



# Costs & Sustainability

[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)



International Collaboration for Data Preservation  
and

Long Term Analysis in High Energy Physics

# Introduction

- Some of what I will present is a very simple “cost model” for “bit preservation” for LHC data
  - An Excel sheet using publically available pricing information / technology predictions
- There are some important “hidden details”
  - Building and operating very large data stores is far from trivial: we have a lot of experience (including “horror tales”) plus an outlook to the 5-10EB region (2030s)
- There are significant cost savings in doing this a small number of times
  - Wide-area access to data over “appropriate” network links has been “proven” in the WLCG context
- But there is still significant resistance
  - “Punched card” versus “cloud storage” syndrome

# Background

- Around 20 years ago, we started investigating a range of (component) solutions to LHC computing
- Complemented by “MONARC” model of late 1990s
  - You have to make technology choices on what is available
  - revise later when possible
- Hierarchical model: 622mbps as “worst-case”
- Led to “massive” data migration >10 years ago
- Some 200TB of data: change of medium, change of data format; re-write of “user code”
- Led to on-going “repack & verification” operations

# Before we get lost...

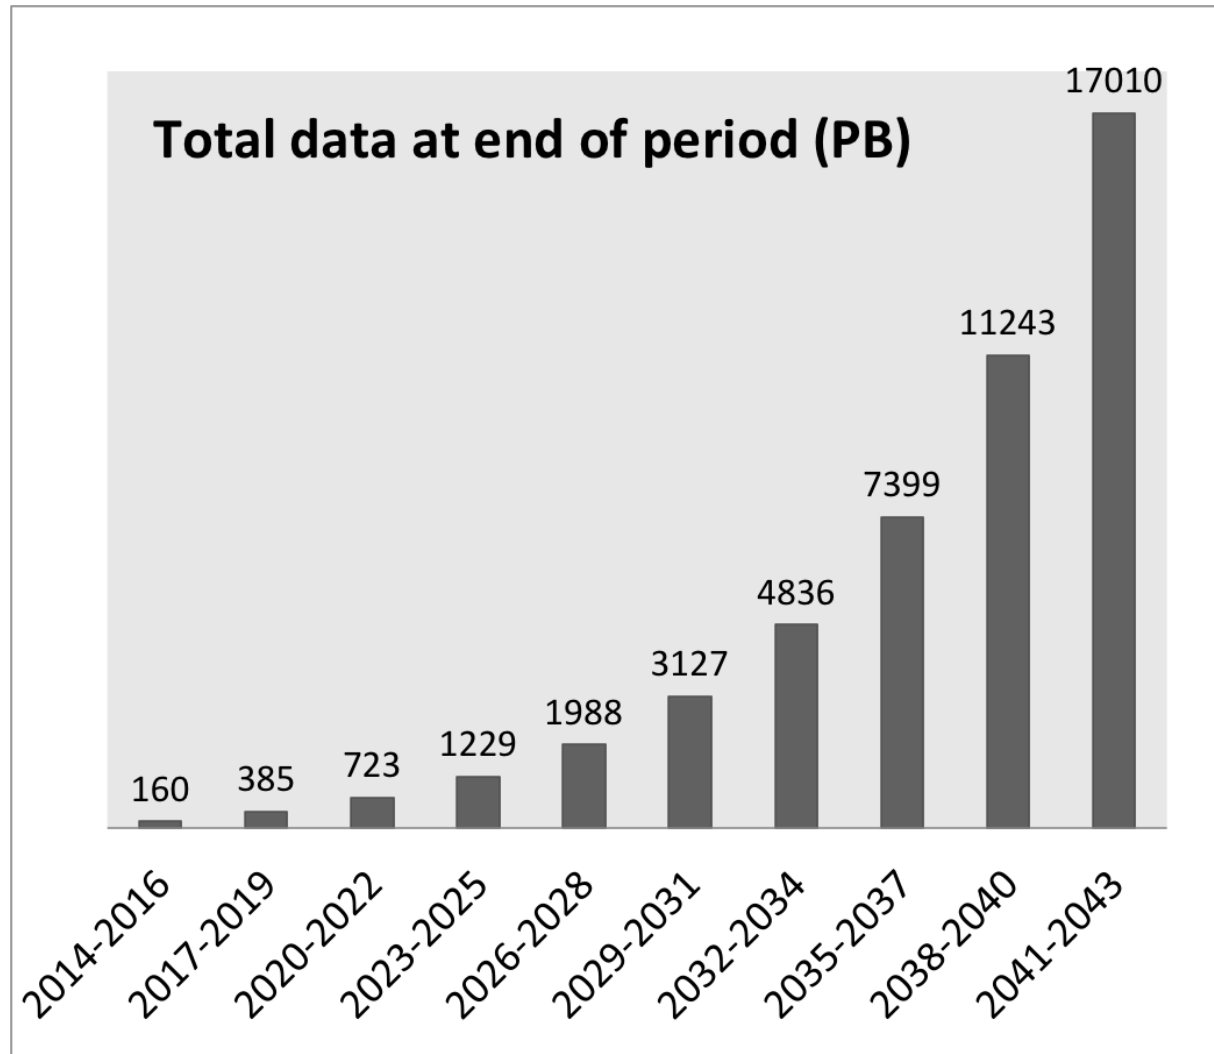
- IMHO, cost of services could be calculated as a “mutuel” – numerous benefits of scale
- “Competitive advantage” for those who opt in
- We do know costs of WLCG globally – with sustainable funding model for ~decades
- Attempts to break these down into individual services also exist

