The upgraded LHCb detector DAQ and trigger: design and first performances

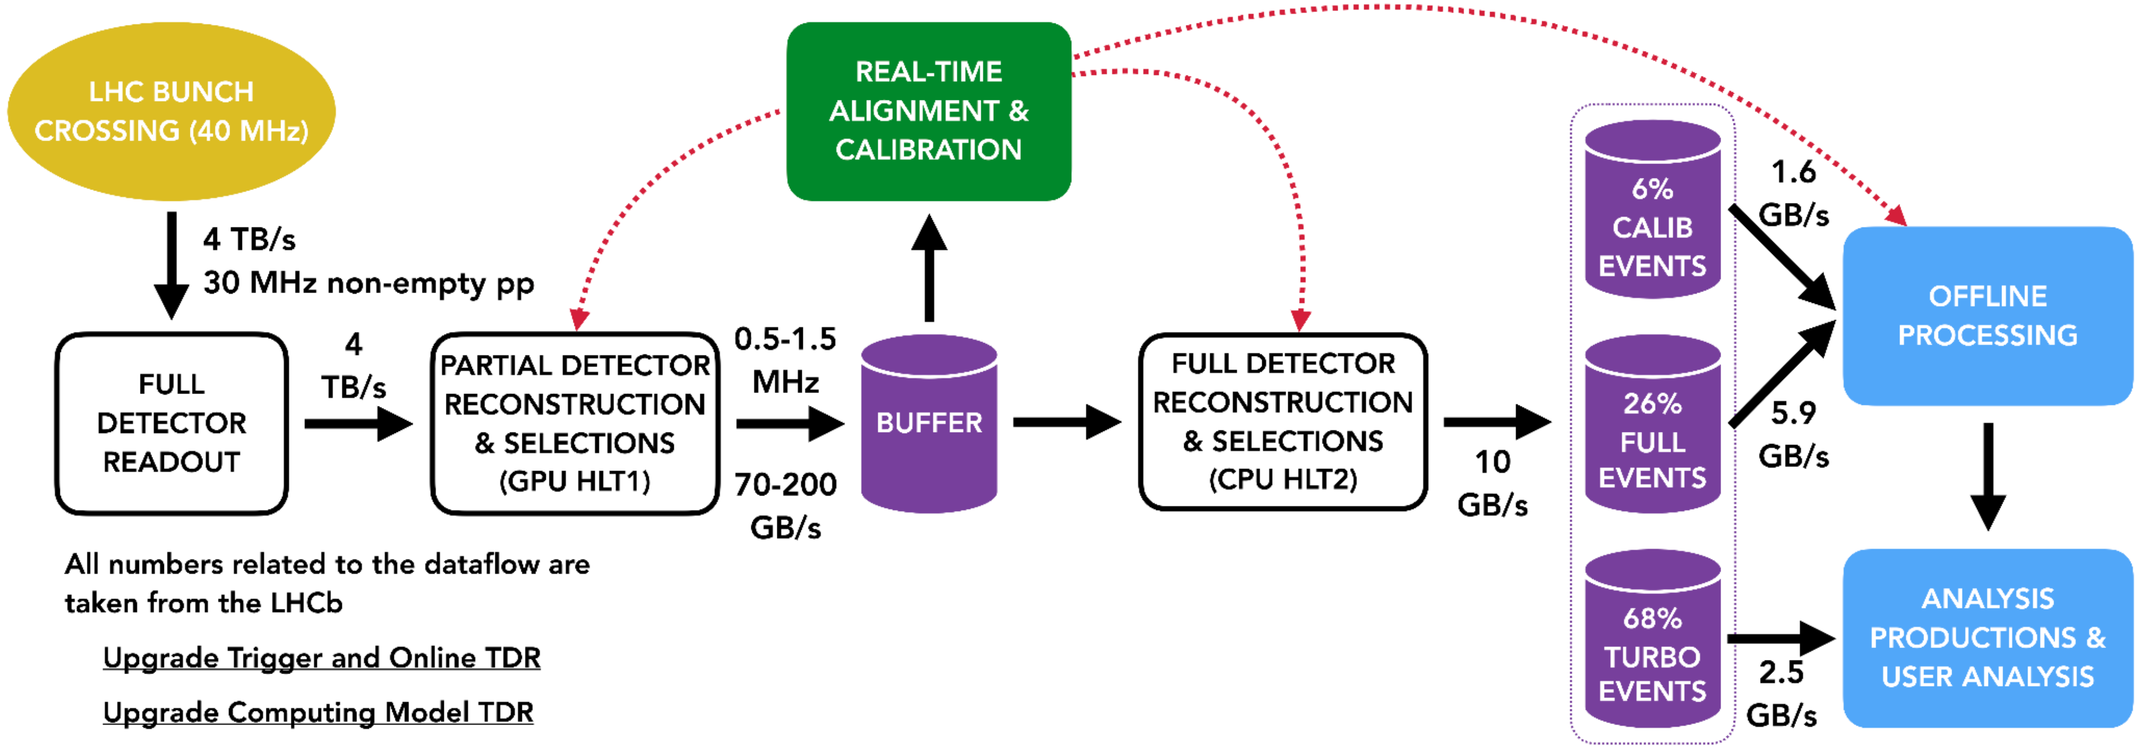

All numbers related to the dataflow are taken from the LHCb

Upgrade Trigger and Online TDR

Upgrade Computing Model TDR

V. V. Gligorov, CNRS/LPNHE

FCC-ee detector concepts meeting, 03.06.2024

The LHCb detector at the LHC

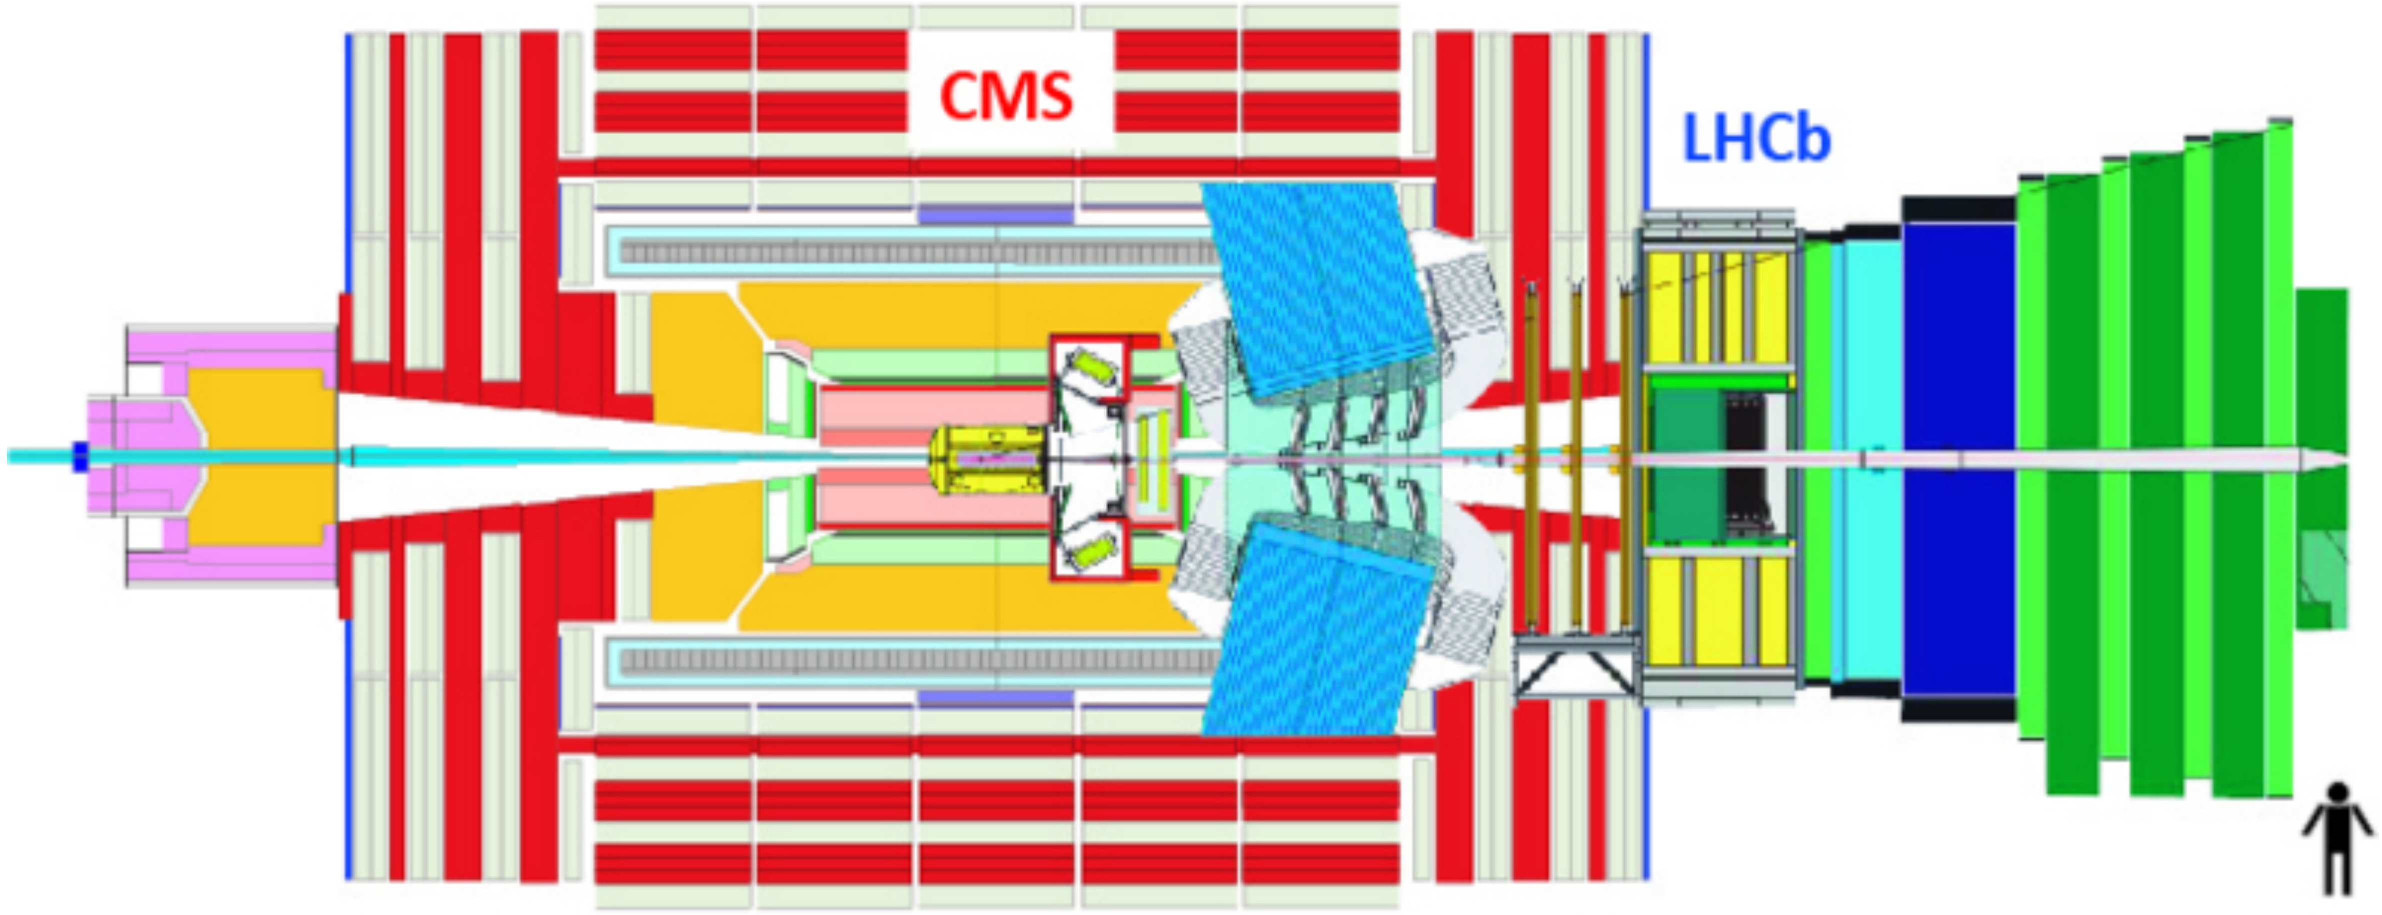

Forward spectrometer optimized for precision physics

The challenge of the LHCb upgrade in one slide

