# Summary of inputs related to projecting computing resources for future facilities

John Chapman
University of Cambridge

Stephane Jezequel
Laboratoire d'Annecy de Physique des Particules (LAPP)

**David Lange**
Princeton University

Thank you to all those that took the time to address questions we had.
Those discussions were essential for putting together this contribution

# Goals for this talk

- High level (and incomplete) summary and comparison of computing needs for potential future collider experiments at CERN
- Computing related themes identified in contributions to ESPP
- Perspective

# Scale of future experiment resources:
# No updates from CMS/Atlas HL-LHC upgrade requirements

CMS CPU projections (2022)

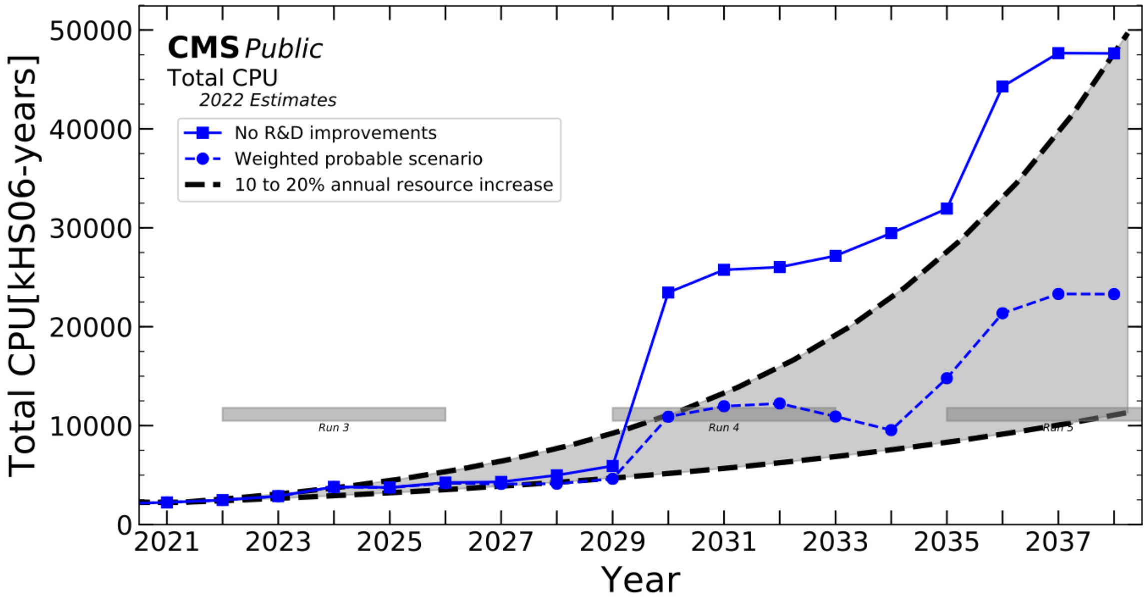


https://cds.cern.ch/record/2815292?ln=en

Atlas CPU projections (2022)

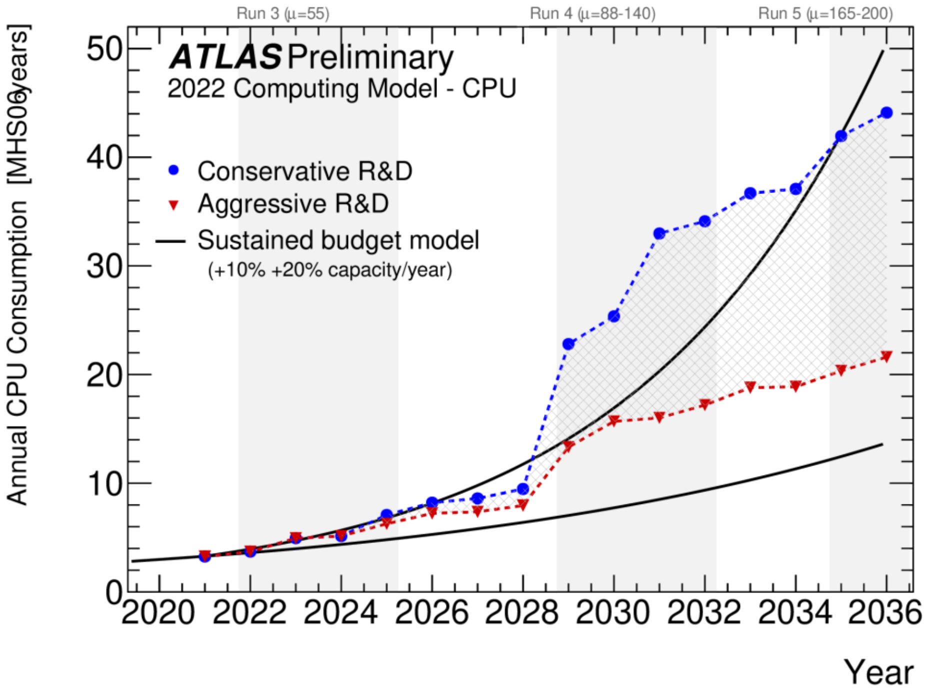


http://cds.cern.ch/record/2802918

# Scale of future experiment resources: FCC-ee

| Run | Process | Number of events | DELPHES Storage (PB) | DELPHES CPU (HS06) | FULL Storage (PB) | FULL CPU (HS06) |
|---|---|---|---|---|---|---|
| Z | $q\bar{q}$ | 1500 G | 12.5 | 2.2 k | 1650 | 2 M |
| | $\ell^+\ell^-$ | 225 G | 0.275 | 12 | 40 | 40 k |
| W | $W^+W^-$ | 60 M | $\sim 10^{-3}$ | | 0.075 | 72 |
| HZ | HZ | 500 k | $\sim 10^{-5}$ | | $\sim 10^{-3}$ | $\sim 1$ |
| | VBFH | 16 k | $\sim 10^{-6}$ | | $\lesssim 10^{-3}$ | |
| top | $t\bar{t}$ | 500 k | $\sim 10^{-5}$ | | $\sim 10^{-2}$ | $\sim 1$ |
| | HZ | 90 k | $\sim 10^{-6}$ | | $\lesssim 10^{-3}$ | |
| | VBFH | 23 k | $\sim 10^{-6}$ | | $\lesssim 10^{-3}$ | |
| Total | | 1725 G | **13** | 2.2 | **1690** | **2 M** |

