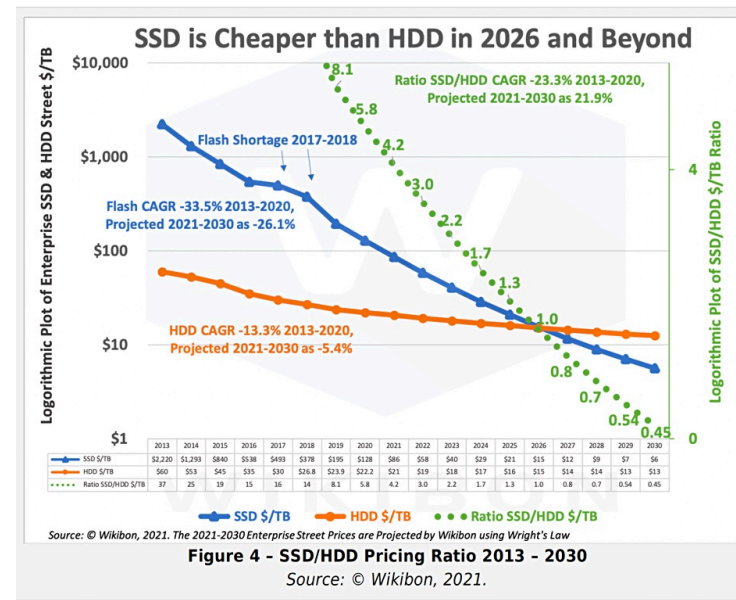
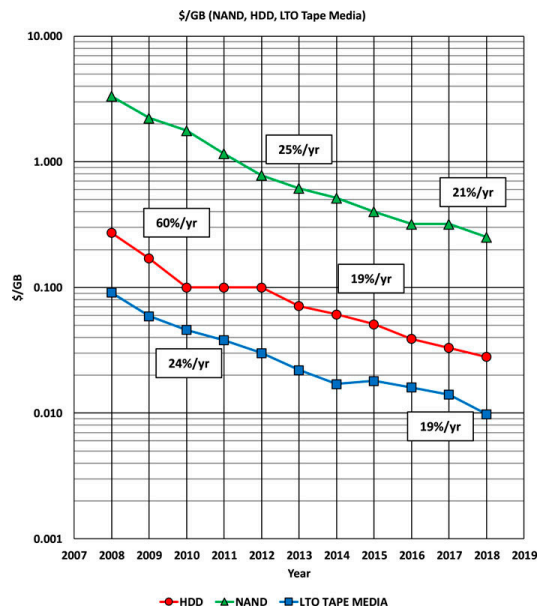
317 Gb/in² demonstrates the sustainability of the INSIC Tape Roadmap
34% CAGR in Areal Density for the next decade



<http://www.insic.org/wp-content/uploads/2019/07/INSIC-Technology-Roadmap-2019.pdf>

Storage costs

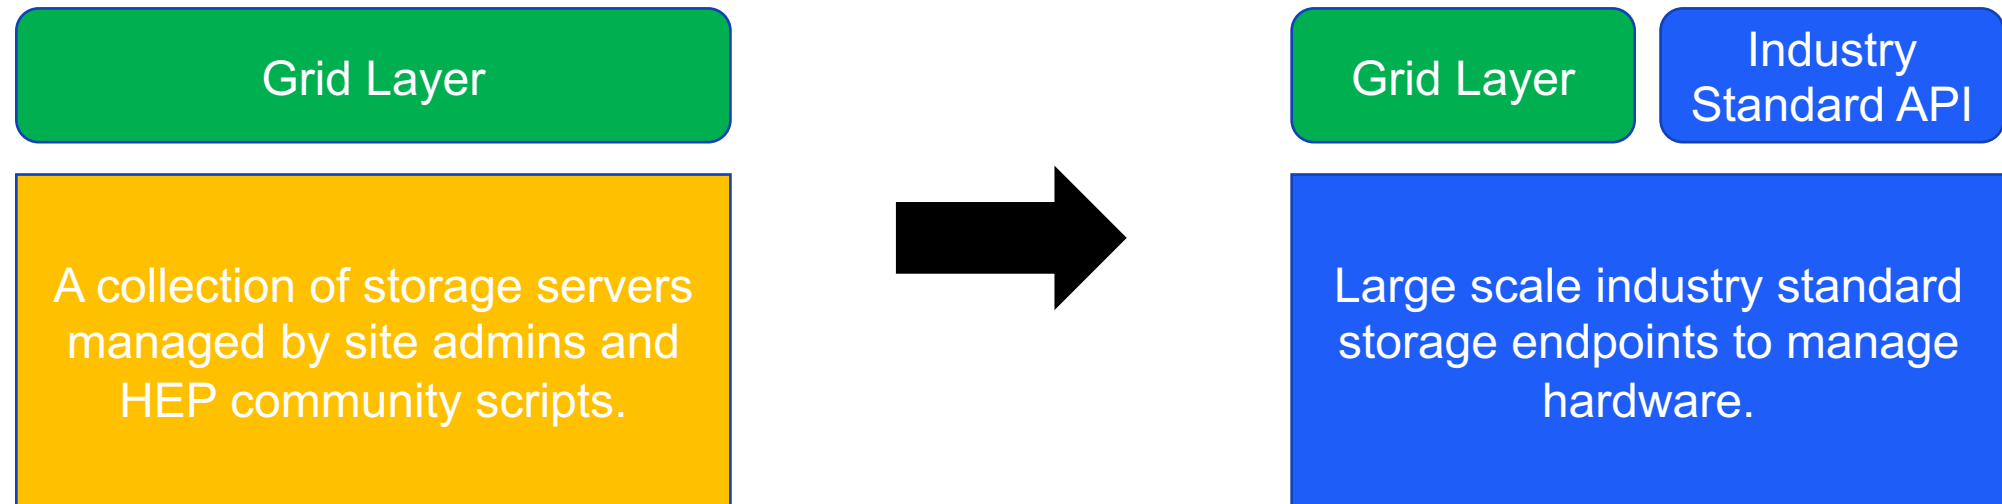
- This year I was quoted ~10 times the cost for SSD compared to HDD storage.
 - Tape is ~1/3 the cost of HDD.
- I believe all 3 technologies will be vital for HL-LHC.