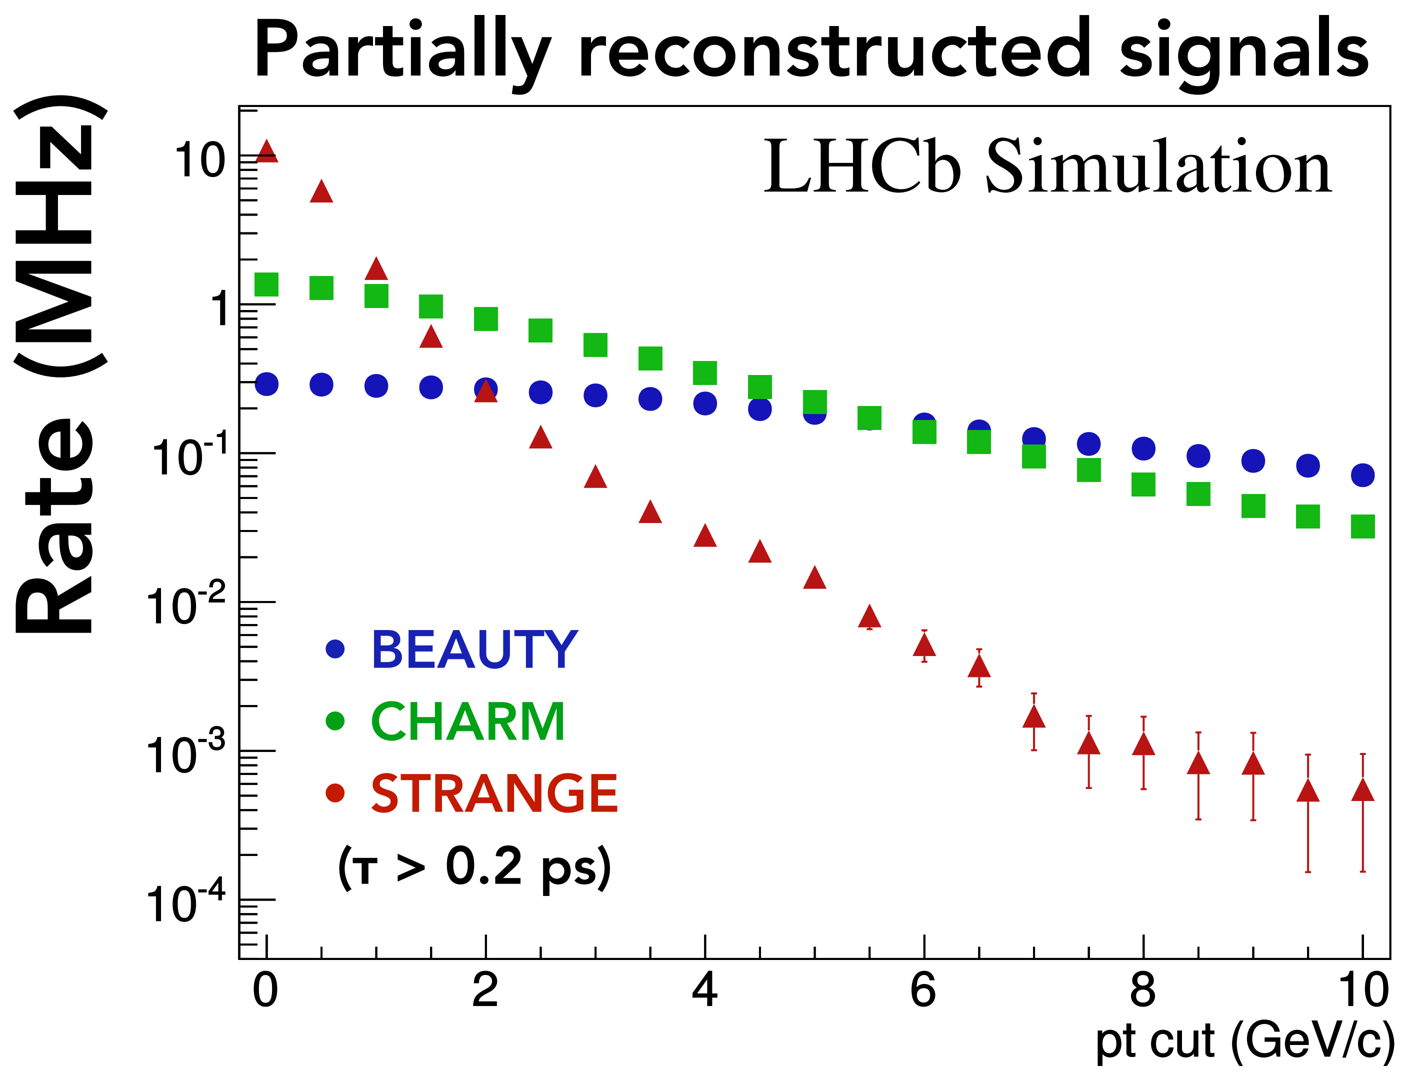

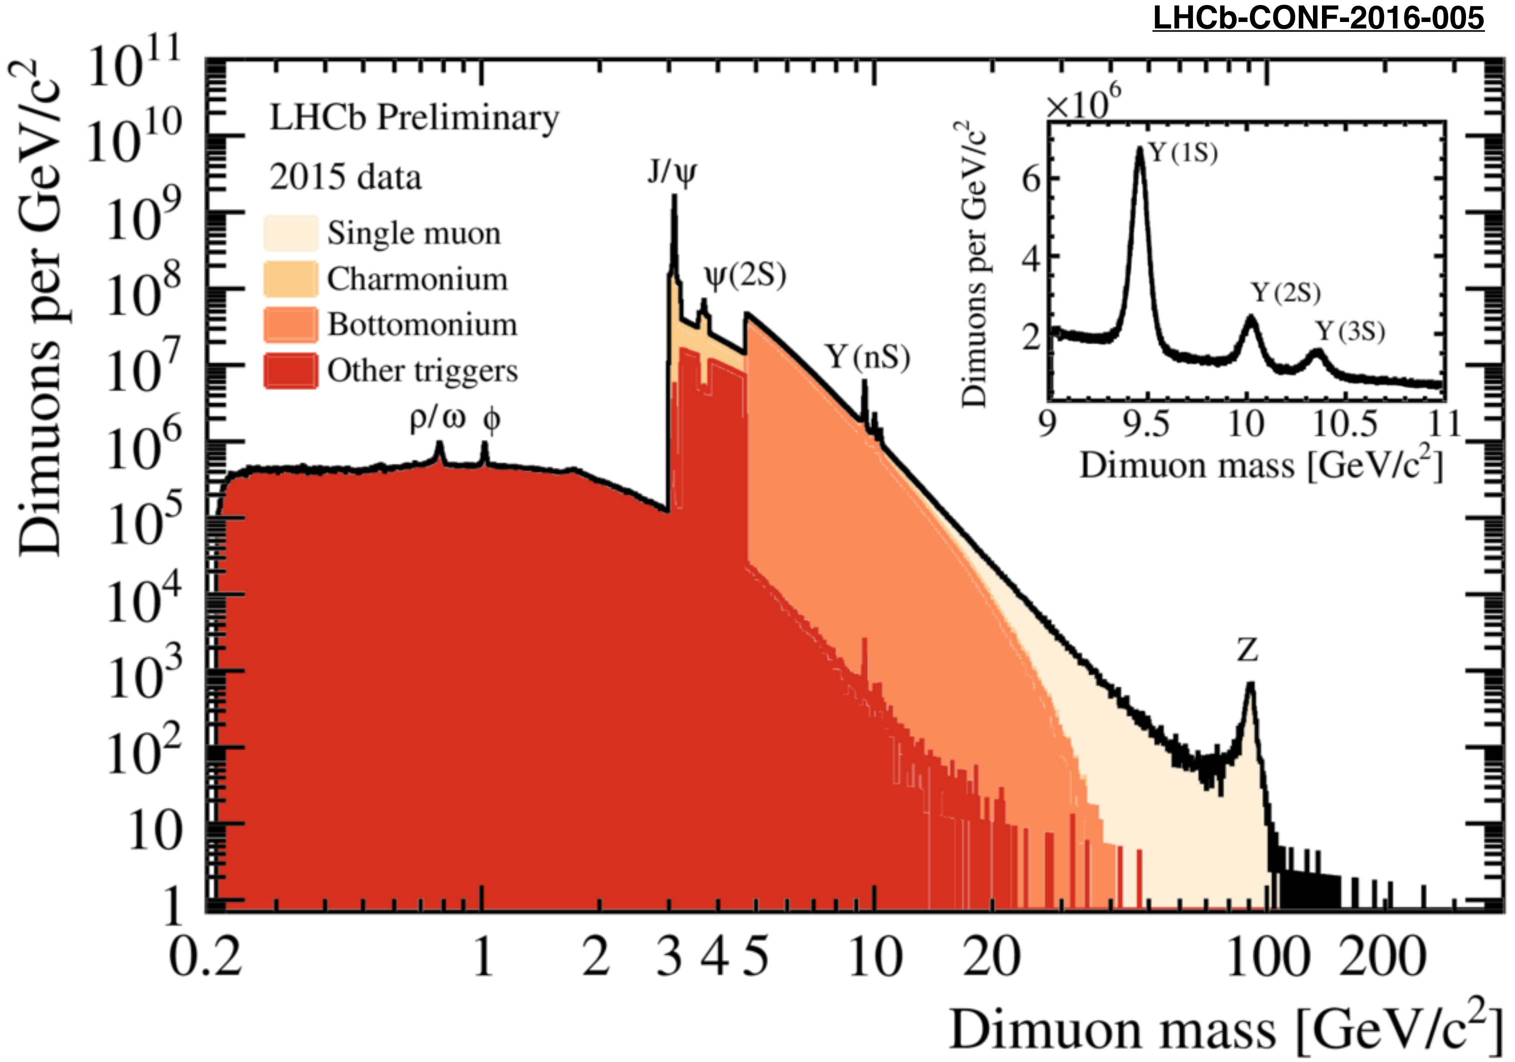

Need to handle signals from O(1) to O(100) GeV in mass
>MHz rates of soft signals, can only afford to fully store detector data for O(100) kHz of events

From this follows the LHCb DAQ design for the upgrade

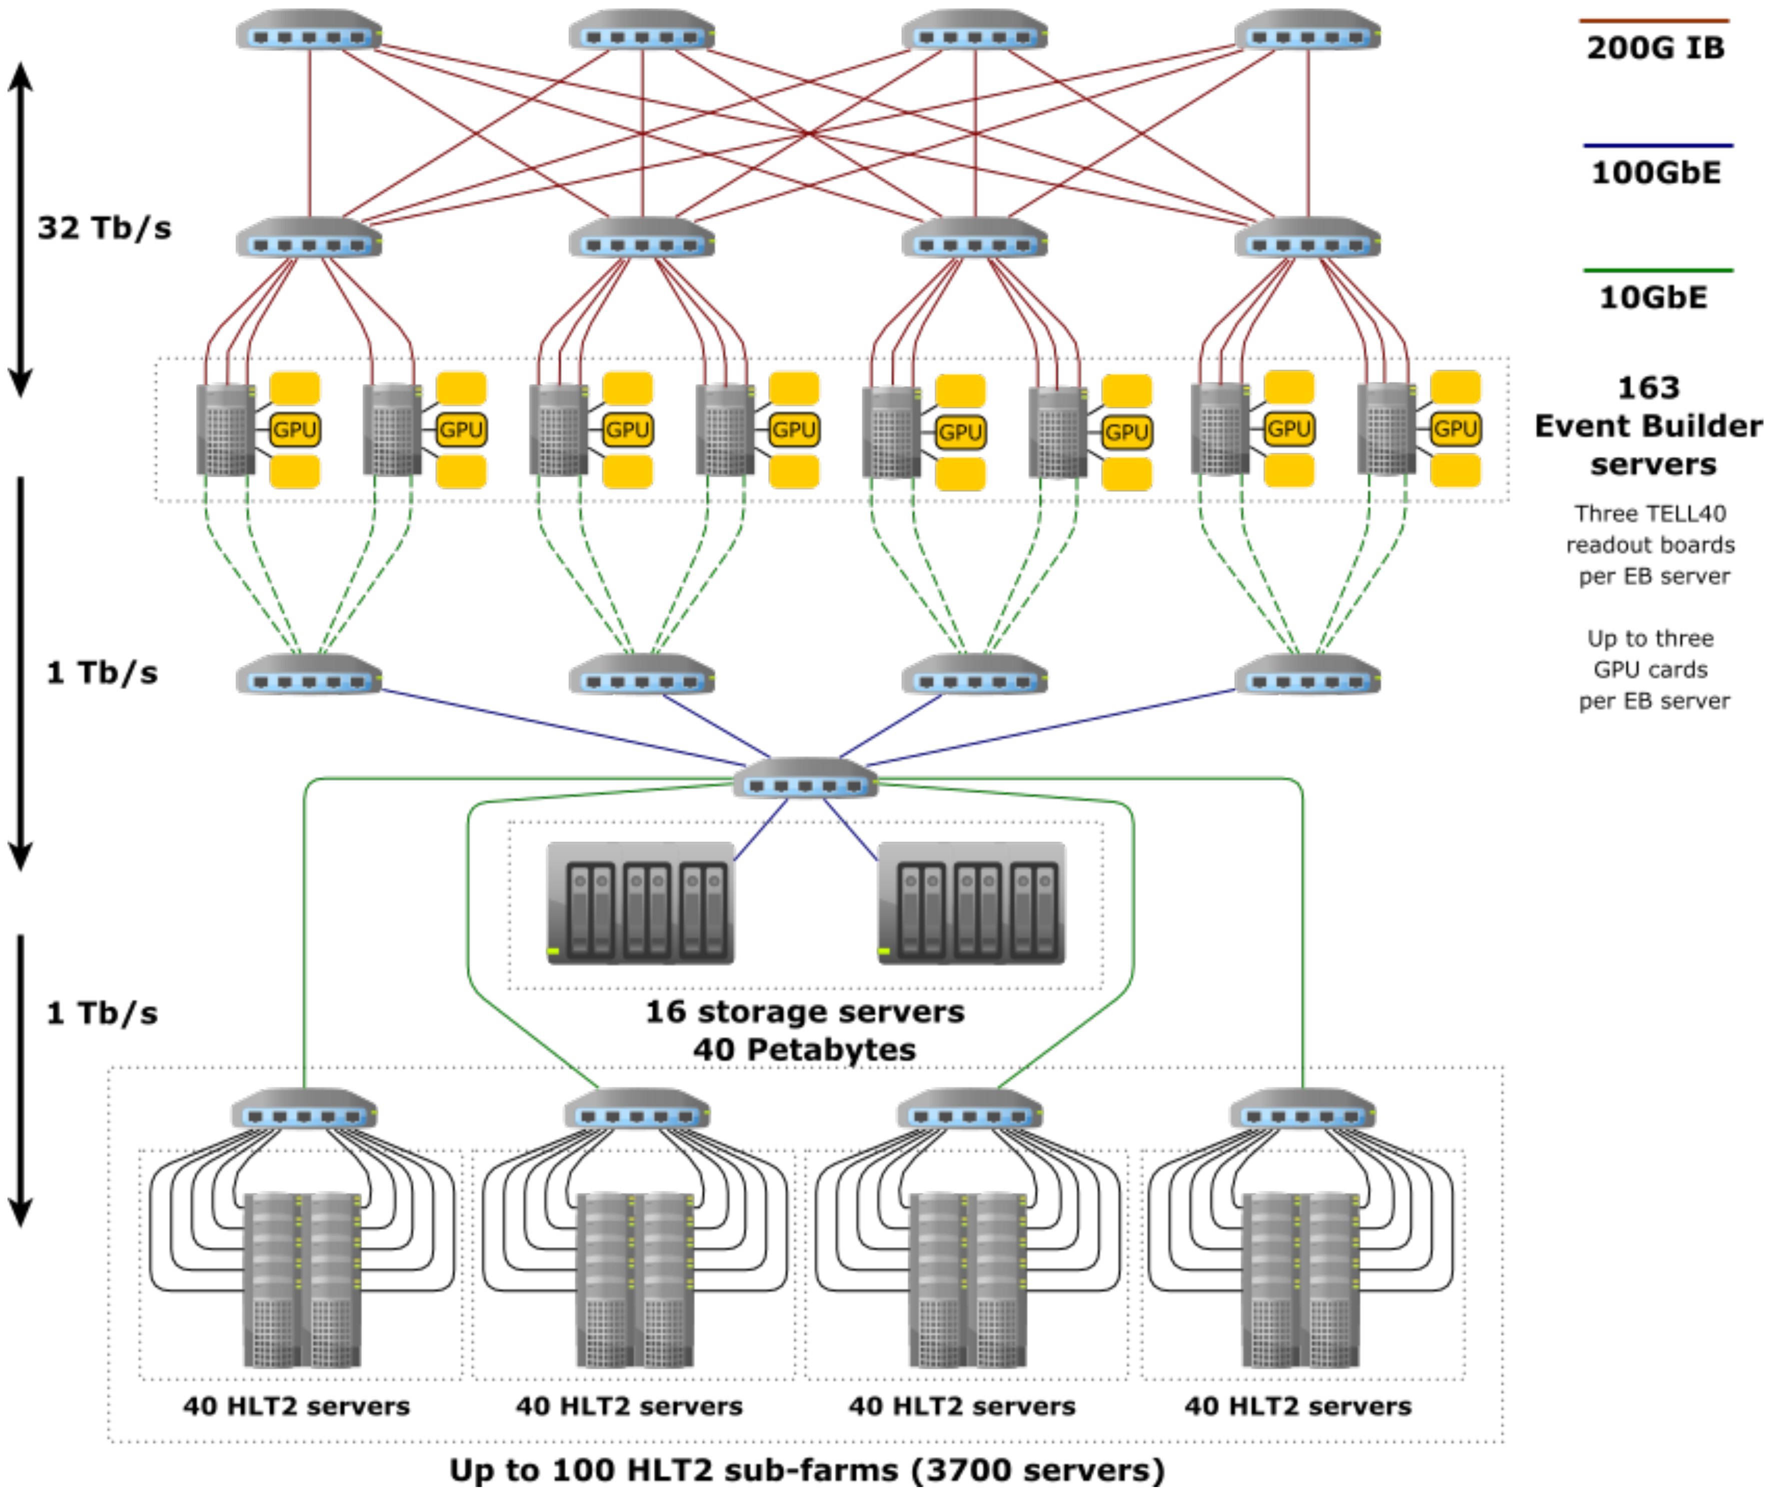

32 Tbit/s full event building & processing in a data centre
Inherent flexibility to choose a processing architecture based on cost/benefit considerations

Event readout (aka DAQ)