https://arxiv.org/abs/2505.00272

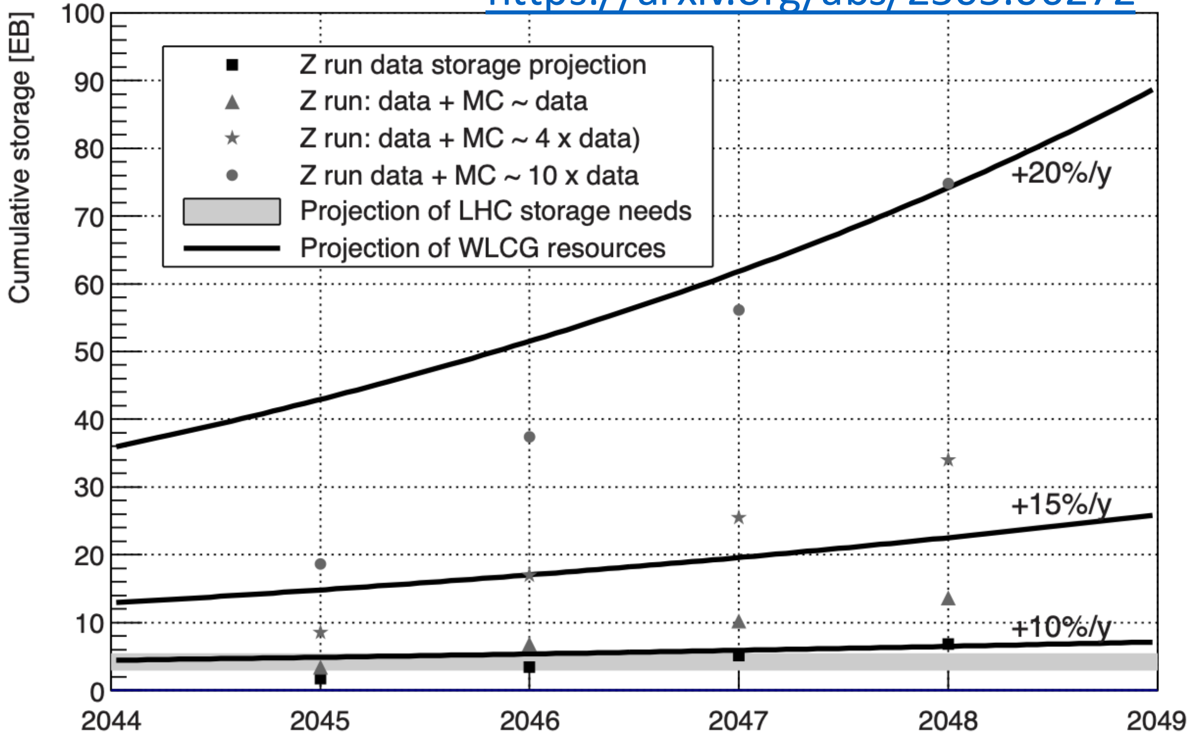


- Considerations/Challenges
  - Computing requirements driven by statistics of giga-Z run. Means that computational challenges come at the start of the physics program
  - Affordable data vs MC ratio and trade-off of storing detailed simulation results (1MB/event) vs recomputing them if needed

# Scale of future experiment resources: LCF4CERN

https://arxiv.org/pdf/2503.24049

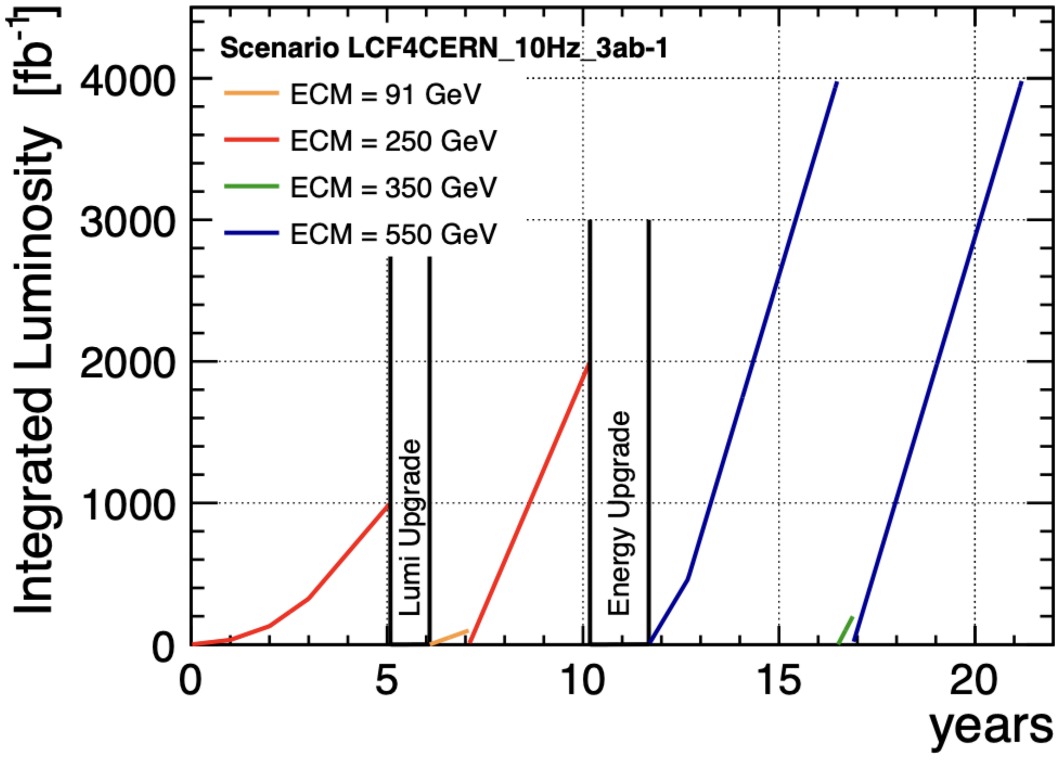


| CM Energy | 250 | 350 | 500 | 250Up | 500Up | GeV |
|---|---|---|---|---|---|---|
| Annual Int. Lumi. | 120 | 80 | 144 | 240 | 288 | $fb^{-1}$ |
| Nb. of Signal + Bhabha / BX | 0.53% | 0.47% | 0.83% | 0.53% | 0.83% | |
| Raw data size/sec | 1106 | 624 | 892 | 2213 | 1784 | MB/sec |
| Data size at Campus per ILD year | | | | | | |
| Raw Data (RD) | 8.85 | 4.99 | 7.13 | 17.71 | 14.27 | PB |
| On-line Processed Data (OPD) | 0.36 | 0.20 | 0.29 | 0.72 | 0.58 | id. |
| Off-line Reconstructed Data (ORD) | 0.54 | 0.30 | 0.43 | 1.07 | 0.87 | id. |
| MC Data (Sim+REC+DST) | 1.27 | 0.74 | 1.74 | 2.55 | 3.48 | id. |
| Sub Total | 11.02 | 6.24 | 9.59 | 22.05 | 19.19 | id. |
| Data size global per ILD year | | | | | | |
| Raw Data (RD) | 17.70 | 9.99 | 14.26 | 35.41 | 28.54 | PB |
| Online Processed Data (OPD) | 3.59 | 2.03 | 2.90 | 7.19 | 5.79 | id. |
| Offline Reconstructed Data (ORD) | 5.36 | 3.03 | 4.32 | 10.73 | 8.65 | id. |
| MC data(Sim+Rec) | 1.27 | 0.74 | 1.74 | 2.55 | 3.48 | id. |
| MC DST with copy | 0.19 | 0.11 | 0.26 | 0.38 | 0.51 | id. |
| Total Data size per ILD year | 26.86 | 15.16 | 21.76 | 53.74 | 43.55 | PB |
| CPU | | | | | | |
| MC production | 658 | 455 | 1259 | 1315 | 2519 | k CPU days |
| Online Processed Data | 1096 | 759 | 2099 | 2192 | 4198 | id. |
| Offline reconstructed (*) | 11 | 7.6 | 21 | 22 | 42 | id. |
| Total CPU days | 1764 | 1222 | 3379 | 3529 | 6759 | id. |

https://zenodo.org/records/4659571

- Considerations/Challenges
  - Comprehensive analysis of requirements for >=250 GeV running (including replication and other factors)
  - Z-pole running estimates extrapolated from FCC-ee estimates

# Scale of future experiment resources: LEP3 and LHeC

- LEP3 considerations/Challenges
  - As with FCC-ee, computing needs of LEP3 are driven by the scale of the Z run
  - Projected based on integrated lumi. relative to FCC-ee in Z-pole running

| | LEP3 | | | | LEP2 |
|---|---|---|---|---|---|
| No. of IPs/Xpts | 2 | | | | 4 |
| com energy √s (GeV). For final state( X) | 230 (ZH) | 160 (WW) | 91.2 (Z) | | 45.60 |
| Circumference/Length [km] | 26.659 | | | | 26.659 |
| Bending Radius (m) | 2958 | | | | 3026 |
| Crossing Angle at IP [mrad] | 30 | | | | 0 |
| SR Energy Loss/turn [GeV) | 5.4 | 1.3 | 0.13 | | 0.1 |
| Total RF Voltage [GV] | 6 | 1.5 | 0.5 | | 0.2 |
| SR Power /beam (MW) | 50 | | | | 0.33 |
| Beam Current [mA] | 9 | 39 | 371 | | 3 |
| Number of bunches/beam | 20 | 220 | 800 | | 8 |
| Bunch Intensity (10E11) | 2.5 | 1 | 2.6 | | 1.8 |
| RF frequency. [MHz] | 800 | | | | 352 |
| Beam Lifetime (Bhabha+Brem) [min] | 17 | 19 | 28 | | 1390 |
| Inst. Luminosity/IP [10^34 cm-2s-1] | 1.6 | 6.2 | 40 | | 0.002 |
| Integrated L/IP [ab-1/yr] | 0.192 | 0.7 | 4.8 | | 0.001 |
| Years of Operation** | 6 | 4 | 5 | | |

