

LHCb Computing

Offline resources in
Run 3 and Run 4

*Concezio Bozzi
CERN and INFN Ferrara*

*9th LHCb Computing Workshop
May 18th 2017*

- Rather than rate (kHz), the amount of resources needed for offline is driven by trigger bandwidth (GB/s)
 - Notable exception: CPU needed for Monte Carlo production scales with the number of collected events, therefore rate
- In the following, 5GB/s is assumed
- Let's be brave: 5GB/s are 100% TURBO
 - At an event size of 50kB the event rate would be 100kHz
- In practice, a (small) fraction of the bandwidth will also be reserved for the FULL stream
 - Monitoring, data mining, scouting, etc.
 - Assume 100kB per event
- Show also the effect of TURBO / FULL at 75:25, 50:50 and 25:75 (~Run 2)
 - No effect on storage for real data (data volume is given by bandwidth \times LHC live time)
 - Small effect on storage for MC (less events to simulate)


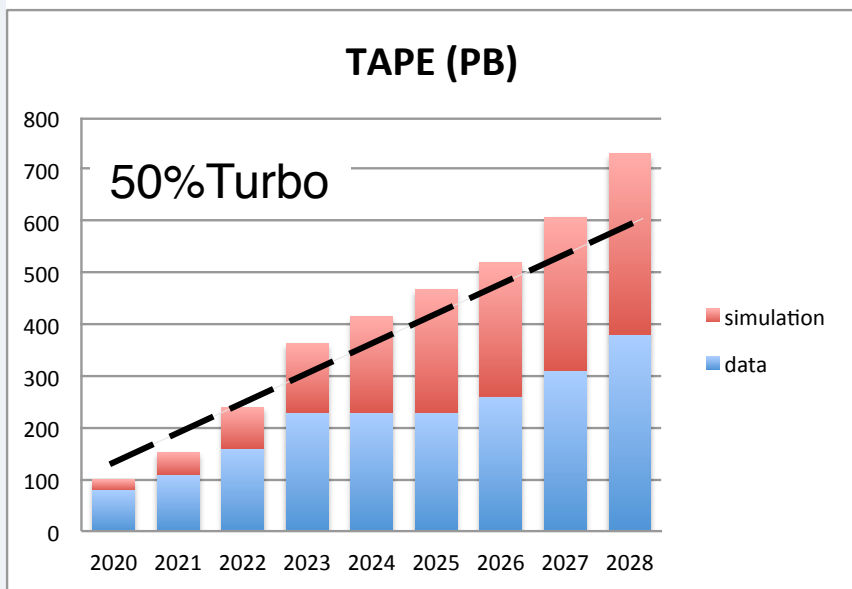


Monte Carlo simulation





- CPU work is dominated by simulation
- In Run1, the amount of simulated events was ~15% that of collected events in data
- What is a reasonable guess for Monte Carlo requests for Run3 and beyond?
 - “every event is interesting” → simulate 100% of real data events
- Being linked to number of collected events, the needed CPU resources depend on TURBO / FULL proportion

- The amount of CPU needed depends on the level of detail of the simulation
 - Several “flavours” of simulation are possible, ranging from FULL Geant 4 (V) to fully parameterised (e.g. DELPHES)
 - Assume that full simulation in Run3 will be faster than the current one by a factor 2
 - Assume a cocktail of “fast” MC is giving a factor 10 less in CPU work with respect to fully simulated MC
- Evaluate the impact on CPU resources when setting the fraction of fastly-simulated to fully-simulated events to 0, 50% and 100%
 - Big impact expected
- Simulation is producing mDST only
- Event size 100kB





Other assumptions

- Live time profile as in Run1
 - 3 / 5 / 7 10^6 seconds in 1st / 2nd / 3rd year of data taking
- TURBO data will never be reconstructed again
- CPU for analysis ~ CPU for stripping in Run2
- 2 disk copies for real data, 1 copy for simulated data
- 2 tape copies for TURBO (and RAW if any), 1 copy for simulation
- Any computing activities related to Run1 and Run2 data?
 - Keep storage space for these data (tape 100PB, disk 60PB)
 - No more stripping after “legacy” stripping of 2019
 - Some level of MC production should be kept...













CPU @ 100% fast sim



CPU @ 50% fast sim



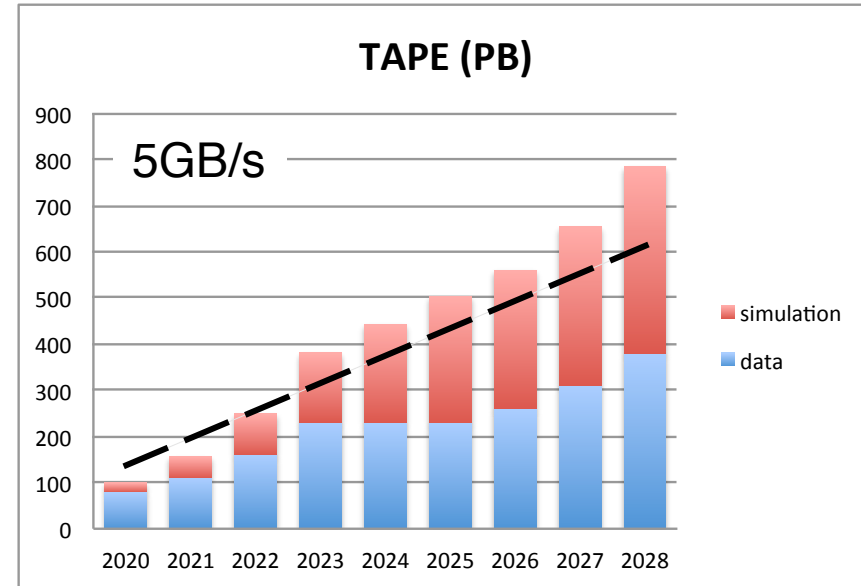


CPU @ 0% fast sim






Change to 2GB/s and 10GB/s: TAPE

- TURBO:FULL = 75:25
- 50% fast sim






Change to 2GB/s and 10GB/s: DISK

- TURBO:FULL = 75:25
- 50% fast sim



Change to 2GB/s and 10GB/s: CPU

- TURBO:FULL = 75:25
- 50% fast sim



When comparing to “likely to be pledged”

- TAPE OK for 2GB/s and 5GB/s, not OK for 10GB/s
- DISK OK for 2GB/s, not OK otherwise
 - Park data?
 - one data replica instead of two ?
 - Give up MC replicas !?
- CPU ~OK even at ~50% fast sim

- “likely to be pledged” == 20% increase per year in DISK and CPU, 25% tape
 - Disk increase is optimistic