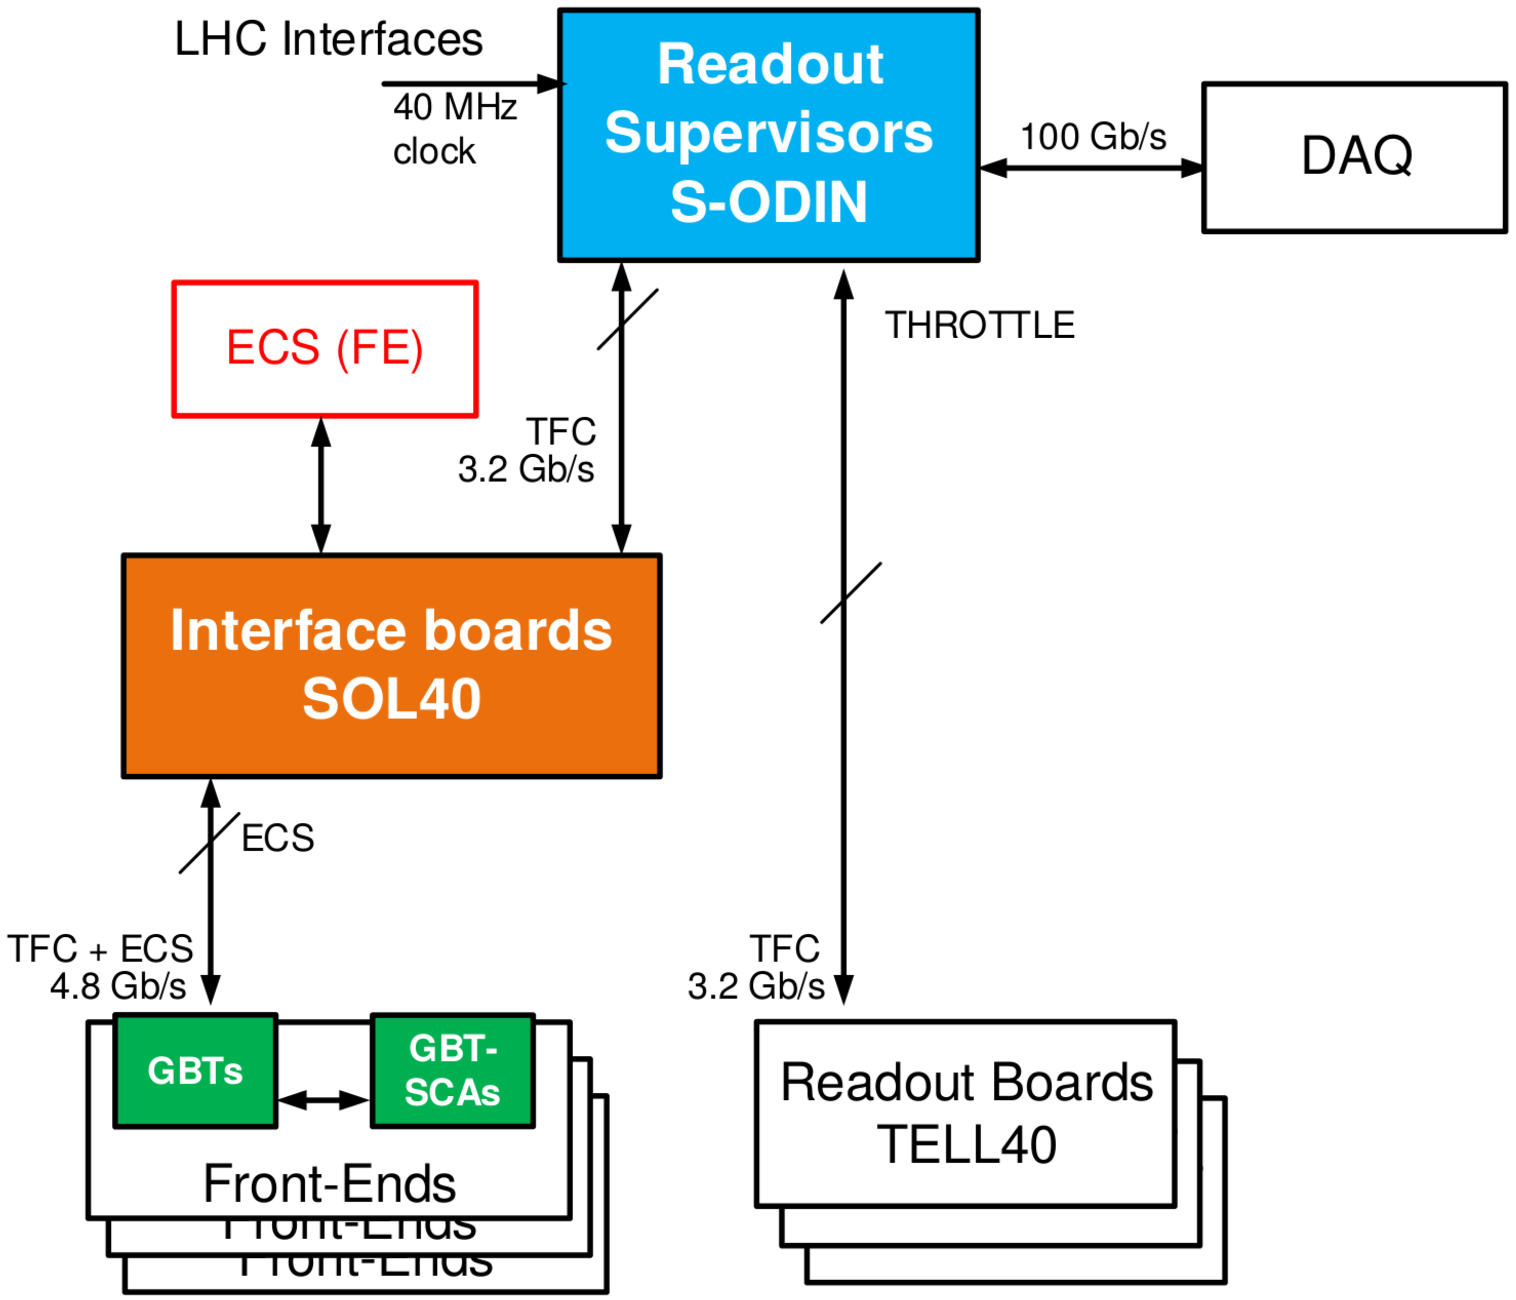

F. Alessio et al.
https://inspirehep.net/literature/1346081

Event readout (aka DAQ)

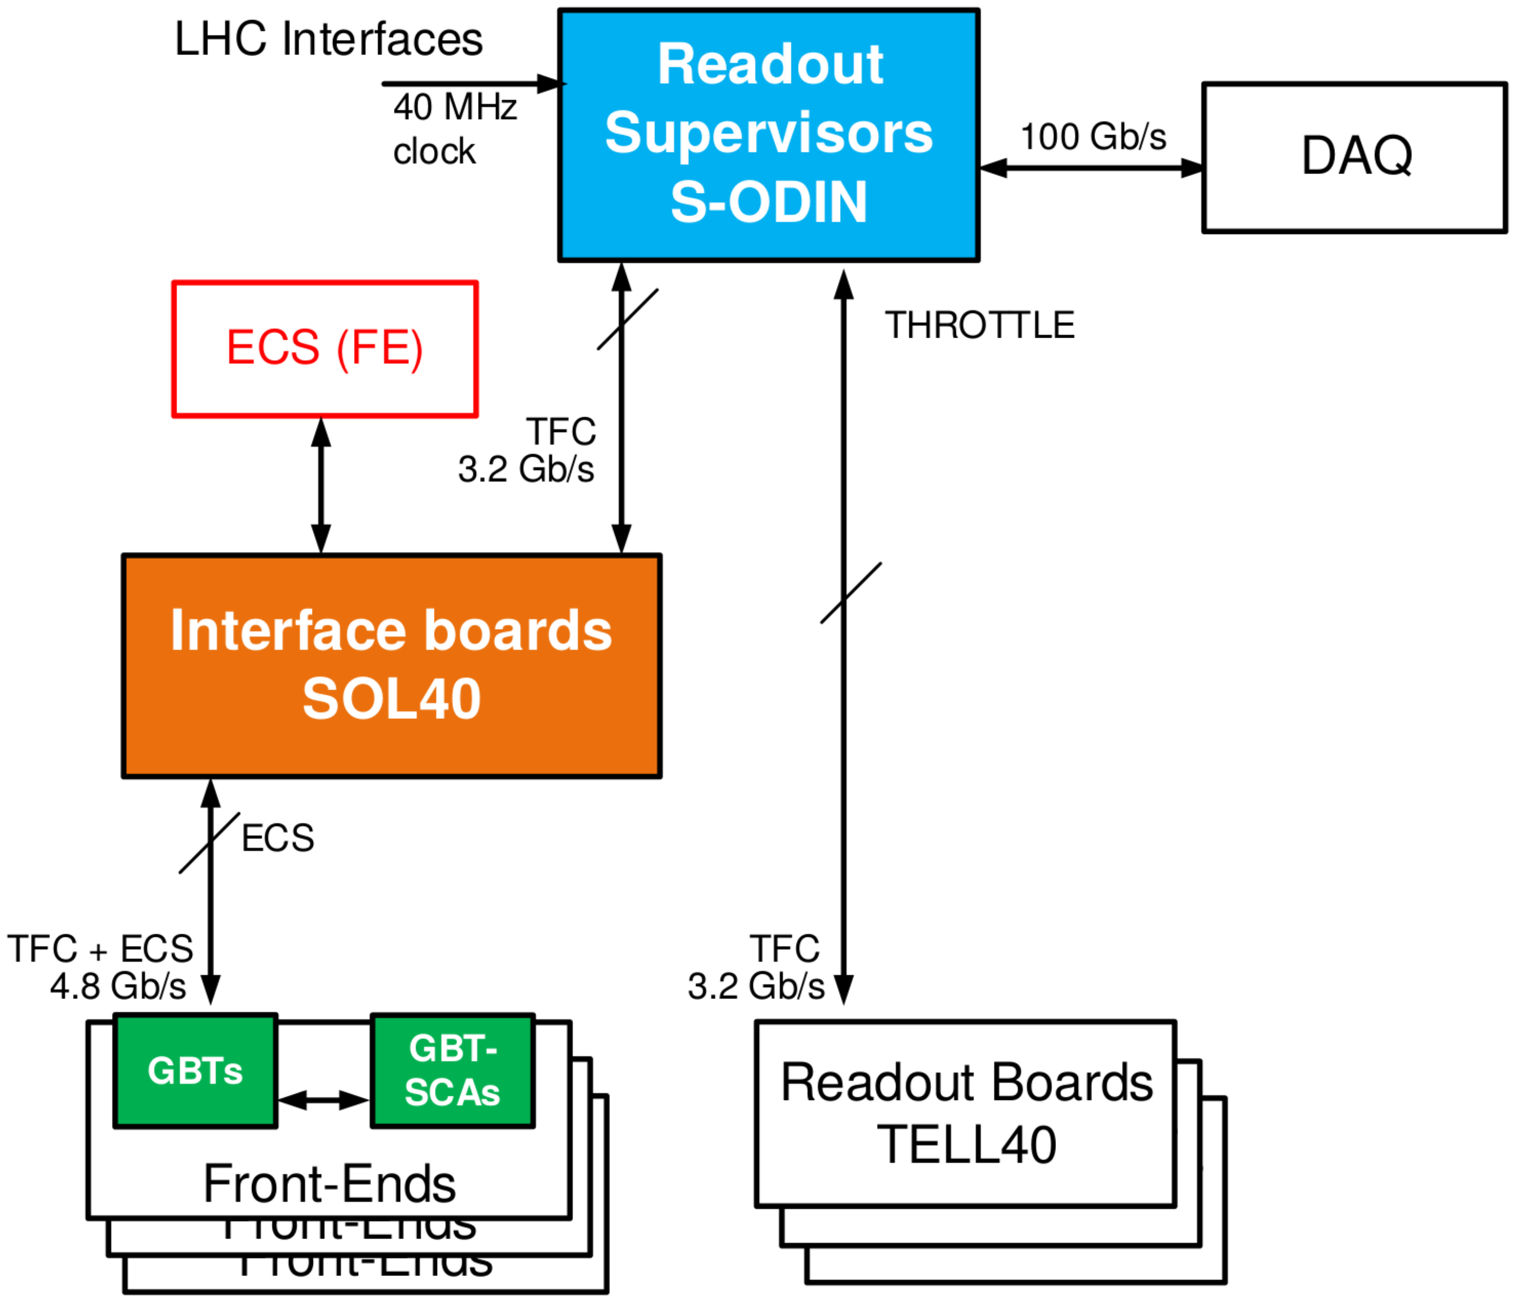

F. Alessio et al.
https://inspirehep.net/literature/1346081

Not a streaming DAQ as such.

Timing is centrally distributed to the front-ends, together with slow control information.

The readout supervisor has a limited ability to apply prescales (for us normally based on the bunch crossing type) to the information being read from the front end.

Event readout (aka DAQ)

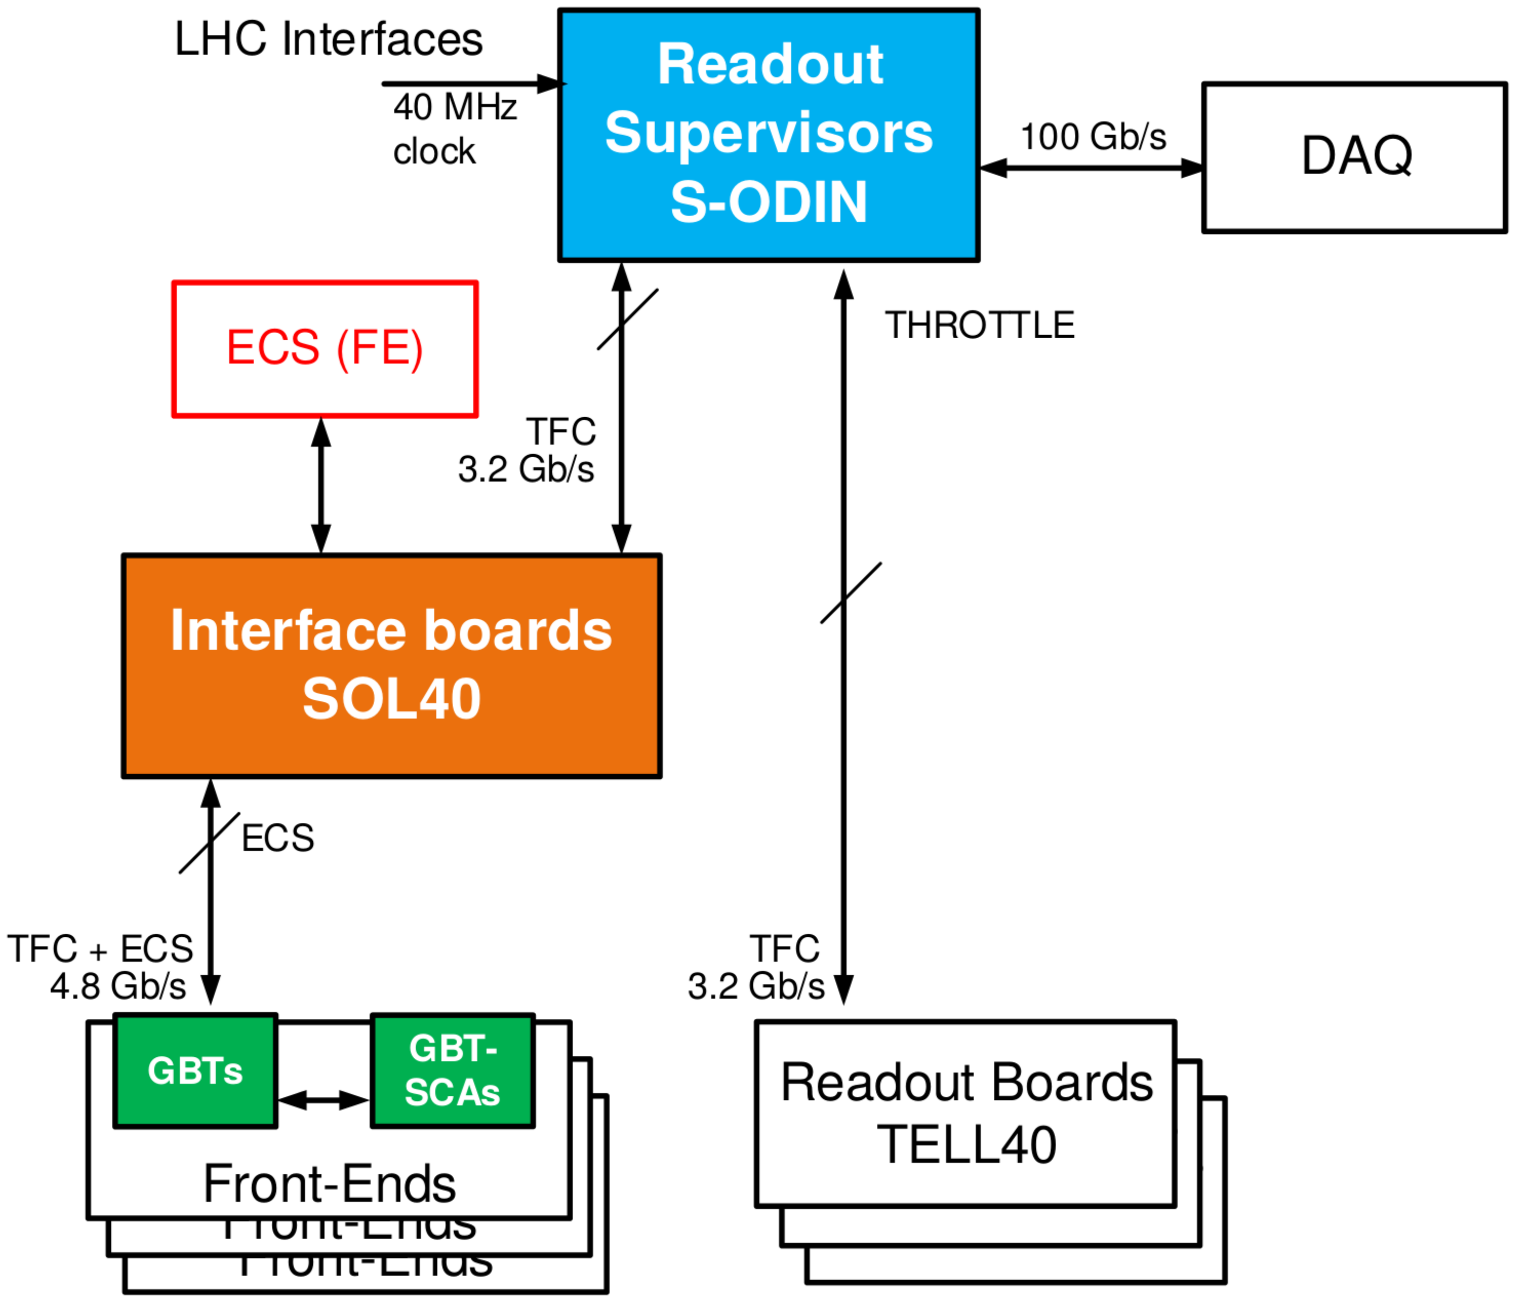

F. Alessio et al.
https://inspirehep.net/literature/1346081

Not a streaming DAQ as such.