https://arxiv.org/pdf/2504.00541

# Scale of future experiment resources: LEP3 and LHeC

- LEP3 considerations/Challenges
  - As with FCC-ee, computing needs of LEP3 are driven by the scale of the Z run
  - Projected based on integrated lumi. relative to FCC-ee in Z-pole running

| | LEP3 | | | | LEP2 |
|---|---|---|---|---|---|
| No. of IPs/Xpts | 2 | | | | 4 |
| com energy √s (GeV). For final state( X) | 230 (ZH) | 160 (WW) | 91.2 (Z) | | 45.60 |
| Circumference/Length [km] | 26.659 | | | | 26.659 |
| Bending Radius (m) | 2958 | | | | 3026 |
| Crossing Angle at IP [mrad] | 30 | | | | 0 |
| SR Energy Loss/turn [GeV) | 5.4 | 1.3 | 0.13 | | 0.1 |
| Total RF Voltage [GV] | 6 | 1.5 | 0.5 | | 0.2 |
| SR Power /beam (MW) | 50 | | | | 0.33 |
| Beam Current [mA] | 9 | 39 | 371 | | 3 |
| Number of bunches/beam | 20 | 220 | 800 | | 8 |
| Bunch Intensity (10E11) | 2.5 | 1 | 2.6 | | 1.8 |
| RF frequency. [MHz] | 800 | | | | 352 |
| Beam Lifetime (Bhabha+Brem) [min] | 17 | 19 | 28 | | 1390 |
| Inst. Luminosity/IP [10^34 cm-2s-1] | 1.6 | 6.2 | 40 | | 0.002 |
| Integrated L/IP [ab-1/yr] | 0.192 | 0.7 | 4.8 | | 0.001 |
| Years of Operation** | 6 | 4 | 5 | | |

- LHeC considerations/Challenges
  - A streaming readout without any hardware triggering is expected to be possible due to the relatively modest event rate.
  - Permanently save only 10kHz, independently of the delivered luminosity
    - Corresponds to $Q^2 > 10\ \mathrm{GeV}^2$ at a luminosity of $10^{33}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$
  - Event size of around 100kB and 100 HS06·s of processing time

https://arxiv.org/pdf/2504.00541

# Scale of future experiment resources: 10 TeV Muon Collider

Compute and data tier level estimates for 5 years running of a 10 TeV Muon collider

| | HEPscore[1] s / event | events | total CPU [kHEPscore[1]-years] |
|---|---|---|---|
| Generation[2] | 640 | | $2.03 \cdot 10^3$ |
| Simulation[3] | 550 | $10^{11}$ | $1.74 \cdot 10^3$ |
| Reconstruction (sim)[4] | 1130 | | $3.58 \cdot 10^3$ |
| Reconstruction (data)[4 5] | | $5 \cdot 10^{12}$ | $1.79 \cdot 10^5$ |
| Re-processing (sim) | 1680 | $2 \cdot 10^{11}$ | $1.07 \cdot 10^4$ |
| Re-processing (data) | 1130 | $2 \cdot 5 \cdot 10^9$ | $7.17 \cdot 10^2$ |
| | | | $\mathbf{1.98 \cdot 10^5}$ |

Assumes full reconstruction of all events in trigger system

| | size [MB] / event | events | total size [PB] |
|---|---|---|---|
| RAW[6] | 80 | | 400 |
| RECO / AOD[7 8] | 20 | $5 \cdot 10^9$ | 100 |
| analysis[9] | 0.005 | | 0.03 |
| SIM[10] | 250 | | 25000 |
| SIM RECO[11] | 40 | $10^{11}$ | 4000 |
| SIM analysis[11] | 0.01 | | 1.0 |
| **Total** | | | **29501** |

Tiers typically stored by LHC experiments

- Considerations/Challenges
  - Controlling and reducing beam induced backgrounds.
  - Modeling these backgrounds is resource intensive
  - Trigger design needed that is adequate for reducing data rate by 1000x
  - Understanding need of retaining detailed simulation information vs. regenerating samples when needed

https://arxiv.org/pdf/2504.21417

# Scale of future experiment resources: FCC-hh

https://cds.cern.ch/record/2842569

Considerations/Challenges

- Rate from high-level trigger estimated to be 100kHz **(~10x HL-LHC)** if keeping LHC trigger thresholds
- Expected pileup 5x HL-LHC at 85 TeV. In terms of total tracks this is like ~1700 overlaid 14 TeV interactions **(~10x HL-LHC)**.

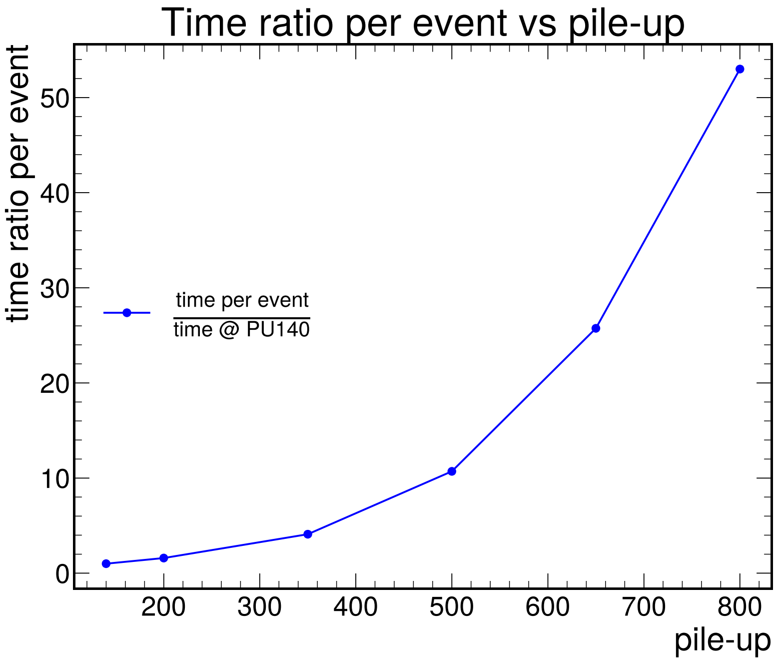


| | LHC | HL-LHC | FCC-hh Initial | FCC-hh Nominal |
|---|---|---|---|---|
| **Physics performance and beam parameters** | | | | |
| Peak luminosity[1] [$10^{34}$ cm$^{-2}$s$^{-1}$] | 1.0 | 5.0 | 5.0 | $<30.0$ |
| Optimum average integrated luminosity / day [fb$^{-1}$] | 0.47 | 2.8 | 2.2 | 8 |
| Target turnaround time [h] | | | 2 | 2 |
| Assumed turnaround time[2] [h] | | | 5 | 4 |
| Peak number of inelastic events / crossing | 27 | 135 levelled | 171 | 1026 |
| Total / inelastic cross section $\sigma$ proton [mbarn] | 111 / 85 | | 153 / 108 | |
| Luminous region RMS length [cm] | | | 5.7 | 5.7 |
| Distance IP to first quadrupole, L* [m] | 23 | | 40 | 40 |
| **Beam parameters** | | | | |
| Number of bunches n | 2808 | | 10400 | |
| Bunch spacing [ns] | 25 | 25 | 25 | |
| Bunch population $N$ [$10^{11}$] | 1.15 | 2.2 | 1.0 | |
| Nominal transverse normalised emittance [µm] | 3.75 | 2.5 | 2.2 | 2.2 |
| Number of IPs contributing to $\Delta Q$ | 3 | 2 | 2+2 | 2 |
| IP beta function [m] | 0.55 | 0.15 (min) | 1.1 | 0.3 |
| RMS IP spot size [µm] | 16.7 | 7.1 (min) | 6.8 | 3.5 |
| Full crossing angle [µrad] | 285 | 590 | 104 | 200[3] |