What is the approximate cost of a data archive over 10, 20 and 30 years?

- Generic archive (as on any HEP site)
- Start from scratch in terms of HW/media, with some initial data to be added
- Consider hardware, media, maintenance and electrical power costs
- 3 base scenarios
  - a) 10 PB initially, growing @ 50PB / year
  - b) 10 PB initially, growing @ 50PB +15% / year
  - c) 100 PB initially, no further data (“stable large archive preservation”)

# Case B) increasing archive growth


Start with 10PB, then +50PB/year, then +50% every 3y (or +15% / year)




# Case B) increasing archive growth



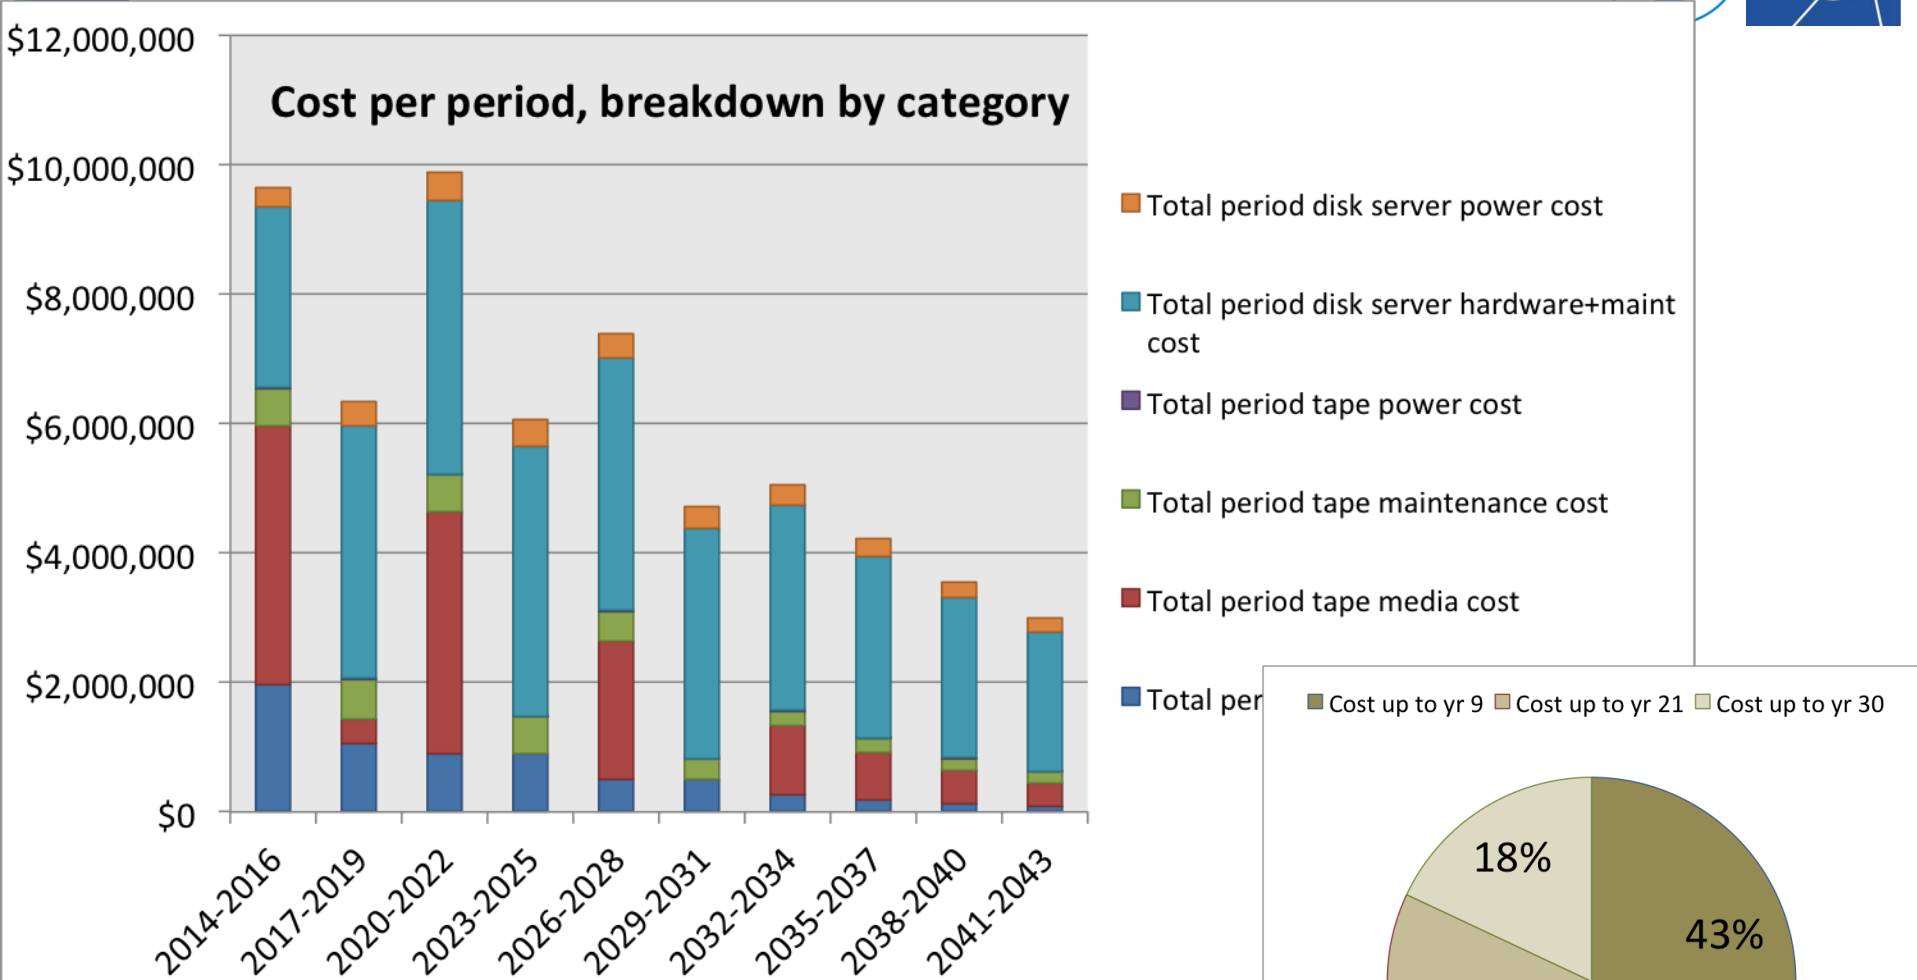
### Total number of disks



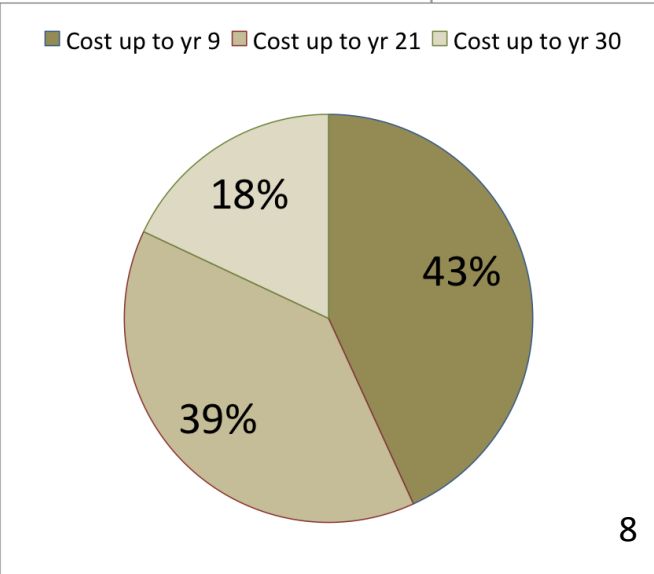
### Total cartridges by end of generation