Timing is centrally distributed to the front-ends, together with slow control information.

The readout supervisor has a limited ability to apply prescales (for us normally based on the bunch crossing type) to the information being read from the front end.

Luminosity determination is based on selecting a fixed rate (30 kHz) of random bunch crossings and propagating these events together with all other physics triggers to the offline processing.

Event readout (aka DAQ)

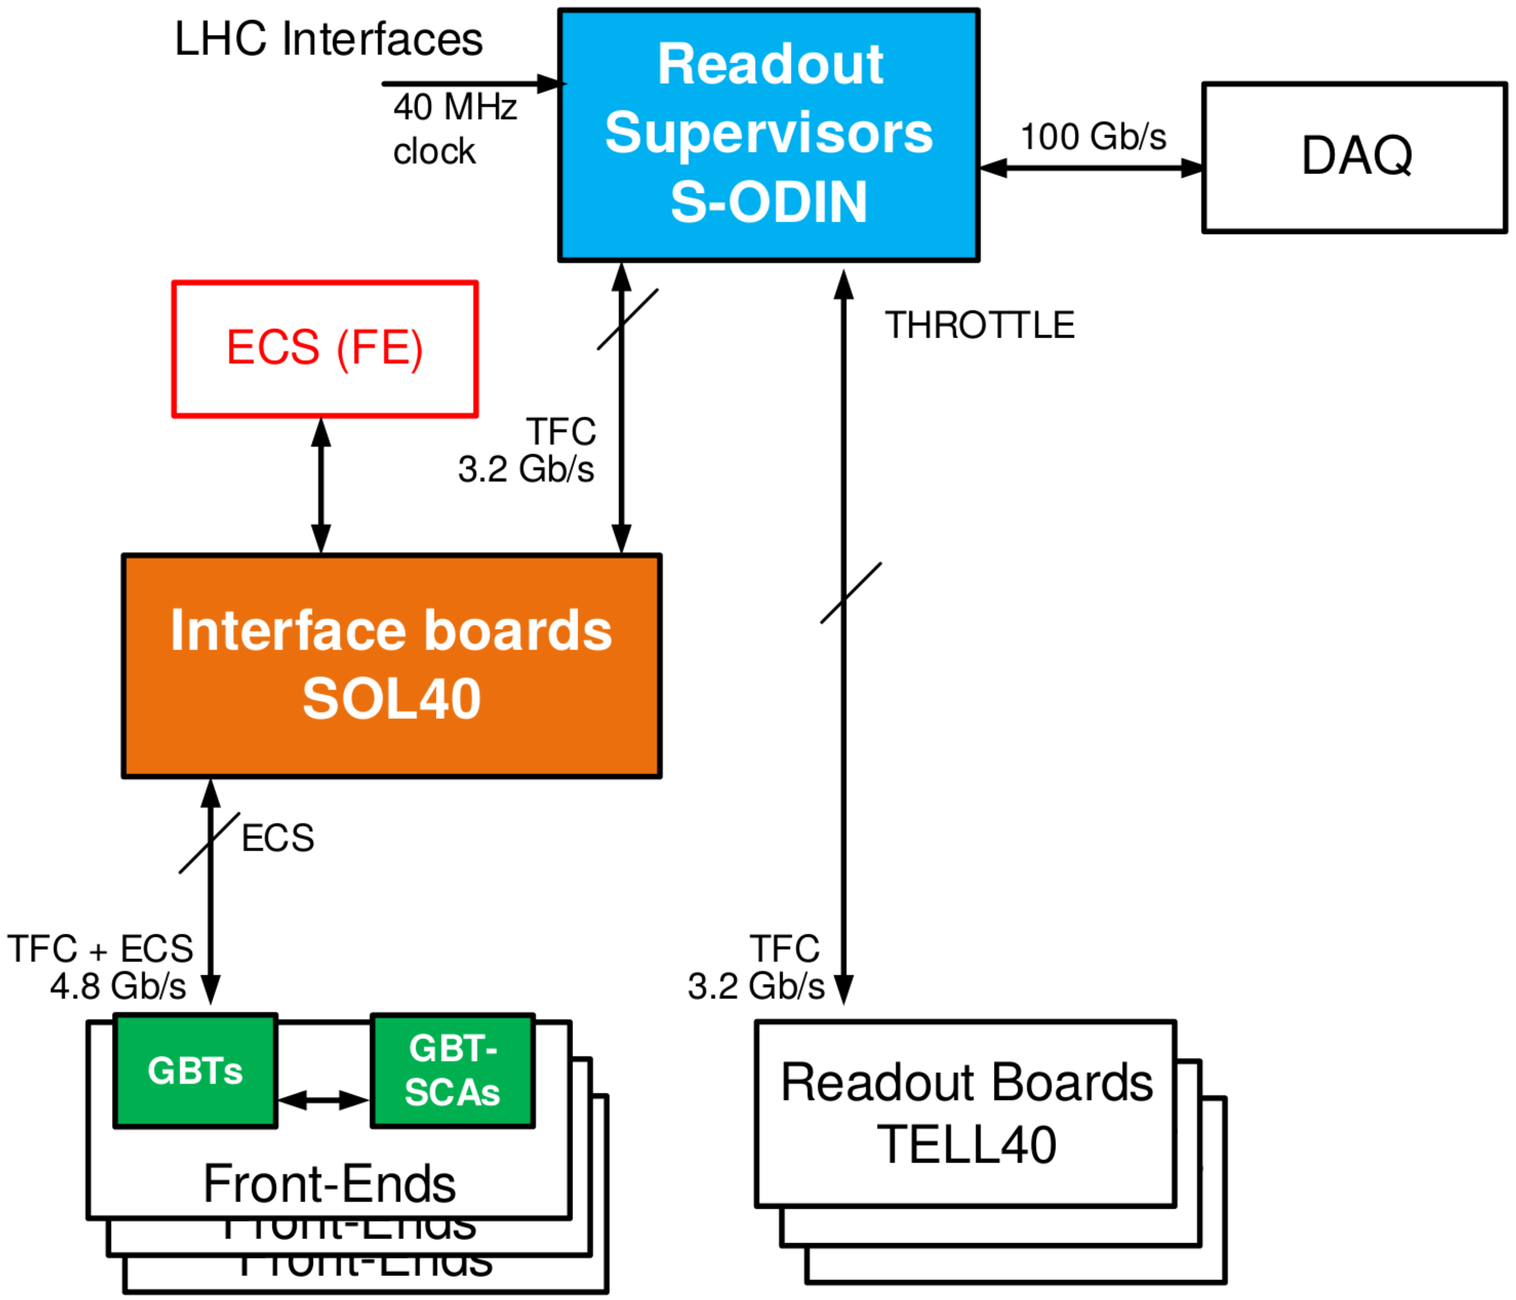

F. Alessio et al.
https://inspirehep.net/literature/1346081

Not a streaming DAQ as such.

Timing is centrally distributed to the front-ends, together with slow control information.

The readout supervisor has a limited ability to apply prescales (for us normally based on the bunch crossing type) to the information being read from the front end.

Luminosity determination is based on selecting a fixed rate (30 kHz) of random bunch crossings and propagating these events together with all other physics triggers to the offline processing.

As lumi events are mixed with all other event types, randomly dropping data packets does not cause a bias. We have never needed to examine if this holds at 10^{-5} however, so some fine print may apply.

Let's look inside an "event building" server

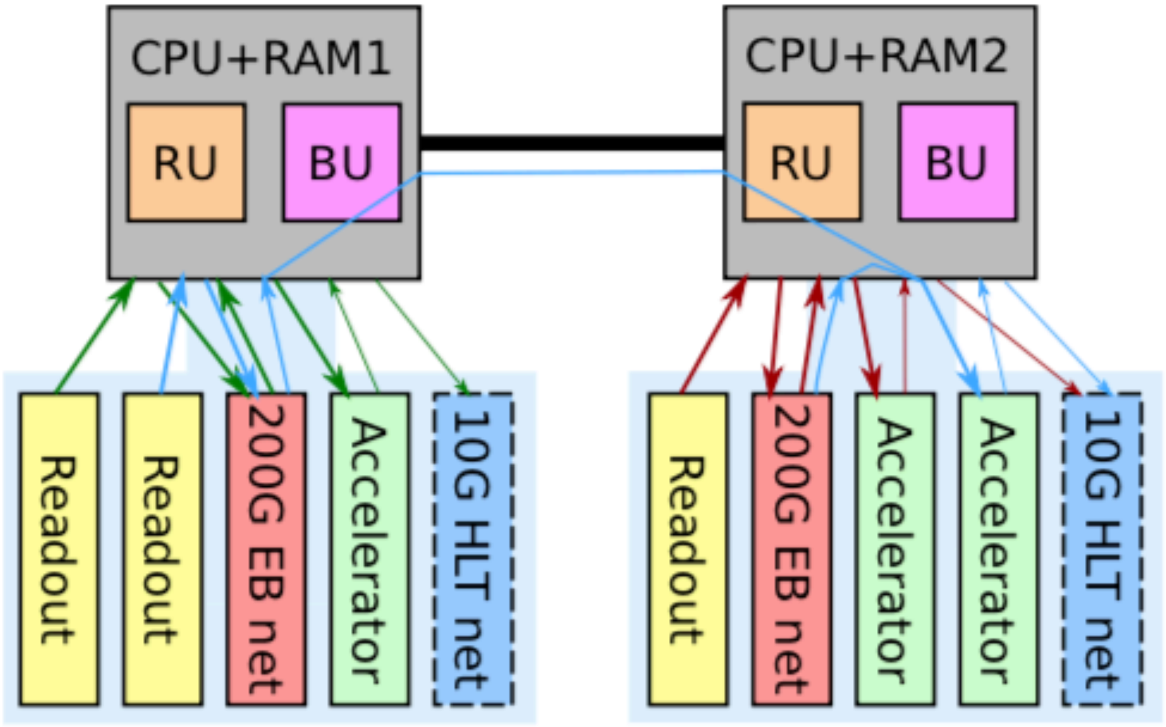

GPU-equipped event builder PC, with traffic of all three readout cards.

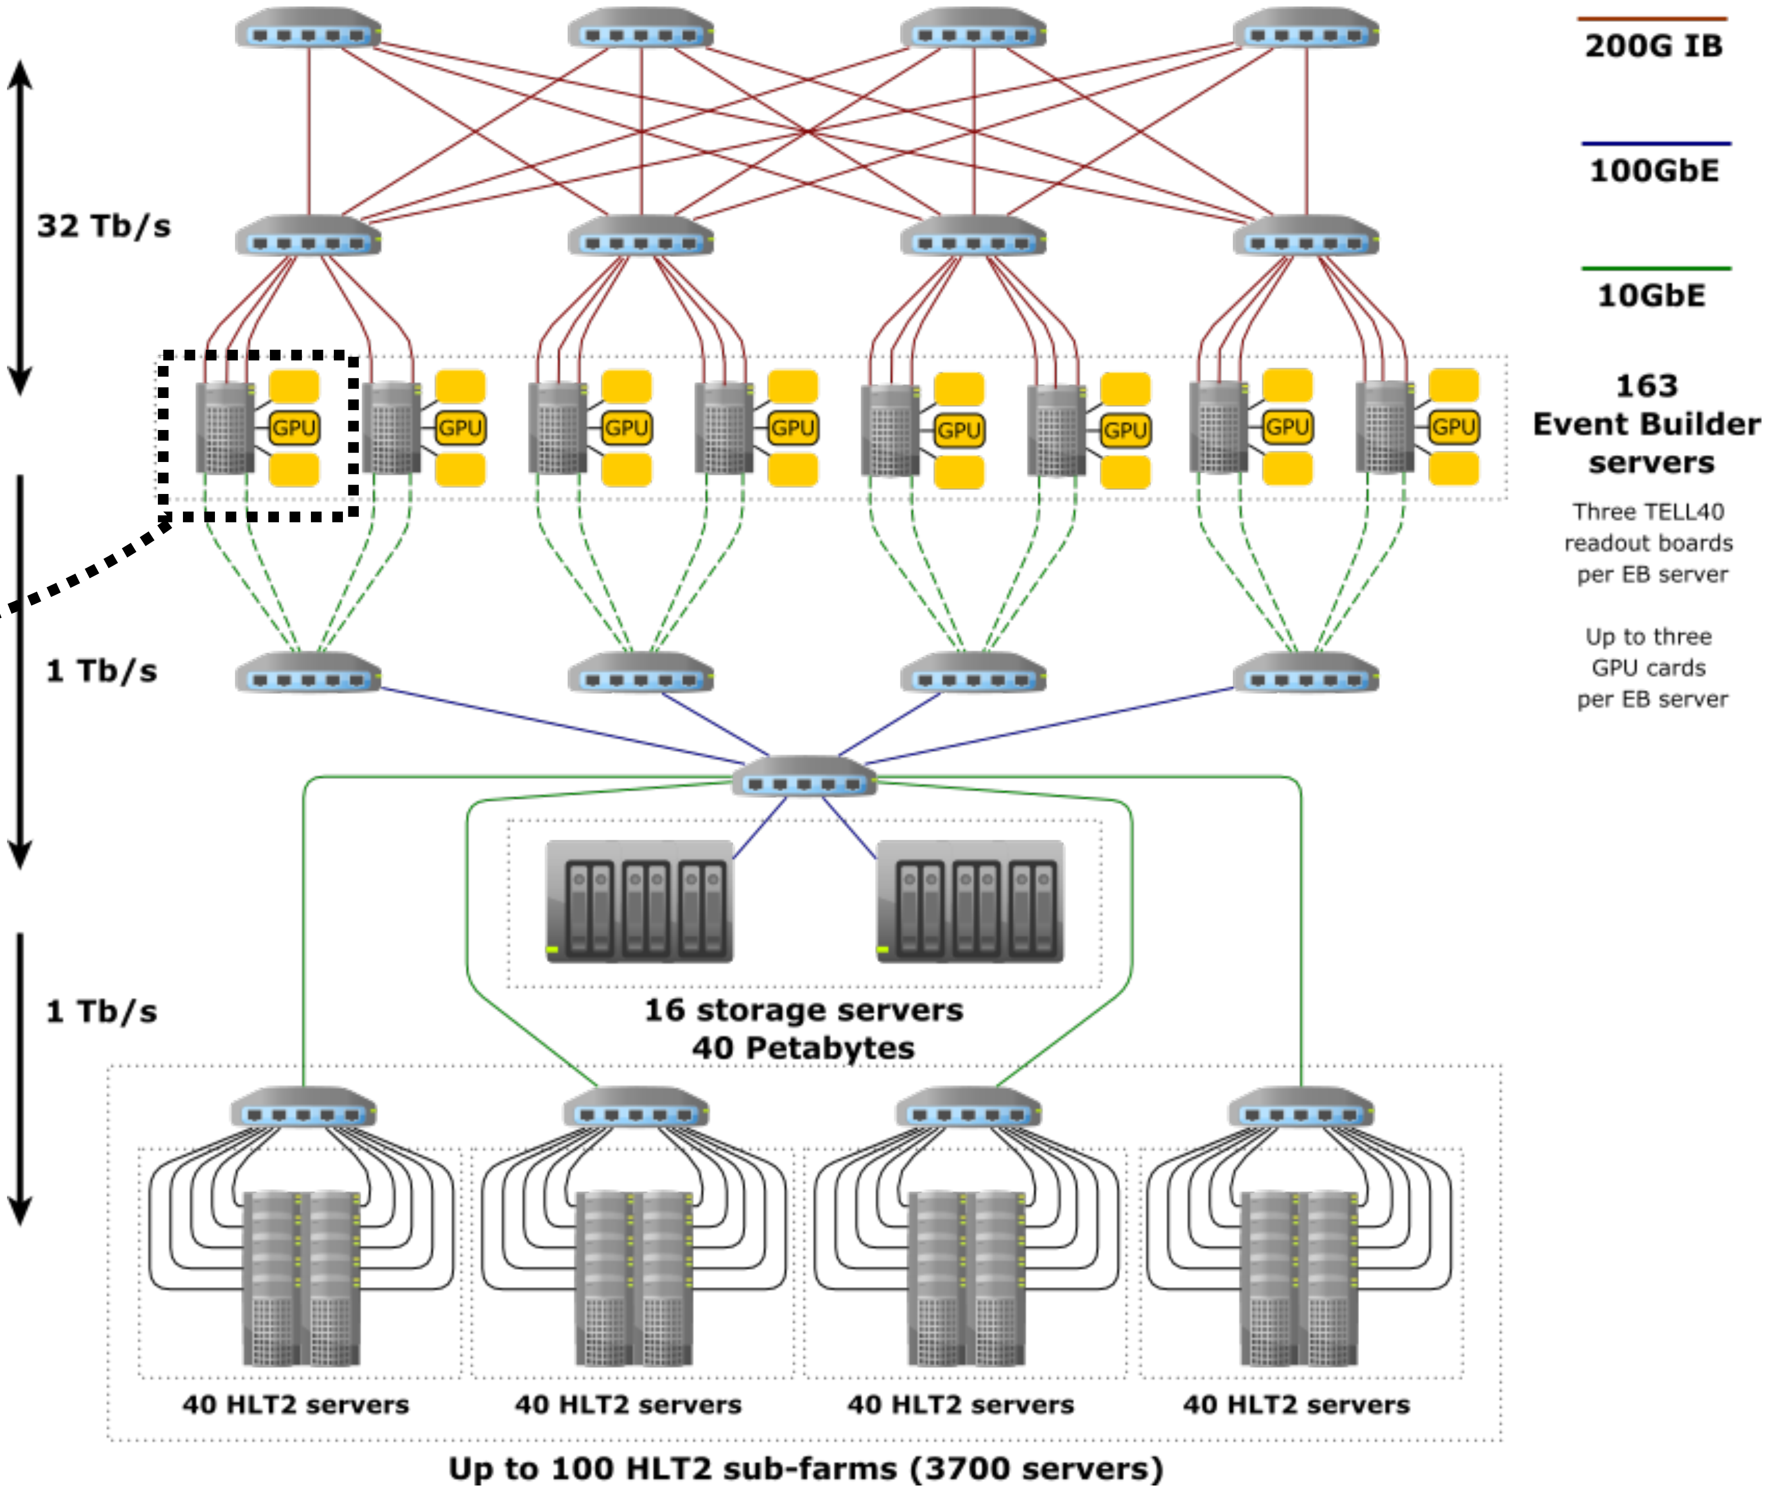

Once data packets arrive, they are assembled in the memory of the event building servers and then fed to the first-level trigger processors (in our case GPUs, but could be anything).