- The compute requirements scale non-linearly with pileup. New approaches to tracking and other reconstruction algorithms are necessary.
- Gains from timing detectors are potential large. To be better understood as a means of removing hits from pileup interactions through HL-LHC operations

# Summary table of planned and potential general purpose experiments at CERN (with apologies to HL-LHC Run 5 upgrades)

| | HL-LHC (GPE) | LCF4CERN | FCC-ee | LEP3 | LHeC | Muon Coll | FCC-hh |
|---|---|---|---|---|---|---|---|
| Potential start | ~2030 | ~2042 | ~2045 | ~2045 | ~2042 | ~2050 | ~2070 |
| Physics program (Energy: [ab$^{-1}$, Events]) | 14: 3,1e12 | Z: 1,3e10<br>ZH: 2,4e5<br>350: 0.2,<br>550: 8, | Z: 205,6e12<br>WW: 19,2.4e8<br>ZH: 11,2.2e6<br>top: 3,2.5e6 | Z: 48,2e12<br>WW: 5.6,4e7<br>ZH: 2.4,4e5 | 1.2: 1,70B | 10 TeV:<br>10,1e10 | 85: 40,2e13 |
| Length [Yrs] | 12 | 25 | 15 | 20 | 6 | 5 | 25 |
| Total CPU [MHS23-y] | 500 | 0.5 | ~110 | ~30 | 4 | 200 | ~100000 |
| Ave annual CPU [MHS23] | 40 | 0.02 | 7<br>(33/yr for Z) | 1.5<br>(0.5/yr for Z) | 1.5 | 40 | 4000 |
| Est total Storage [EB] (including SIM hit tier) | 20 | 1 | ~8 (40) | ~10 | 0.12 | 0.7 (29) | ~4000<br>(in yr~2100) |

- Luminosity, CPU and storage are for all planned IPs.
- All mistakes and misunderstandings are mine..

# Compute modeling themes and challenges: Importance of moving further into a heterogeneous world

Several considerations for making reliable projections

- Evolving estimates of what components / fraction of applications to be offloaded to an accelerator.
- Different workflow types have different requirements and may lead to different choices of resources in a heterogeneous environment.
  - Analysis: Optimize time to insight
  - Production: Optimize cost per processed event
- Owned/bought vs HPC vs opportunistic resources
- Scale of corresponding data preservation and analysis reinterpretation
- Development of benchmarks and accounting

# Scale of envisioned resource growth (WLCG flat budget estimates)

| | Previous (Circa 2021) | New (June 2025) |
|---|---|---|
| CPU | (15+/-5)% | (10+/-5)% |
| Disk | (15+/-5)% | (5+/-5)% |
| Tape | (15+/-5)% | (10+/-5)% |

Most recent WLCG Disk storage projection

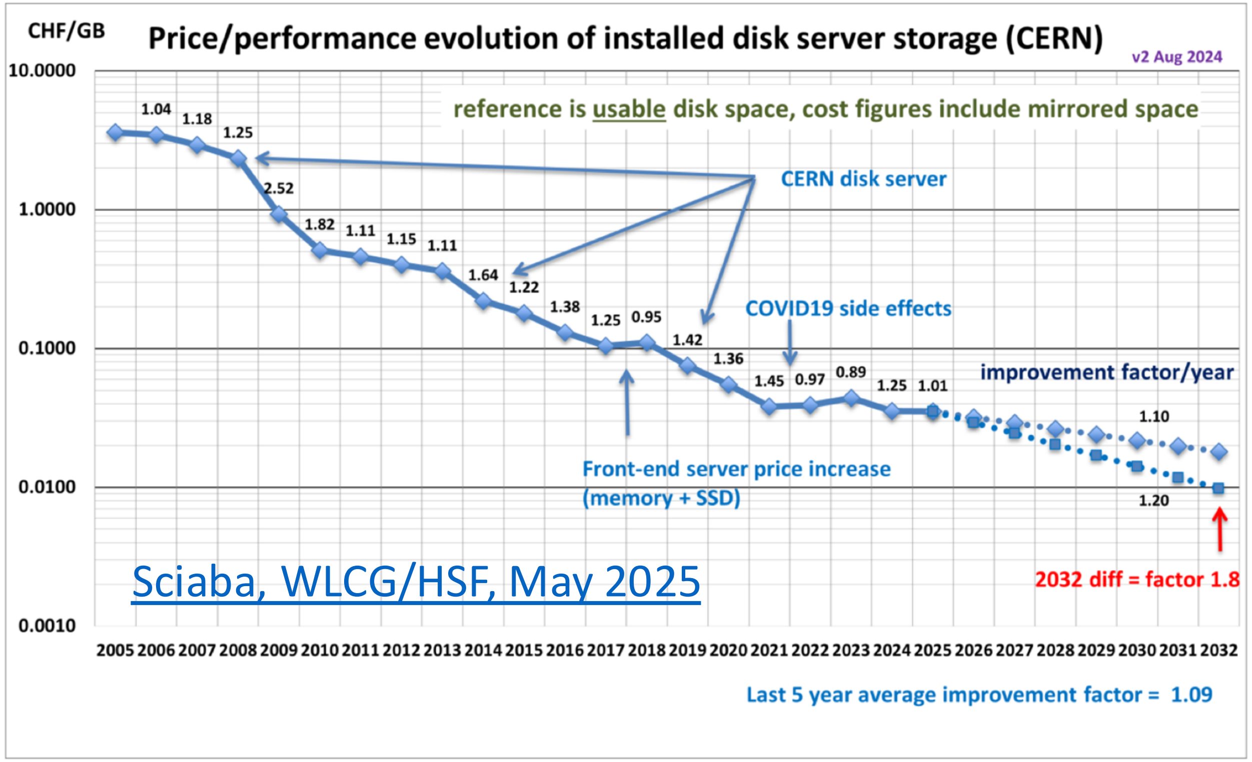


[Sciaba, WLCG/HSF, May 2025]

The storage aspect of future analysis models will need to change considerably if disk costs per unit of usable volume do not continue to reduce over time