# Case B) increasing archive growth

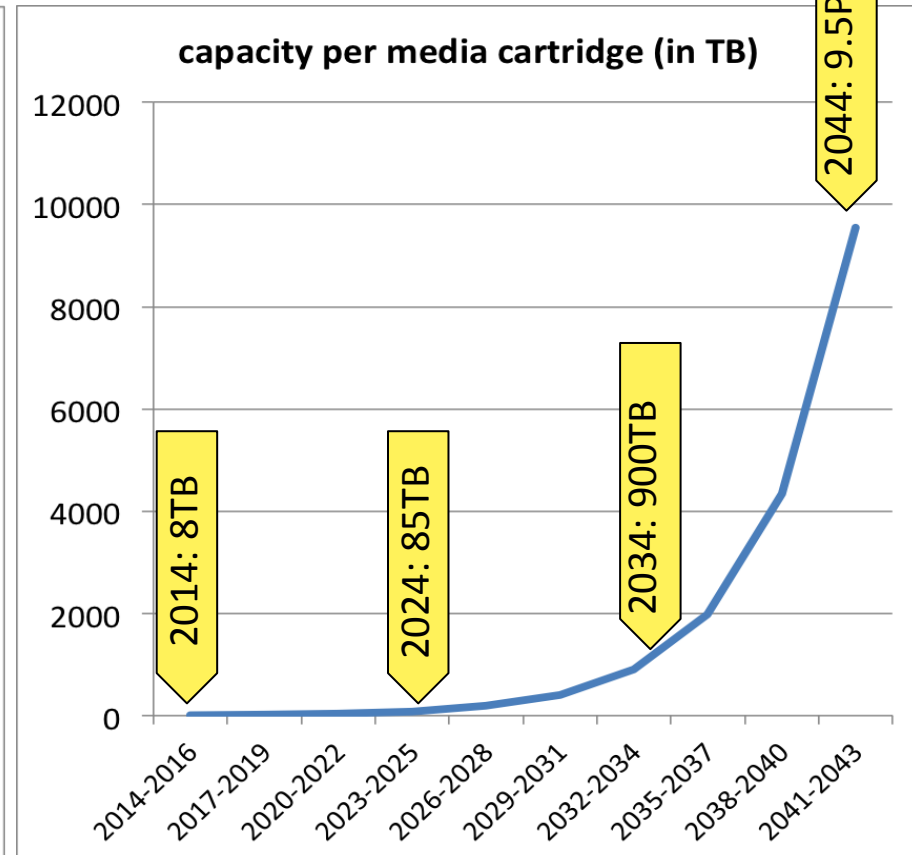
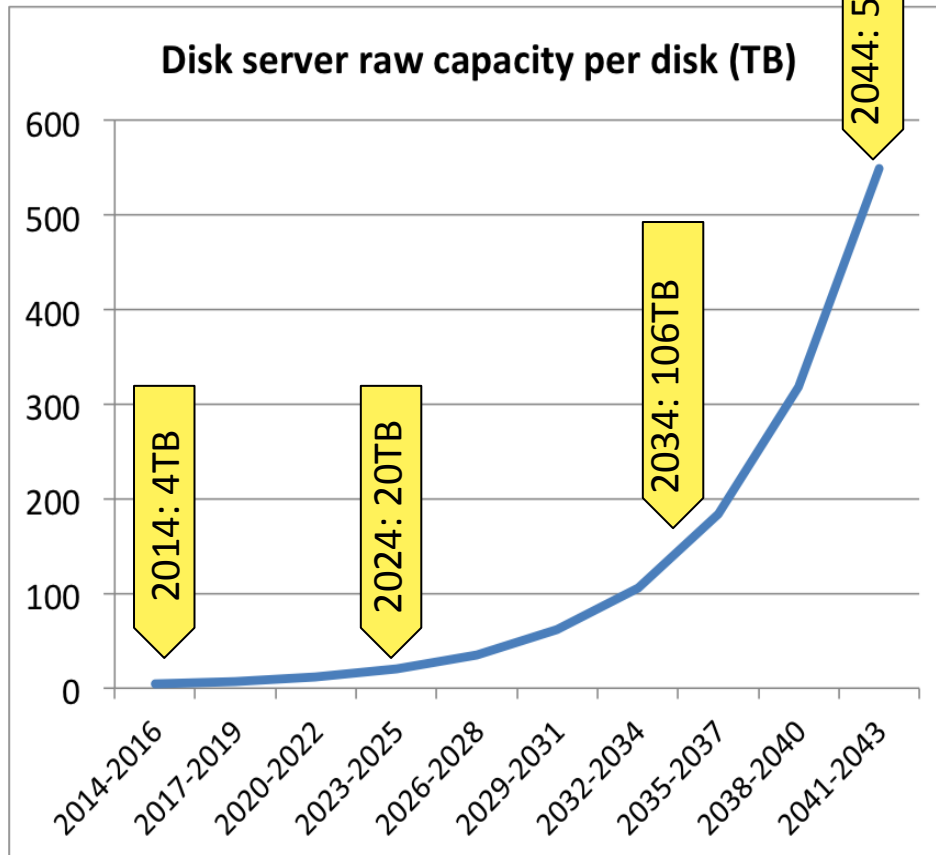


**Total cost: ~\$60M  
(~\$2M/year)**



# Technology evolution

- Assuming
  - +20% yearly disk capacity per constant \$
  - +30% yearly tape capacity per constant \$ (+20%/yr I/O increase)



# Examples

- LEP:  $\sim 100\text{TB} = O(10)$  today's cartridges
- HERA:  $\sim 10\text{PB} = O(10)$  "2030" cartridges
- LHC Run 1:  $\sim 100\text{PB} = O(10)$  "2040" cartridges
- ***LHC total:  $\sim 10\text{EB} = O(??) \text{????}$***