Let's look inside an "event building" server

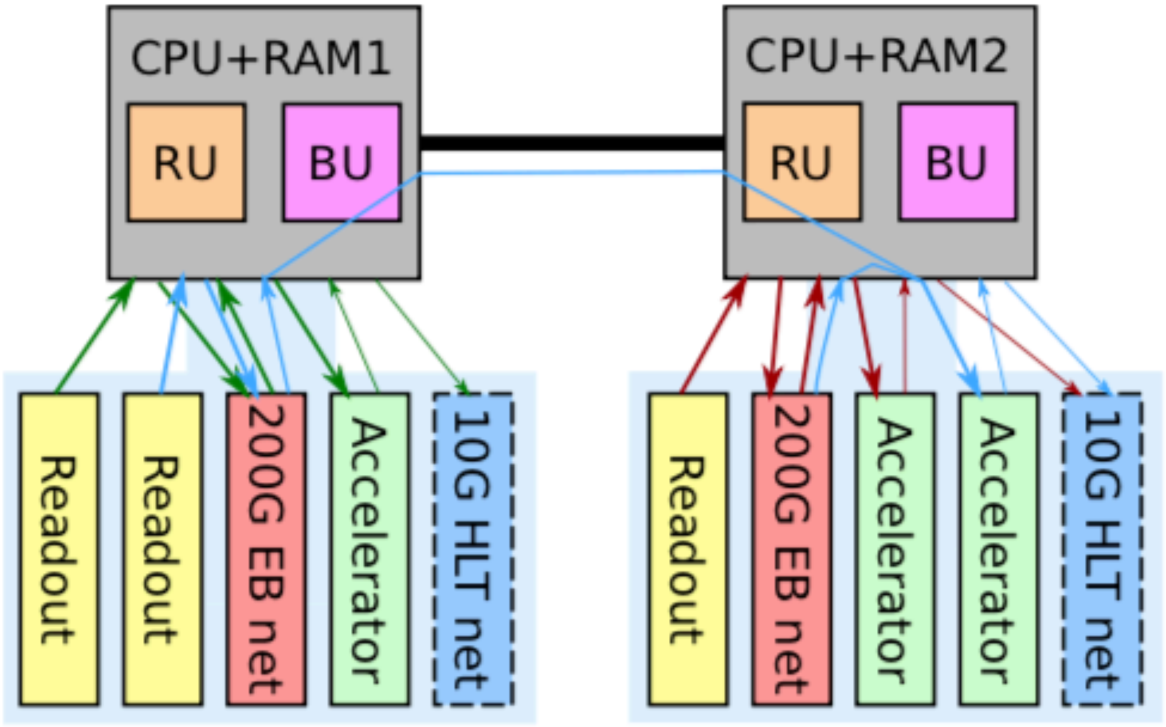

GPU-equipped event builder PC, with traffic of all three readout cards.

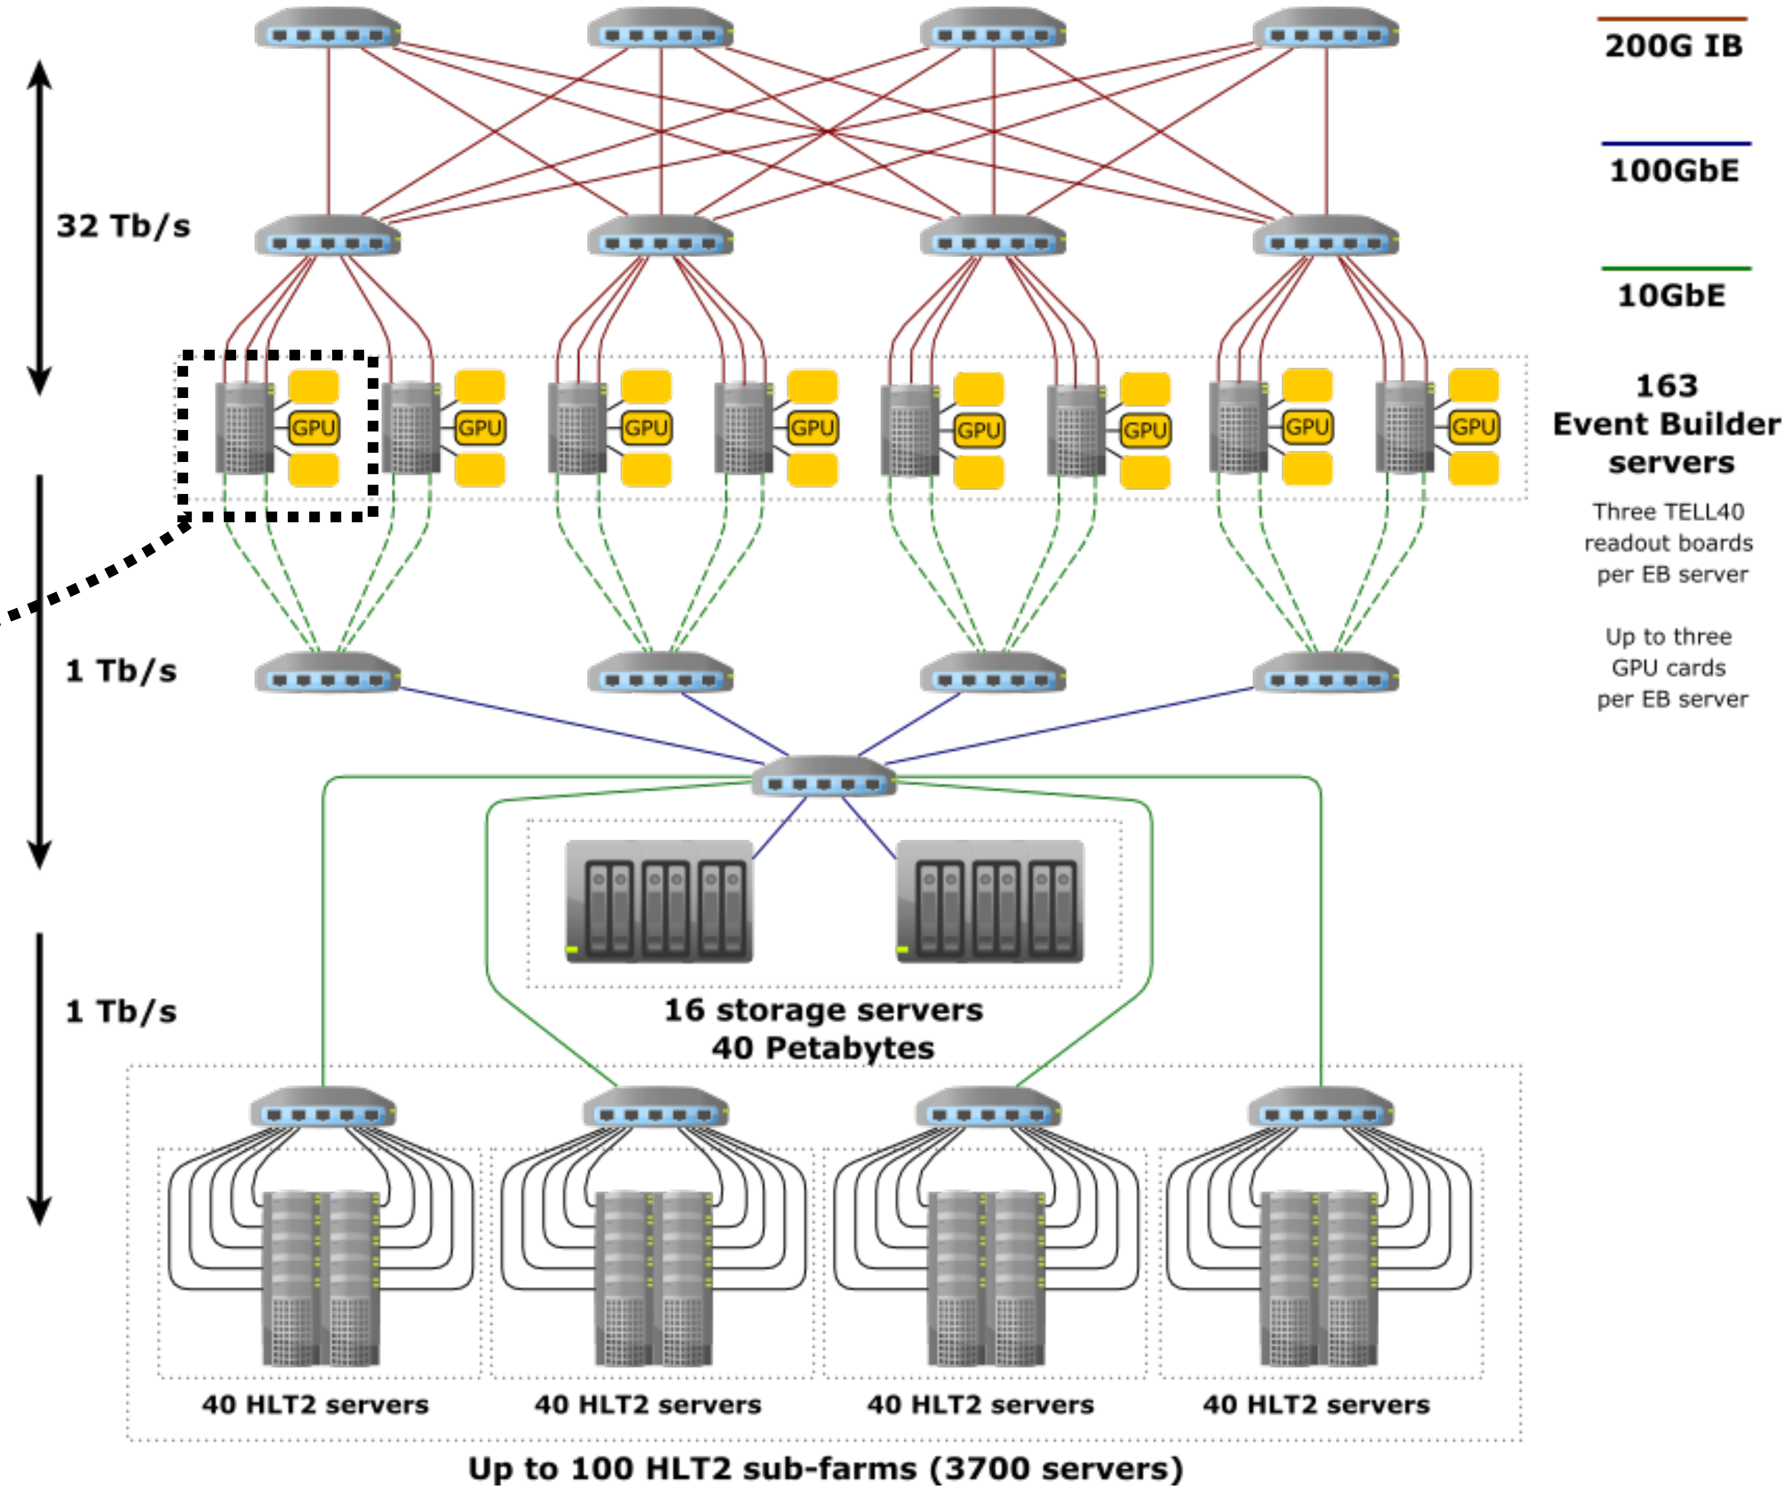

Packets dropped due to I/O issues or backpressure are monitored however because of mixing of luminosity events also cannot cause a bias in the eventual luminosity calculation.

What do we need to reconstruct @30 MHz?