# Thinking about projections of experimental feasibility of future experiments over decades

https://indico.cern.ch/event/1052077/

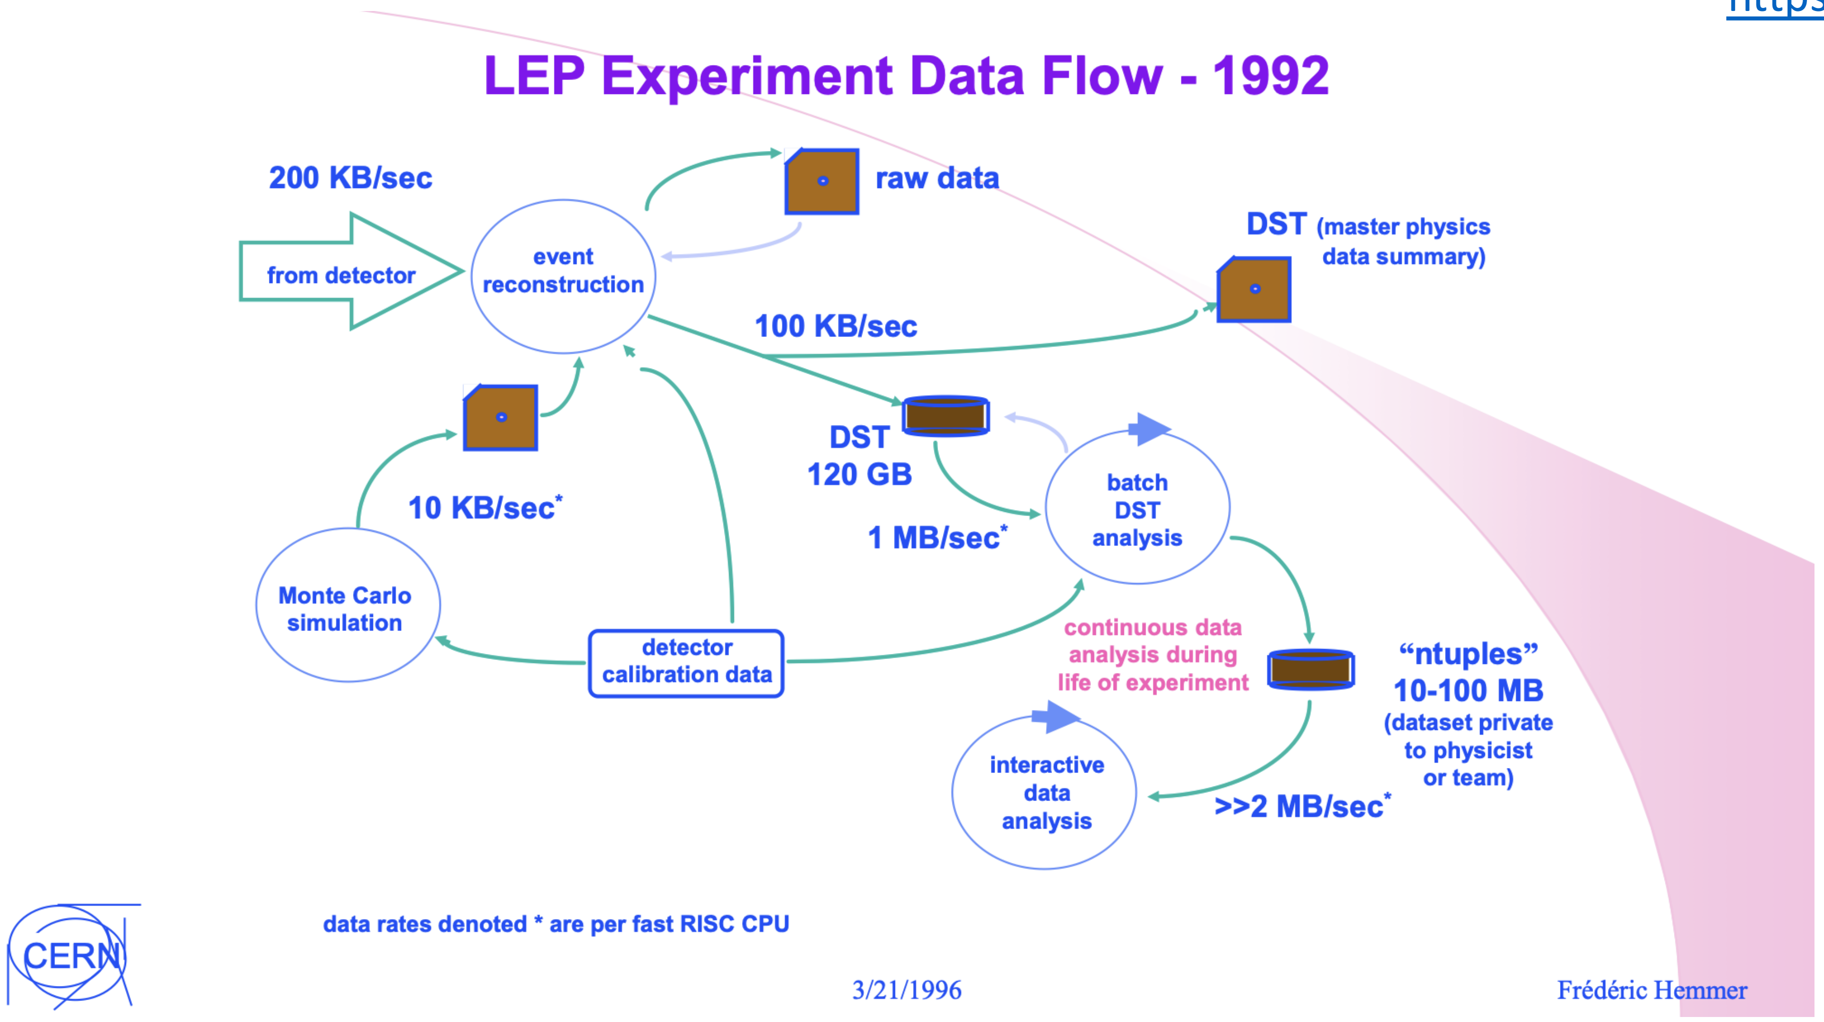

# Thinking about projections of experimental feasibility of future experiments over decades

https://indico.cern.ch/event/1052077/

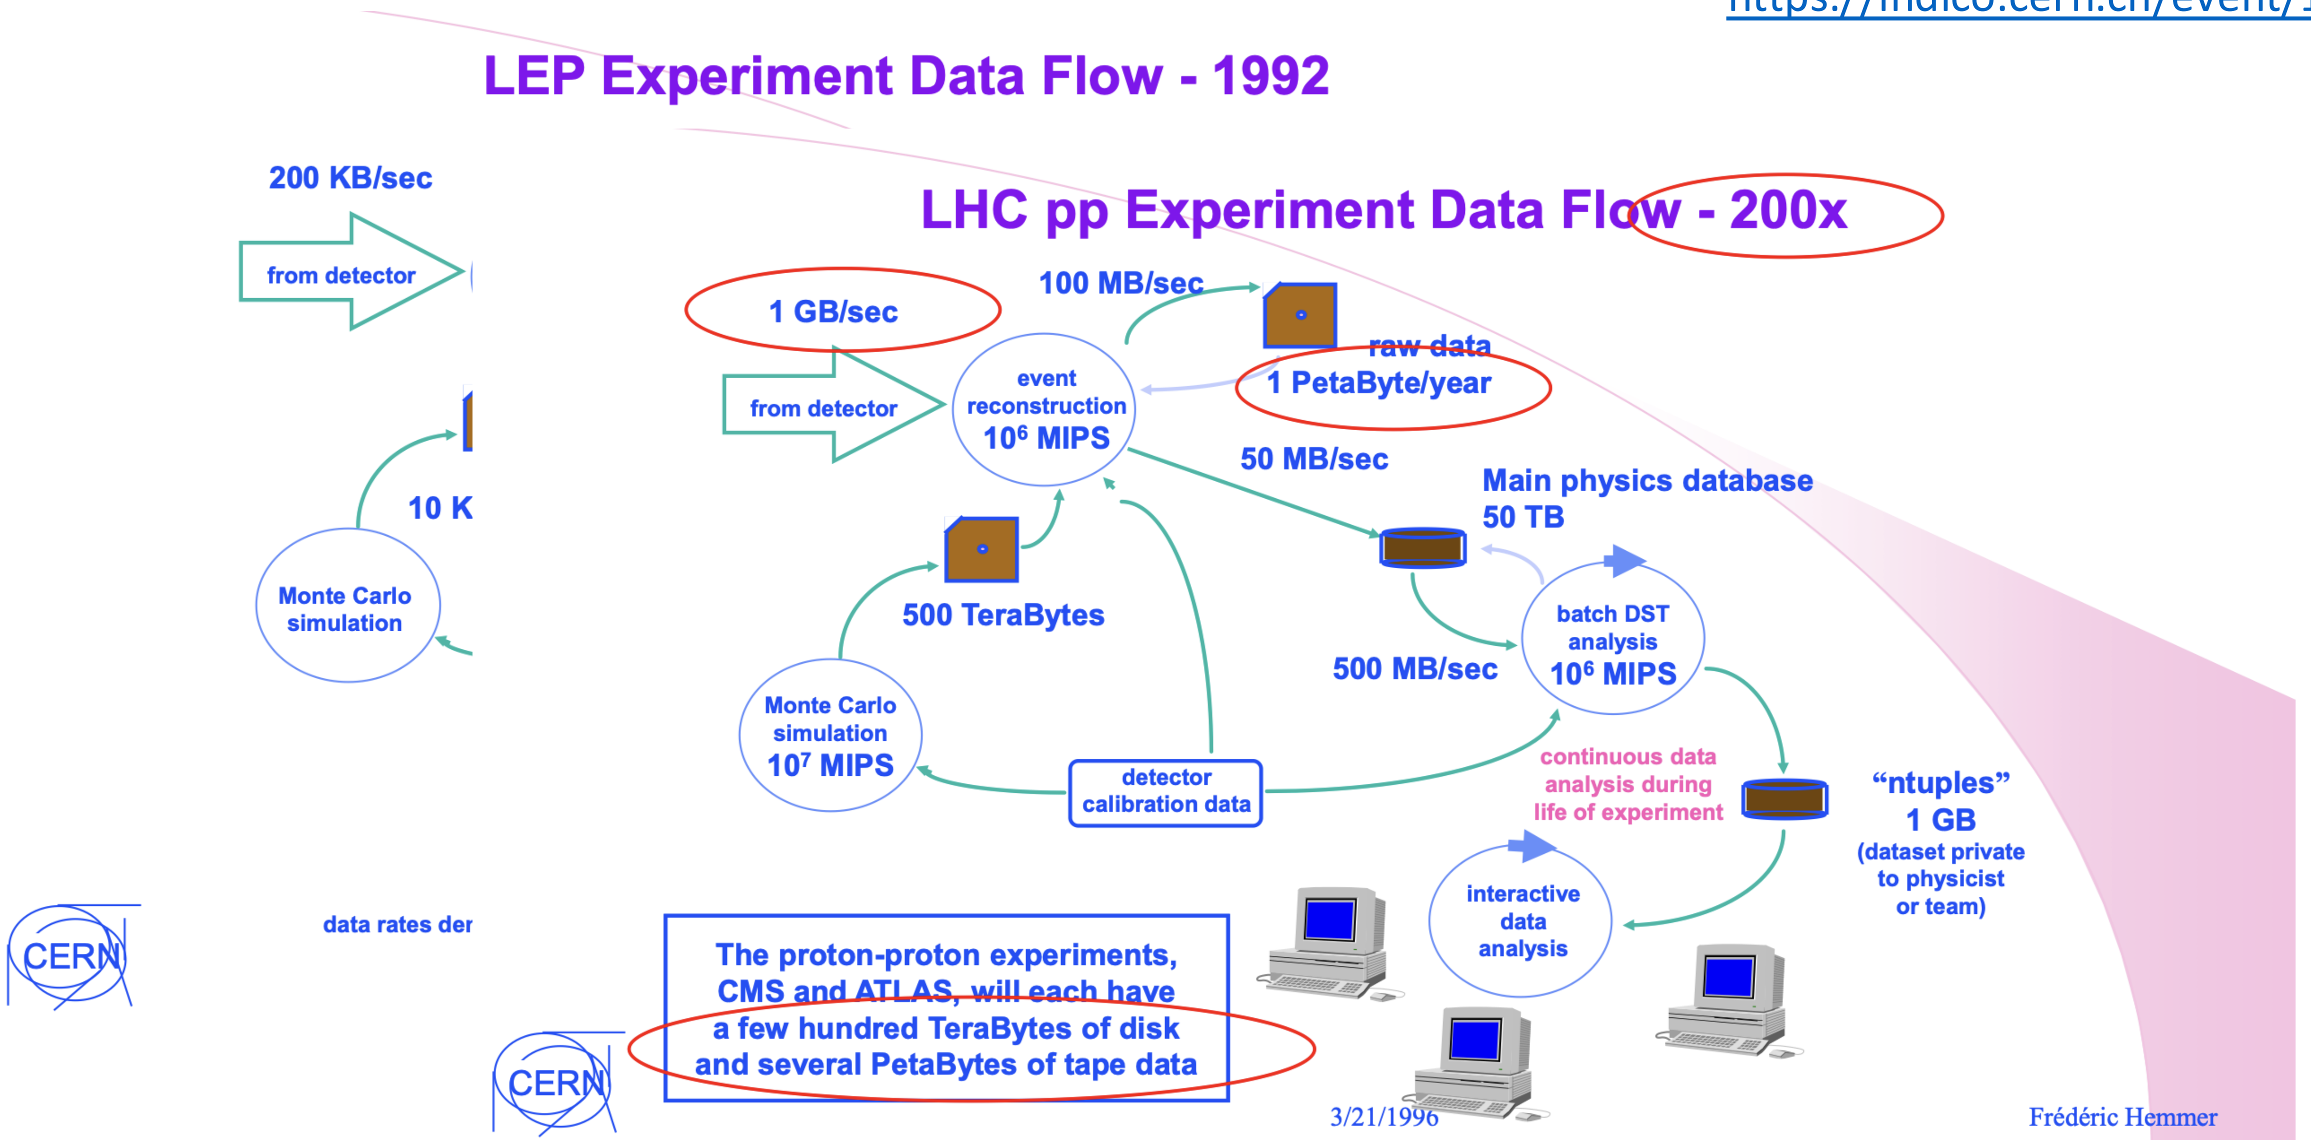

# Thinking about projections of experimental feasibility of future experiments over decades

https://indico.cern.ch/event/1052077/

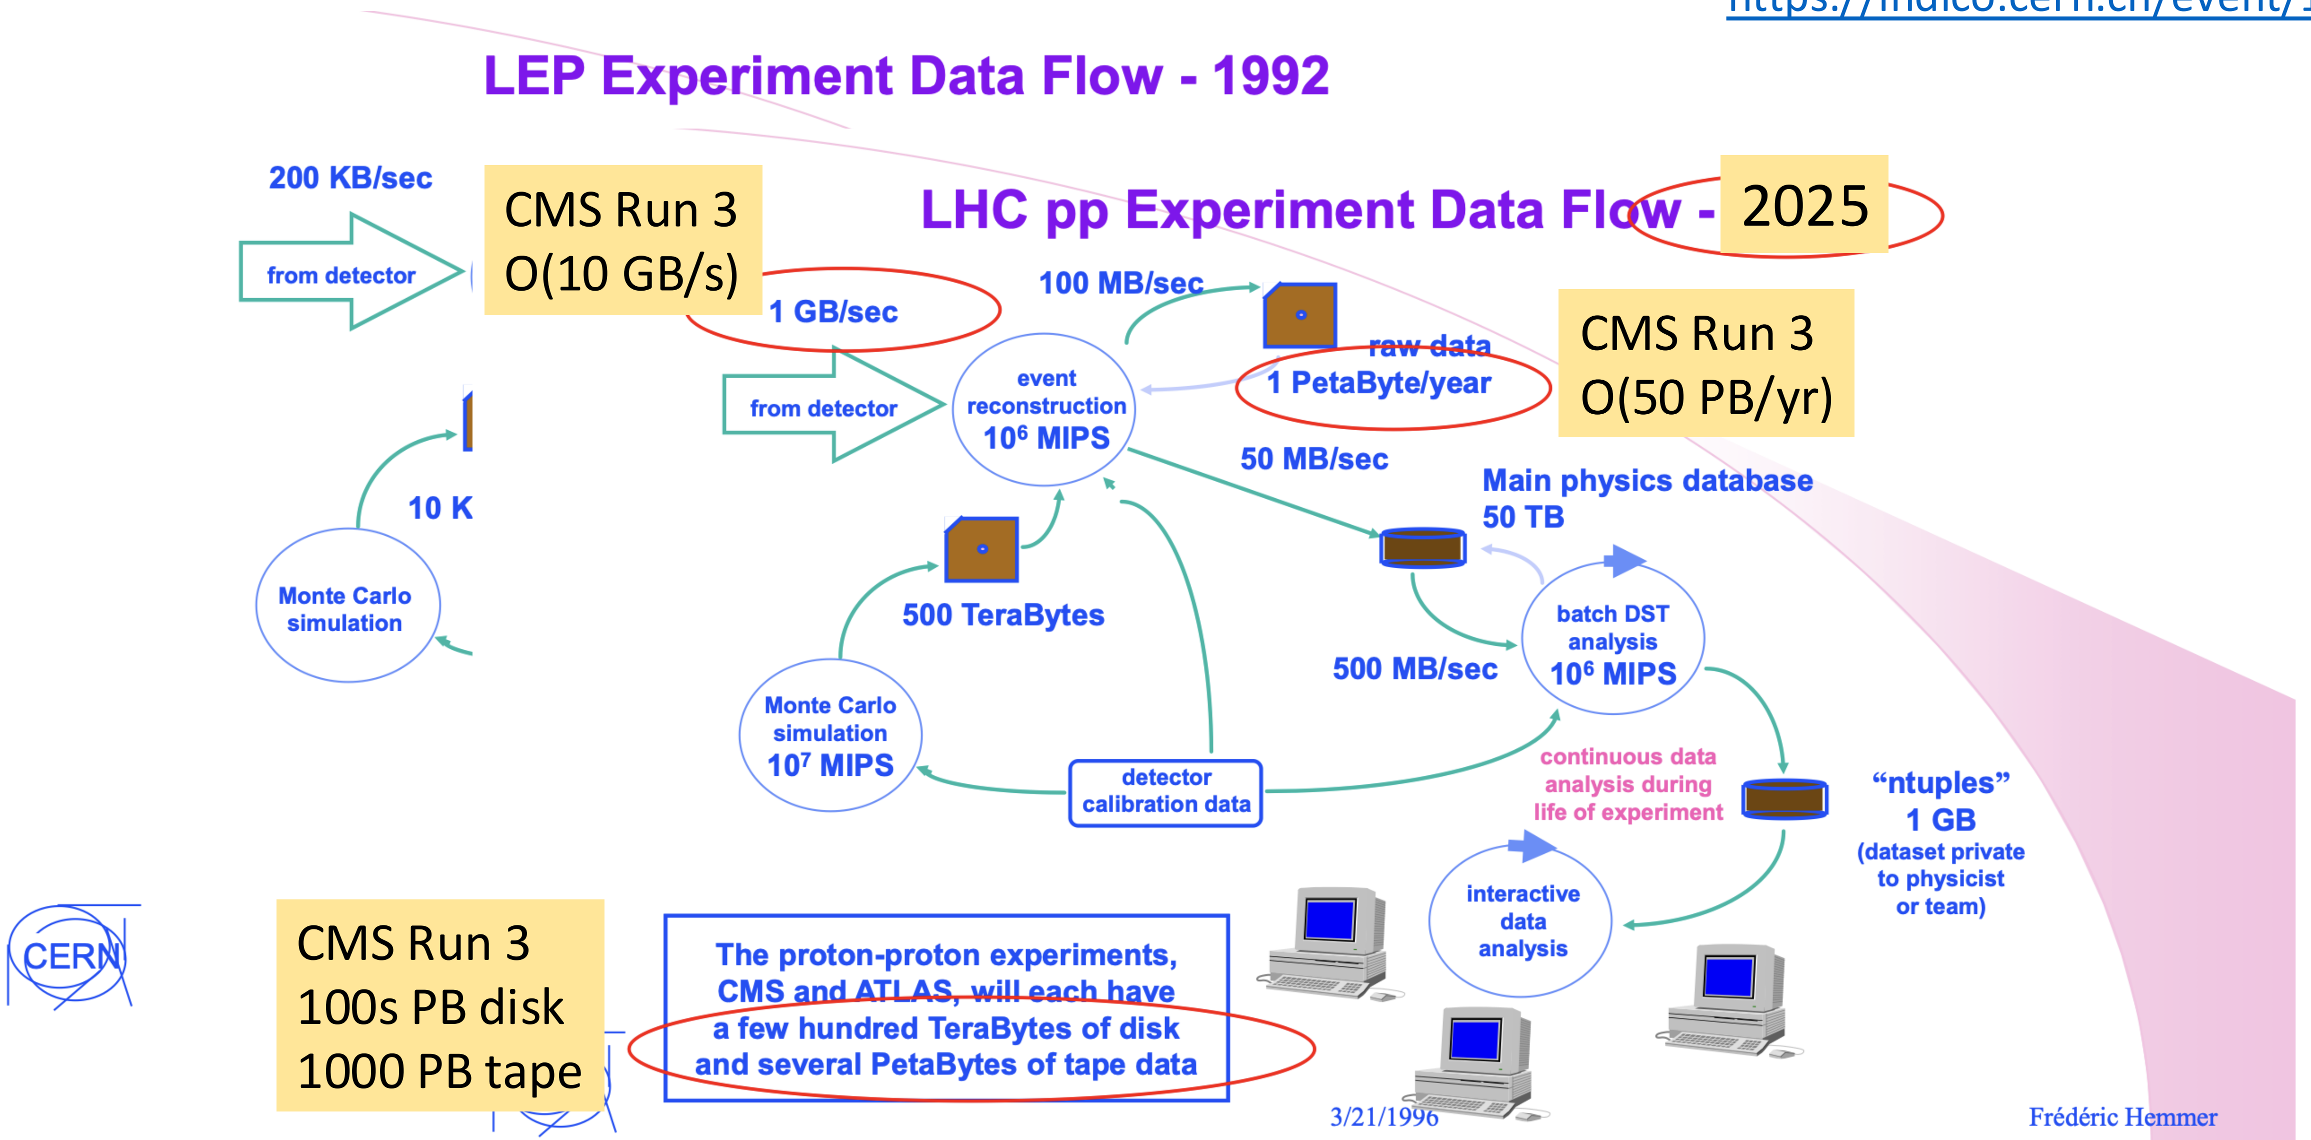

# Thinking about projections of experimental feasibility of future experiments over decades

https://indico.cern.ch/event/1052077/

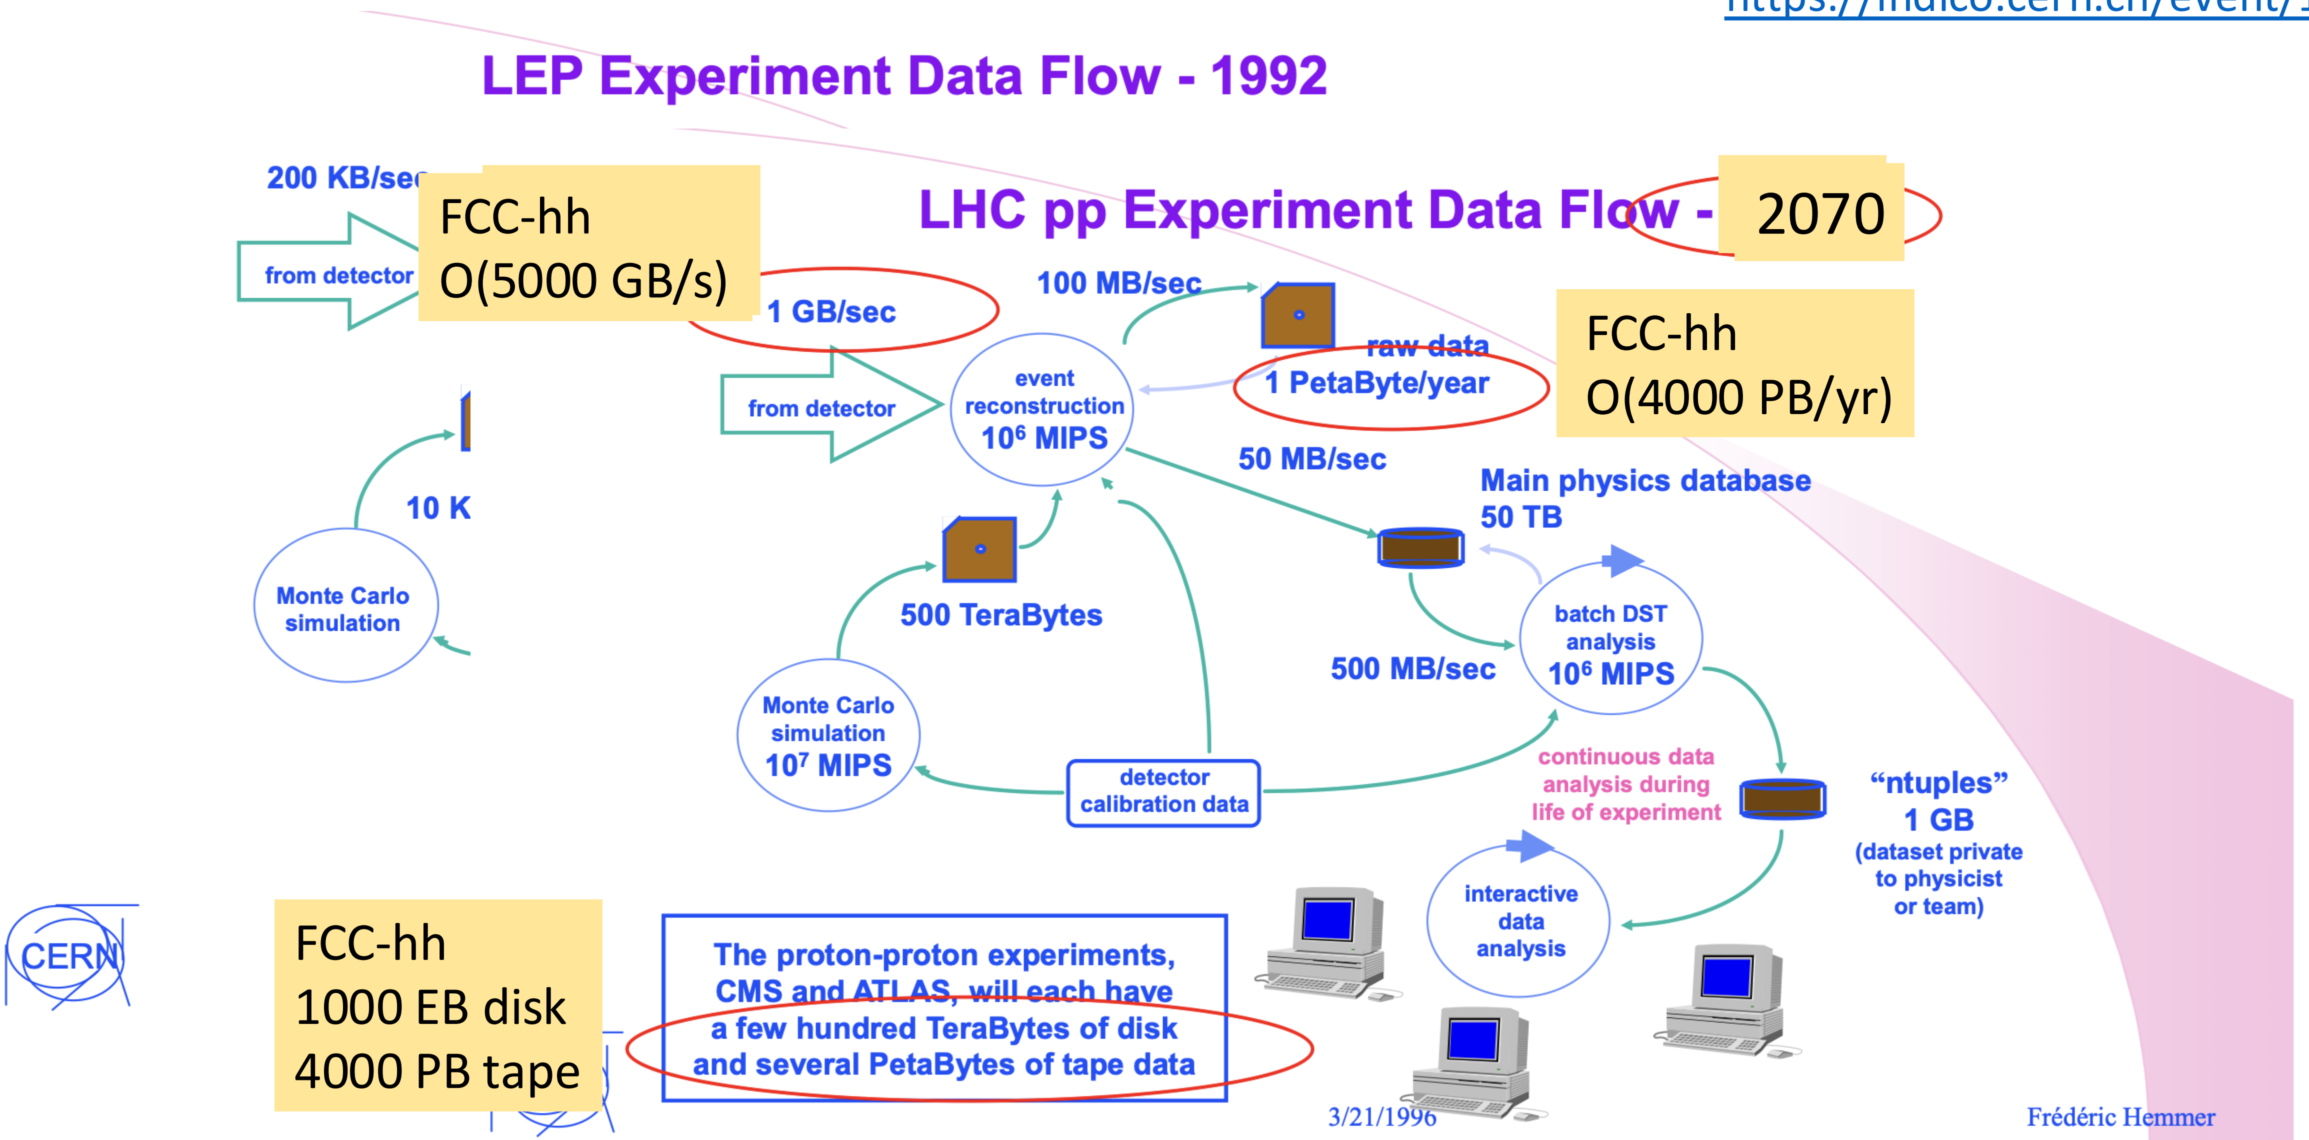

# Outlook and conclusion

- We will not be using the same technologies and approaches of today