<https://aip.scitation.org/doi/10.1063/1.5130404>

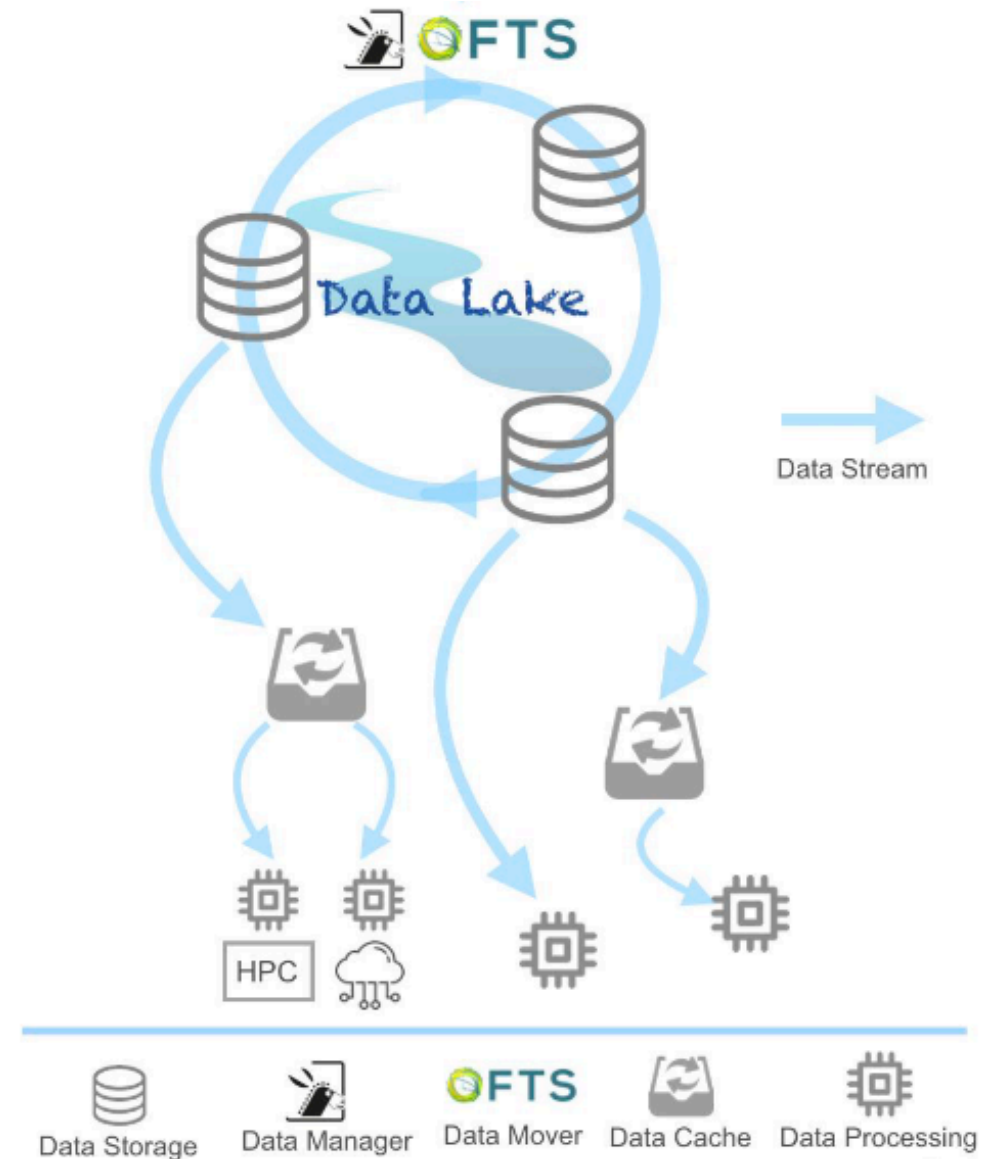
Storage software

- Two things to consider:
 - Managing the underlying storage
 - Managing the middleware
- Grid tools were designed in a different era.
 - Slowly the WLCG is replacing bespoke Grid software with industry standard tools.



Data Lakes

- The WLCG is working on a Data Lake concept.
- I would expect fewer larger scale storage endpoints.
 - Data could be accessed directly from these endpoints.
- Numerous small performant caches (SSD) for data processing.
 - Experiments and Resource providers need to work together to get data in the right place before it is needed.



UK Strengths

- The GridPP project has been in existence for 20 years.
 - Seen as a very reliable partner from an operations point of view.
 - We have many people in leadership positions within the WLCG community.
- UK has pioneered the use of Erasure Coding for data storage.
- UK has additional funding from EGI and Swift-HEP to run / develop FTS, Rucio and CVMFS services.



Science and
Technology
Facilities Council

The background features a large blue triangle pointing downwards, which is partially obscured by a series of overlapping, stylized blue lines that create a sense of depth and movement. The overall color scheme is orange and blue.

Questions?