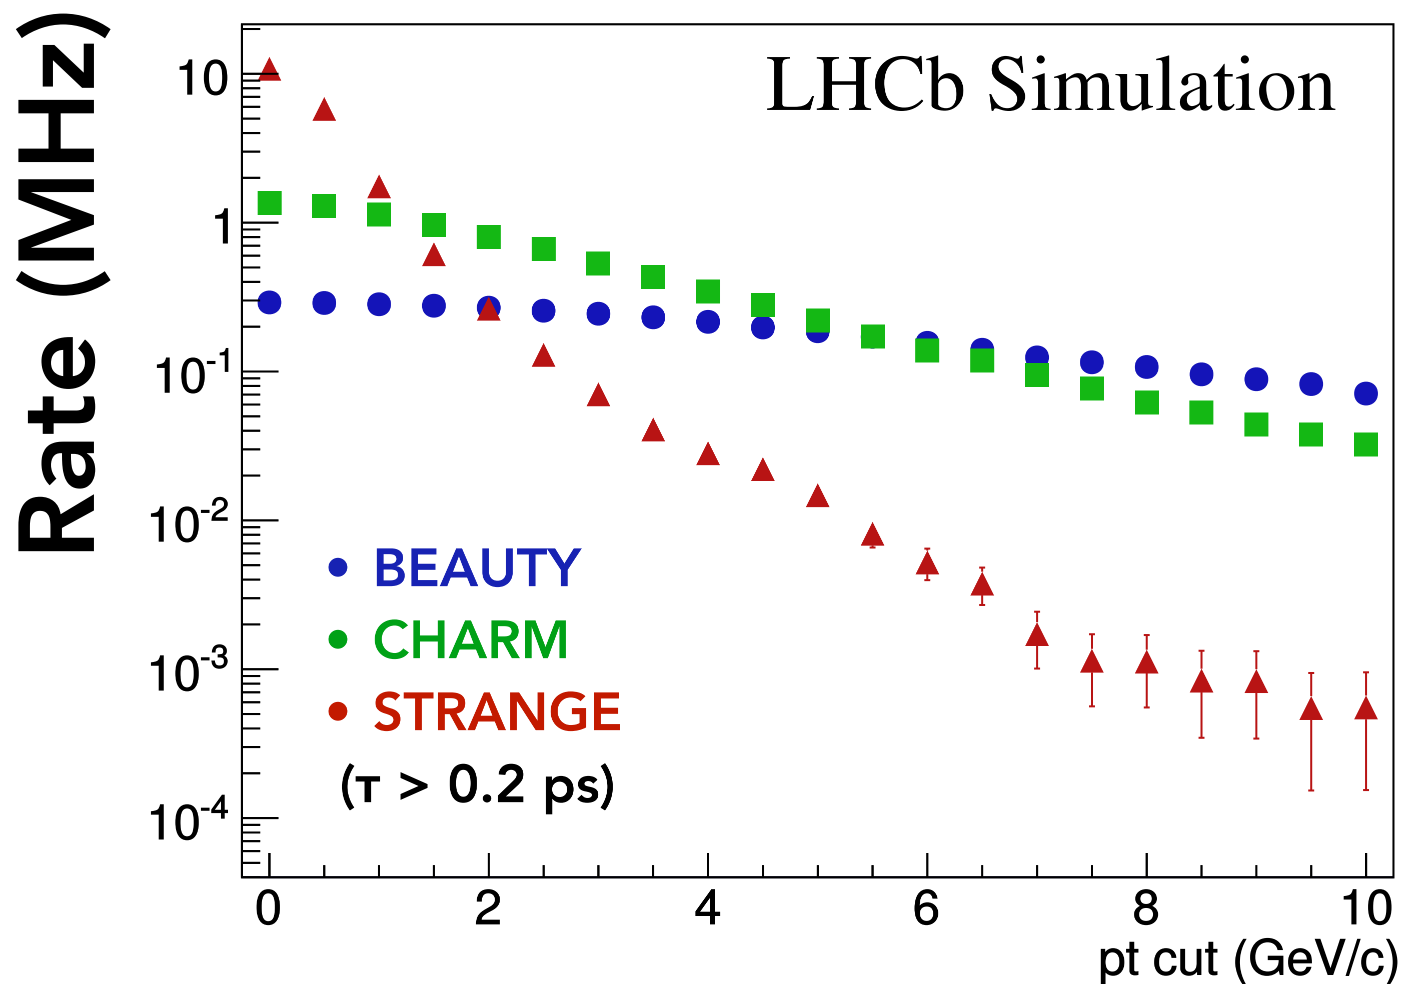

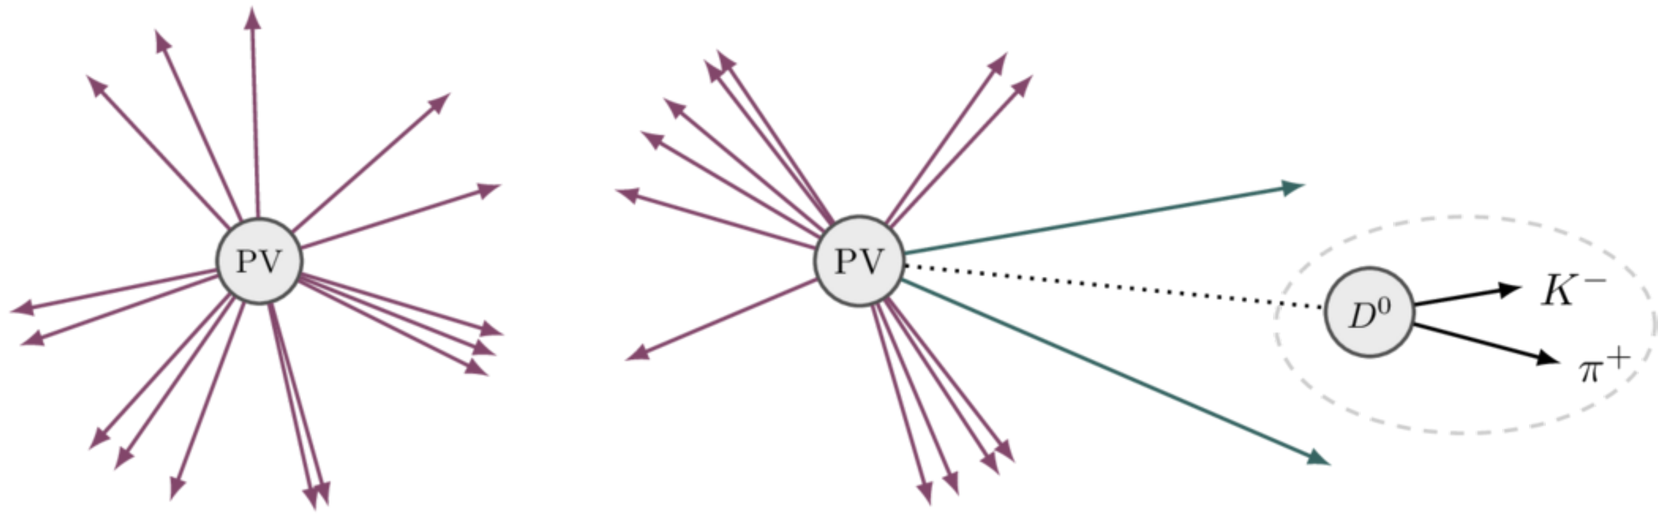

Need to reduce data rate by ~30x while keeping most of the reconstructible charm and a subset of softer physics. If you can do that the b-physics and everything else is "for free"

Key signature is a secondary decay vertex with significant transverse momentum and displacement from the primary *pp* collision. Displacement information is mandatory.

Require charged particle reconstruction at 30 MHz in the full detector

Rough extrapolation of processing complexity to FCCee

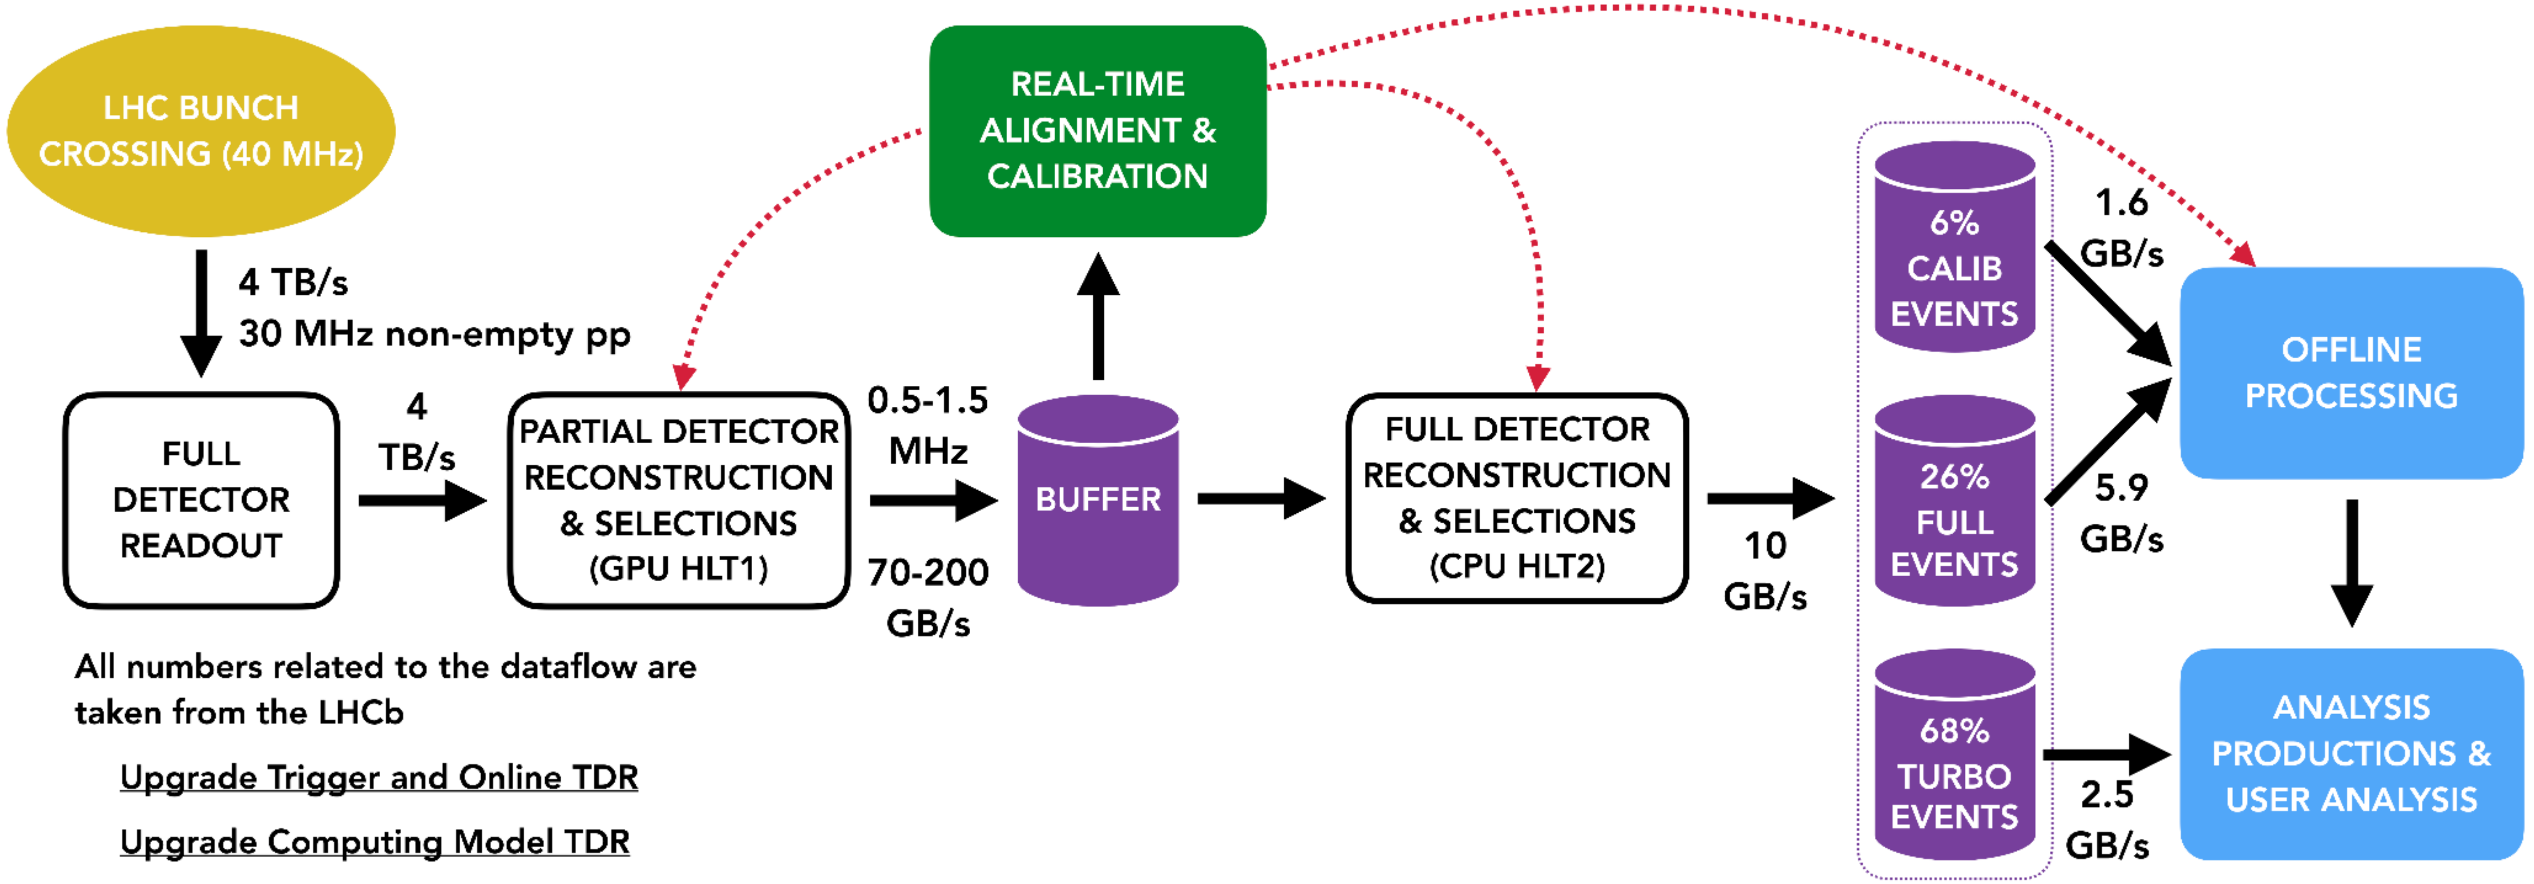

All numbers related to the dataflow are taken from the LHCb

Upgrade Trigger and Online TDR

Upgrade Computing Model TDR

Our full offline processing is performed on O(1) MHz of the most interesting events selected by the first-level processing, and has O(100) times the computational cost.

Even on a pessimistic 10% per annum (HL-LHC experiments typically assume 15-20%) extrapolation, computing technology developments will give 10x price performance gains on FCC-ee timescales. Ergo can process O(10) MHz of complex events (meaning ~5 pp collisions, order 100 charged particles) with ultimate fidelity in real-time without doing anything "intelligent" by the time FCCee starts. More than enough.

LHCb analysis methodology and role of calibration samples

Trigger Efficiency

Tag-and-probe calibration method exists & widely used

Tracking efficiency

Tag-and-probe

Existing	Developing
μ	e,π,K,p

Particle identification

Tag-and-probe

Tag-and-probe calibrations exist for all charged particle species and for π^0/γ, with new sources added over time to improve coverage