  - Advancing technology and our innovation have enabled us to efficiently analyze 10^4 higher data rates in the last 30 years (that’s 25%/year)
- Even so, estimated needs and extrapolations show no large mismatch
  - Given such long timescales we must continue our R&D programs in order to exploit state-of-the-art technologies and approaches to fulfill computation and storage needs
- Expectations for experimental needs (and ambitions) tend to be underestimated.
  - Research that goes into making experiments reality opens new opportunities and creates new demands.
  - Taking advantage of new techniques and technologies (today’s examples AI/ML and heterogeneous systems) can facilitate these more ambitious programs to come as facilities are realized

# Backup…

# Compute modeling themes and challenges: Scale of HPC

Many submissions discussed the importance of the use and availability of HPC resources to our scientific pursuits

Potential next steps or input to discussion:

- To better understand the scale of HPC resources that would be useful to different communities and how this is likely to evolve in the future
- Activities where HPC resources are more inherently more capable than HTC. Eg, where HPC are something beyond an opportunistic resources or replacement for dedicated resources

# Compute modeling themes and challenges: New dimensions to resource projections/modeling

- Evolution in the necessary scale of AI/ML "production" and analysis resources: both in context of inference and training requirements

- Estimation (benchmarks) and optimization of energy footprints and environment impact

- Required I/O and network requirements such as tape bandwidth, WAN, transatlantic, and LAN

- Estimating and benchmarking capabilities of heterogeneous computing

# Notes

- MuonColl: Doc, Tables 22.1.1 and 22.2.1 storing RAW and AOD/analysis
- FCC-ee: Doc, Tables 21,22 and Figure 133 are for detector simulation. Then we assume 4x MC/data and that the reconstruction time is 30% of the sim time based on the doc. Raw and simulated raw are stored (data raw size from this doc)
- FCC-hh: Doc and Doc. Effective pileup and event rate suggests scaling of 100x HL-LHC and 2x longer program
- LHC: CMS Doc and ATLAS Doc together with extrapolation to end of planned running. Storage includes RAW
- LEP3: Estimated as 1/4 to 1/3 of FCC-ee (Doc)
- LHeC: Doc
- LC: https://arxiv.org/pdf/2503.24049 https://zenodo.org/records/4659567 https://zenodo.org/records/4659571 Z run extrapolated from FCC-ee