Data driven efficiency calibration key to precision physics

How do we align and calibrate our detector in real time

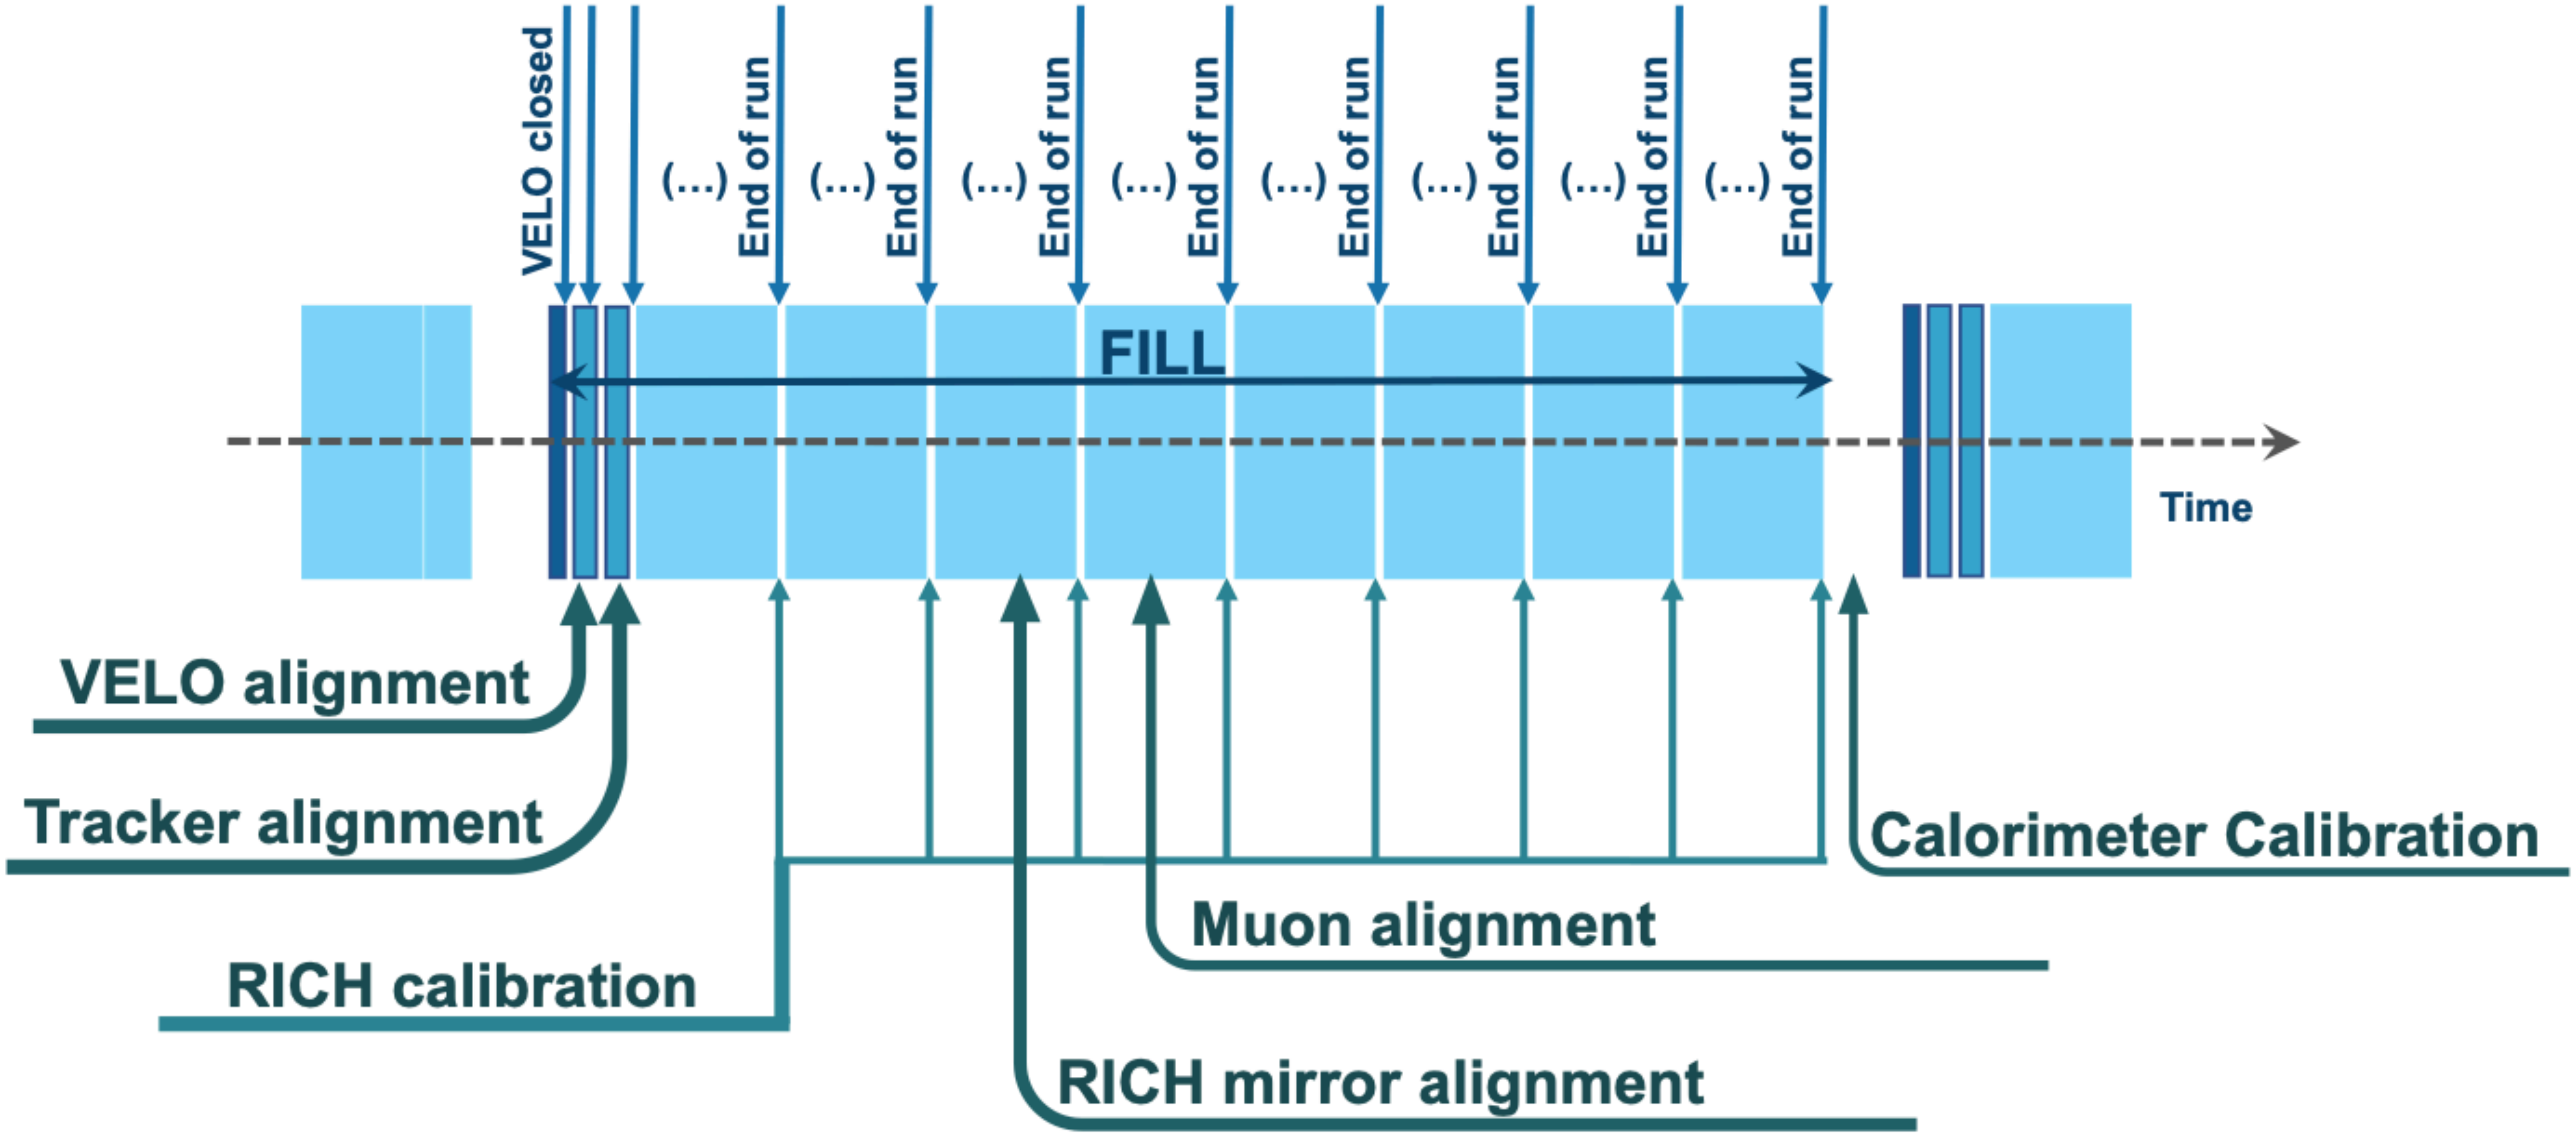

Early 2024 alignment and calibration results

Vertex detector residuals after alignment

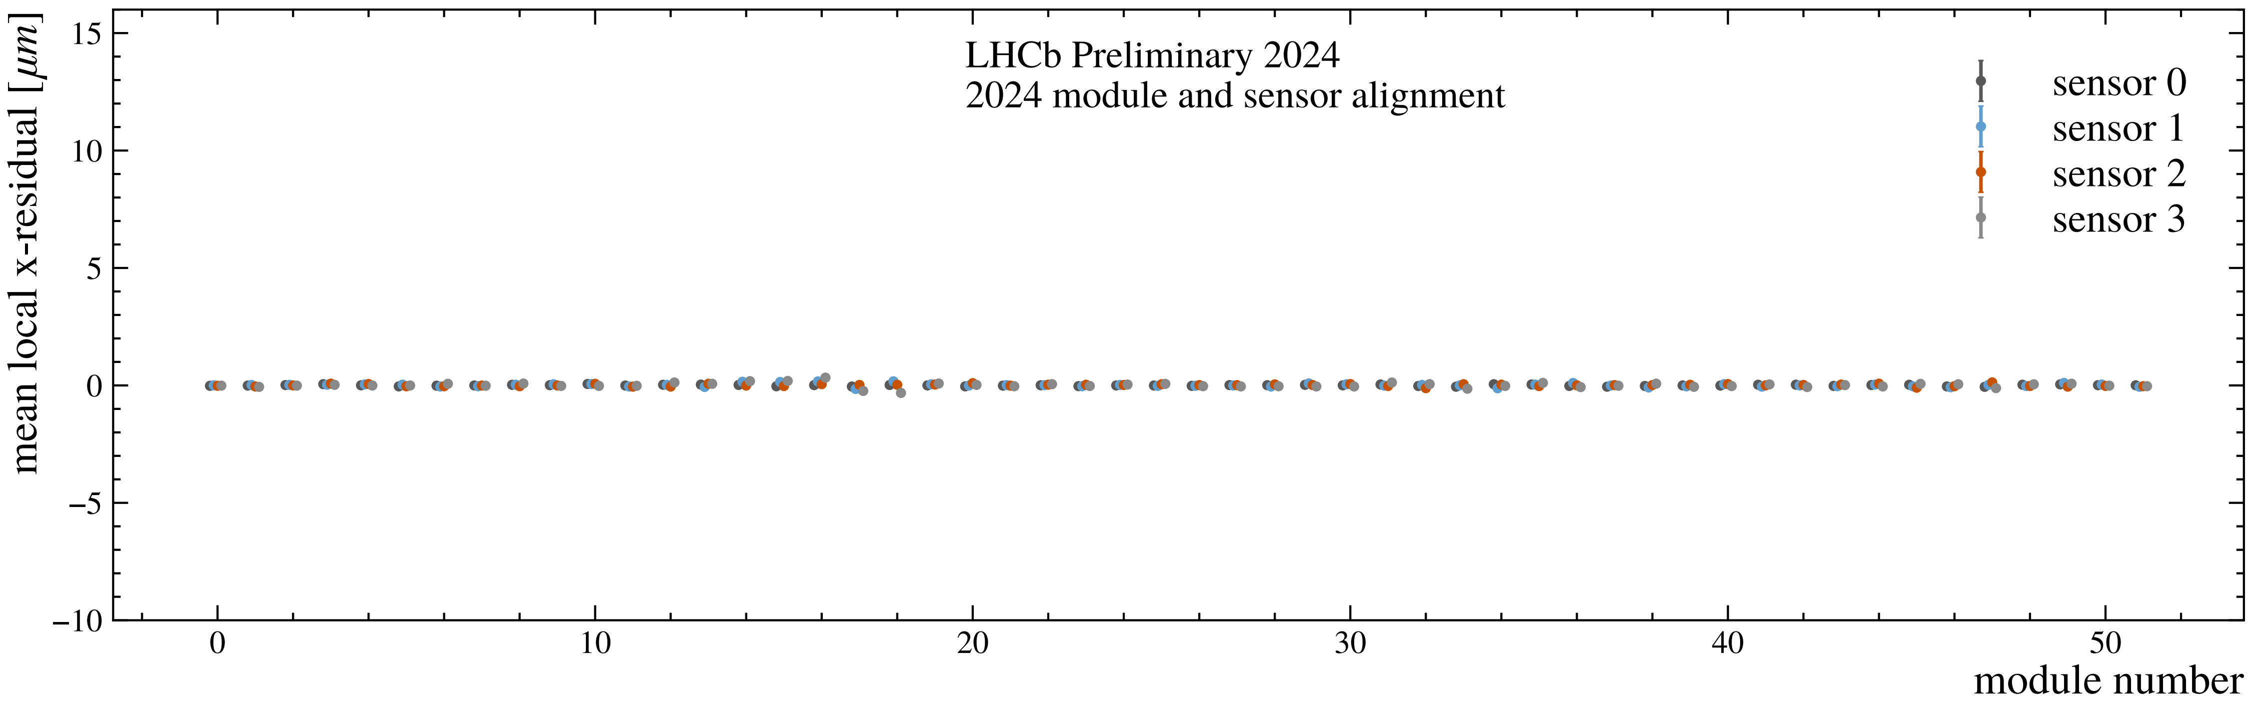

Impact of tracker alignment on track χ^2

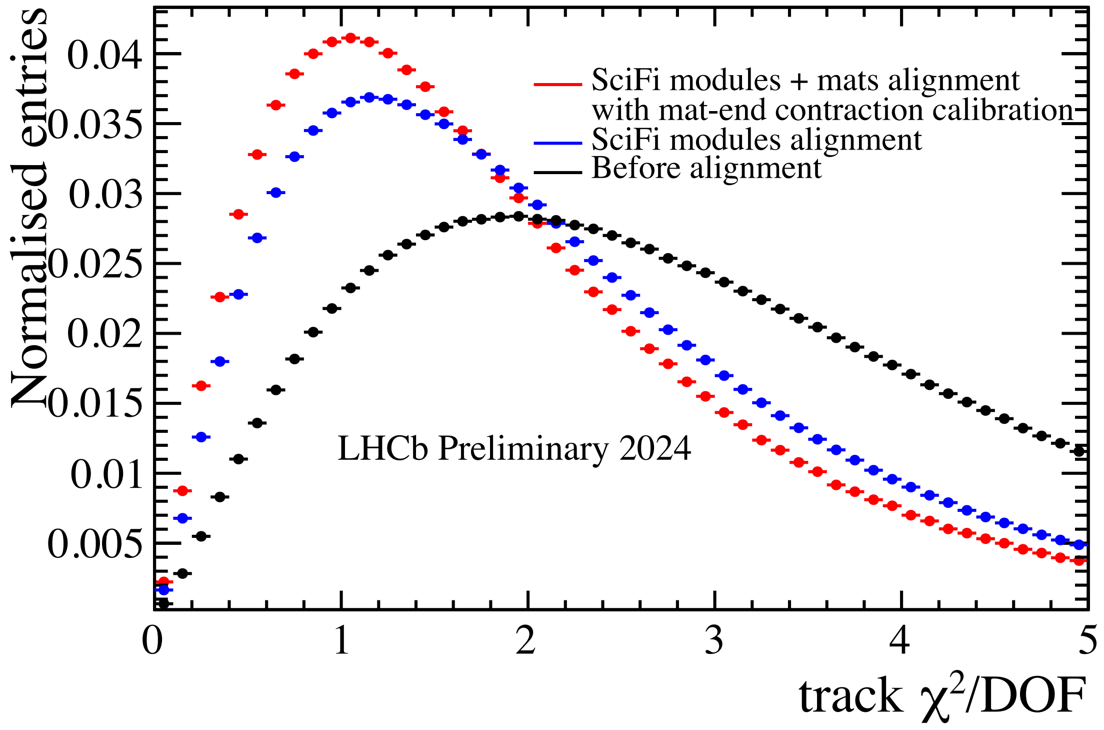

π^0 mass after calorimeter calibration

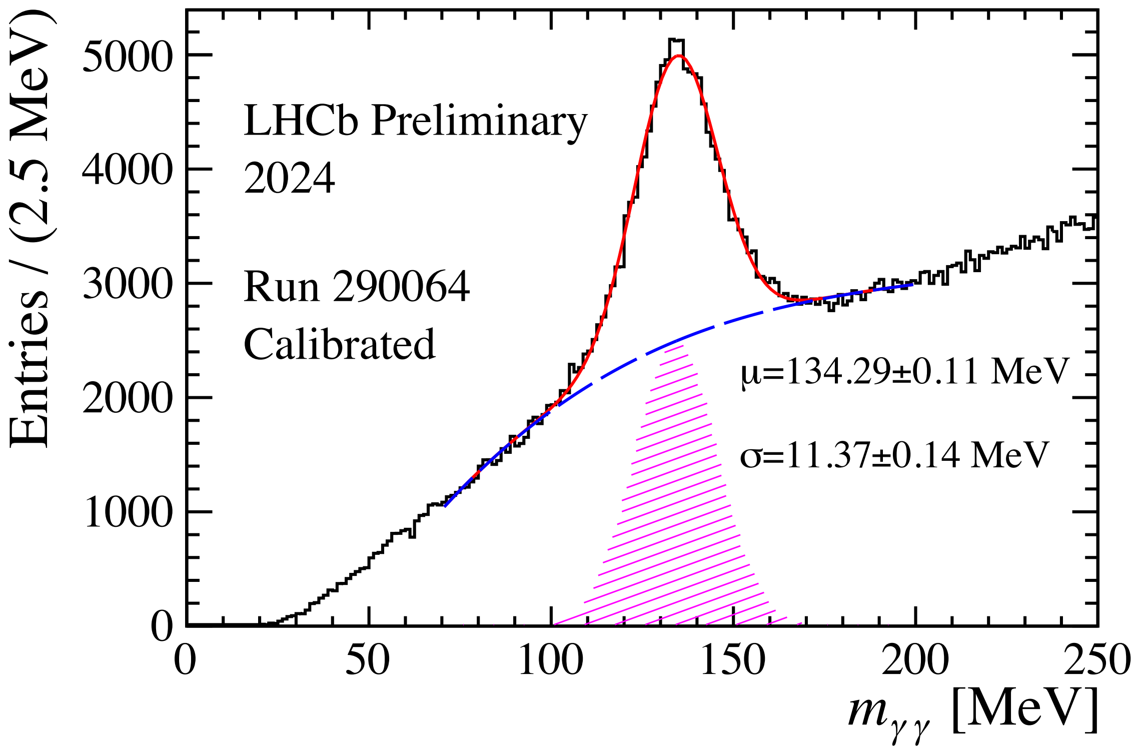

Alignment making good progress in 2024, hope to achieve nominal performances soon

Early 2024 reconstruction performance results

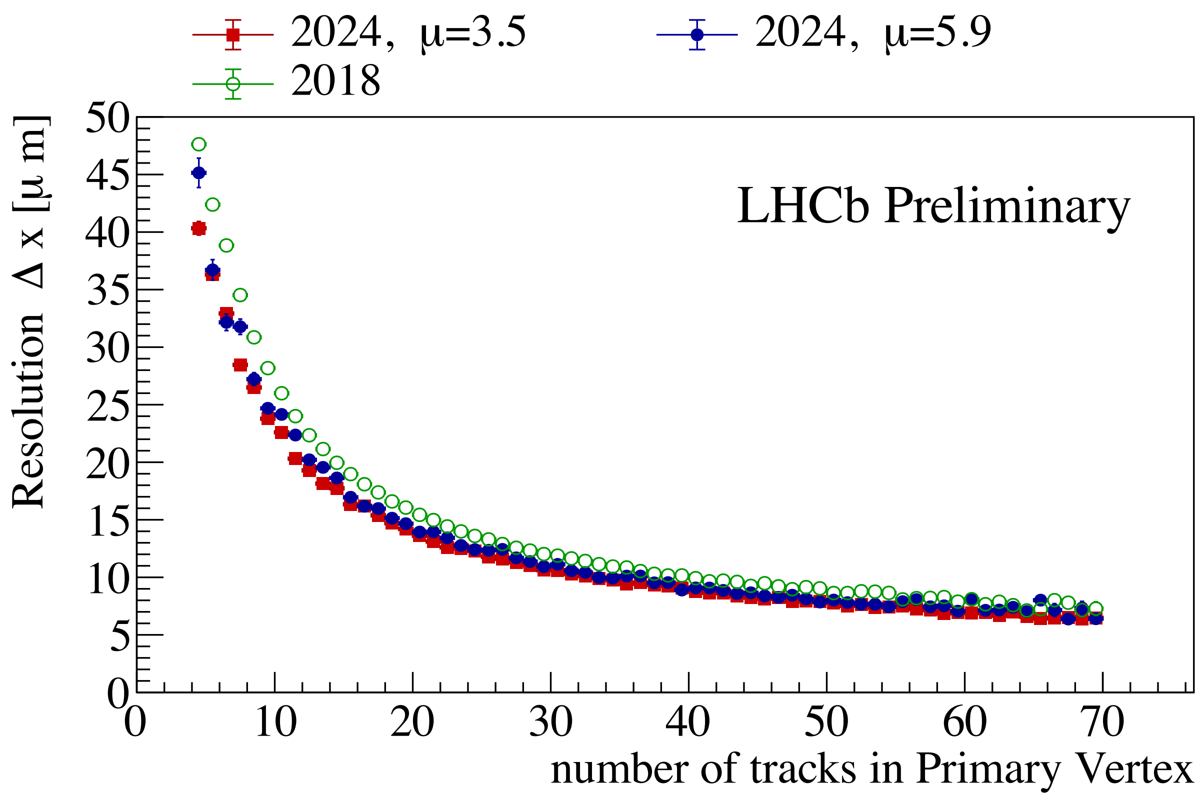

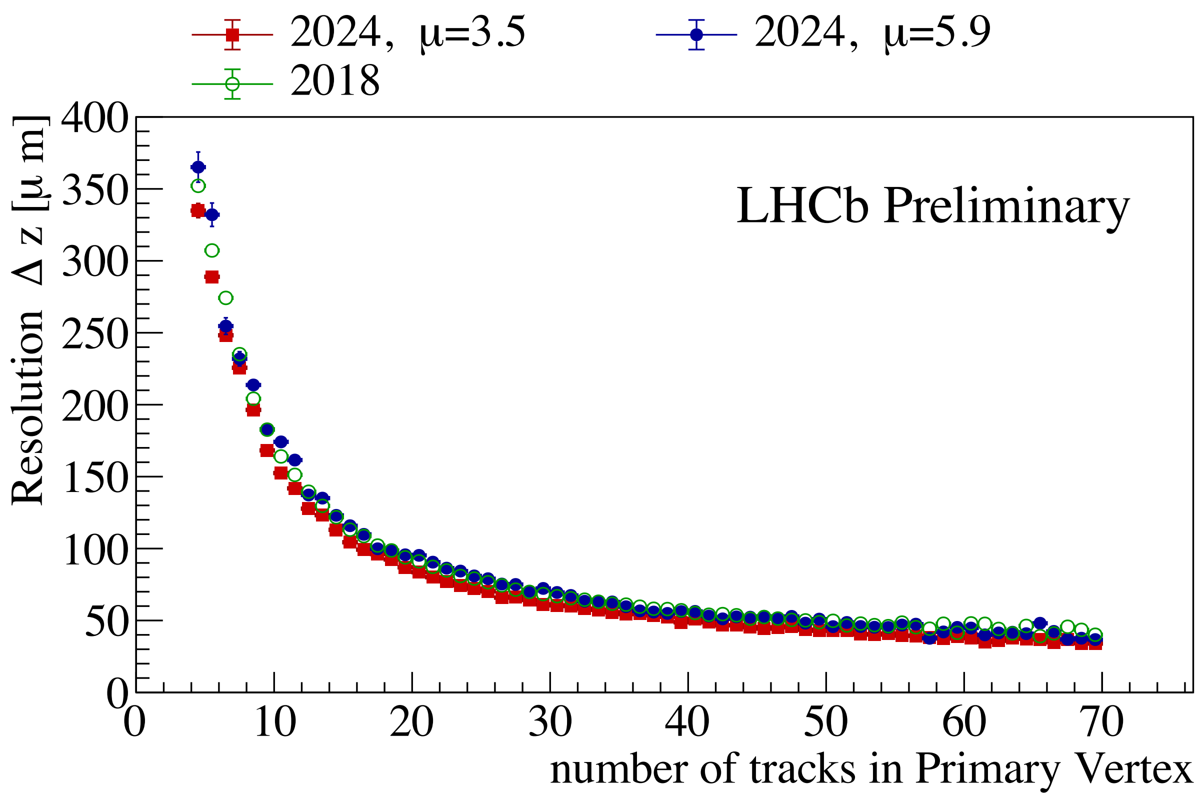

Primary vertex resolutions better than Run 2 and as expected from detector configuration

Early 2024 reconstruction performance results

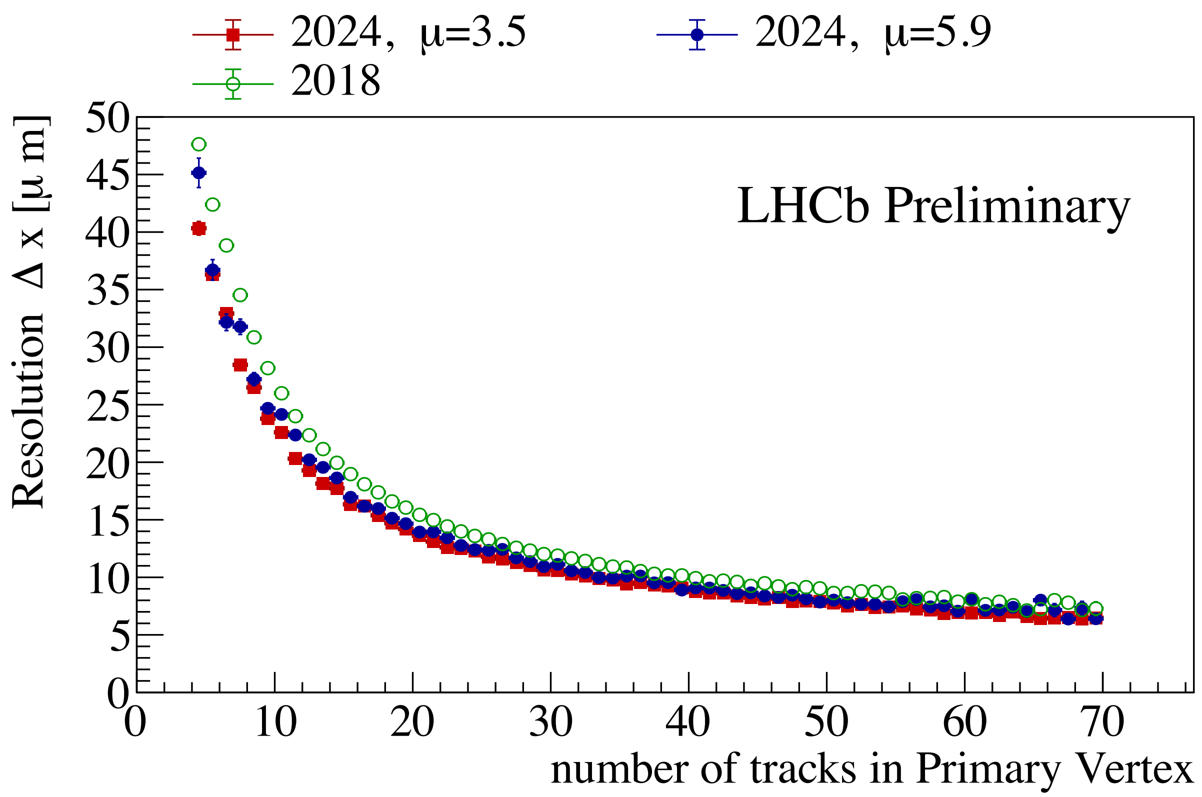

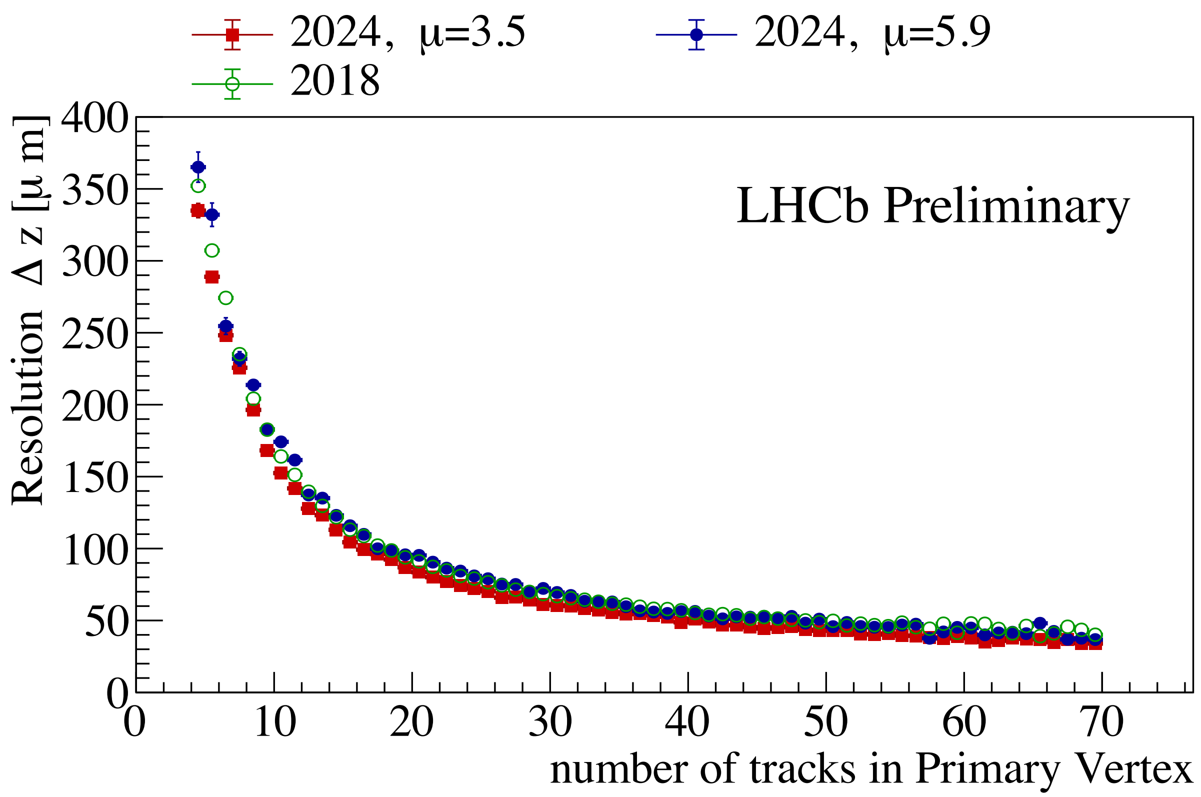

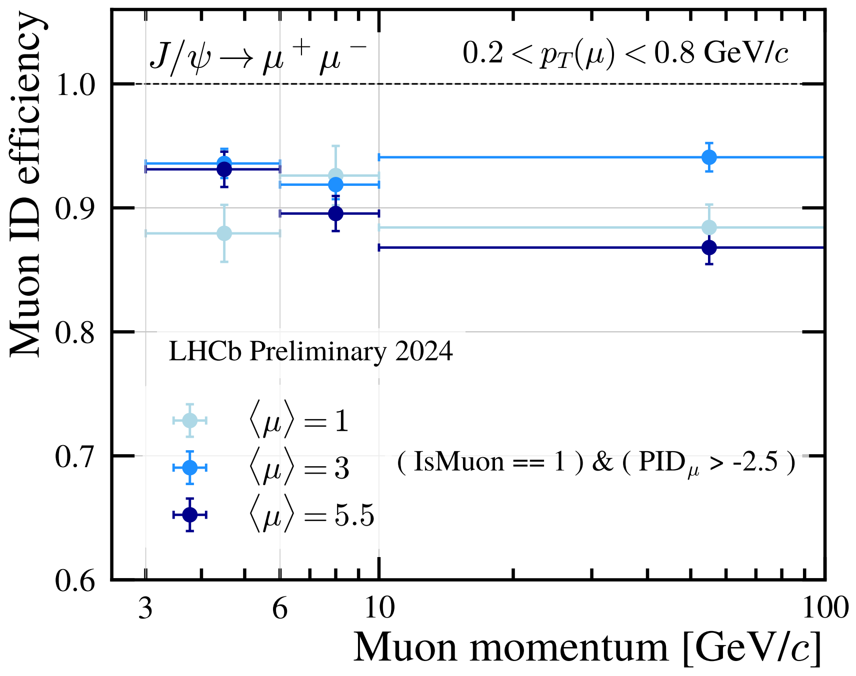

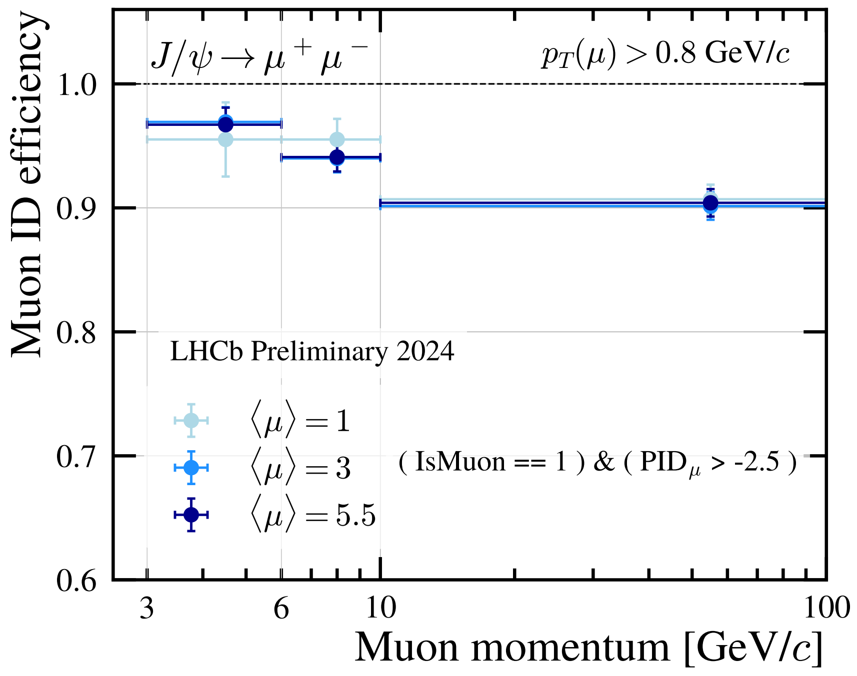

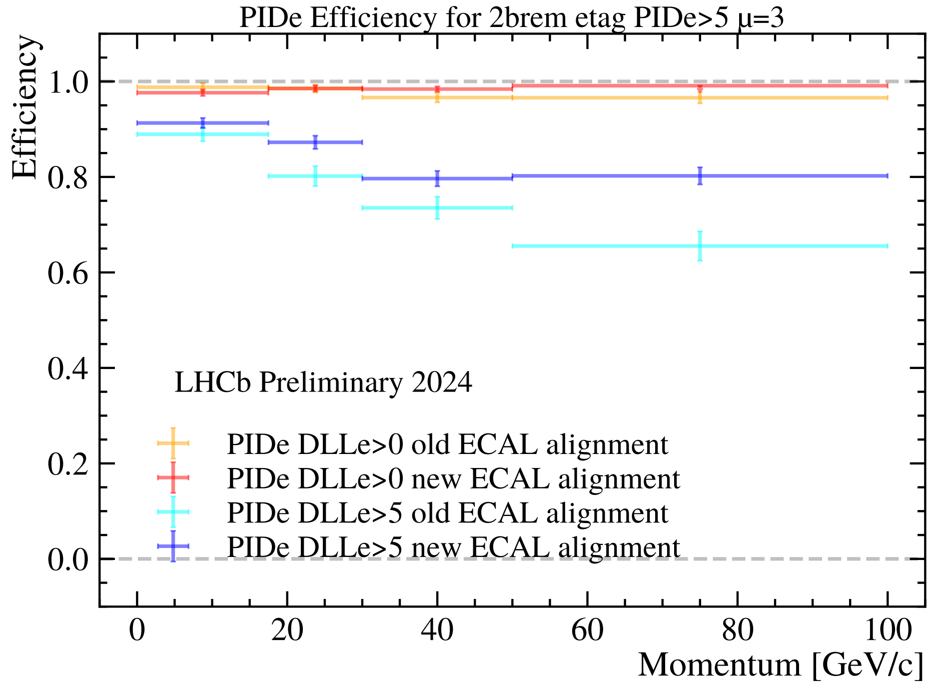

Muon and electron identification performing as expected

Early 2024 trigger performance results

$B^0 \to D\pi$

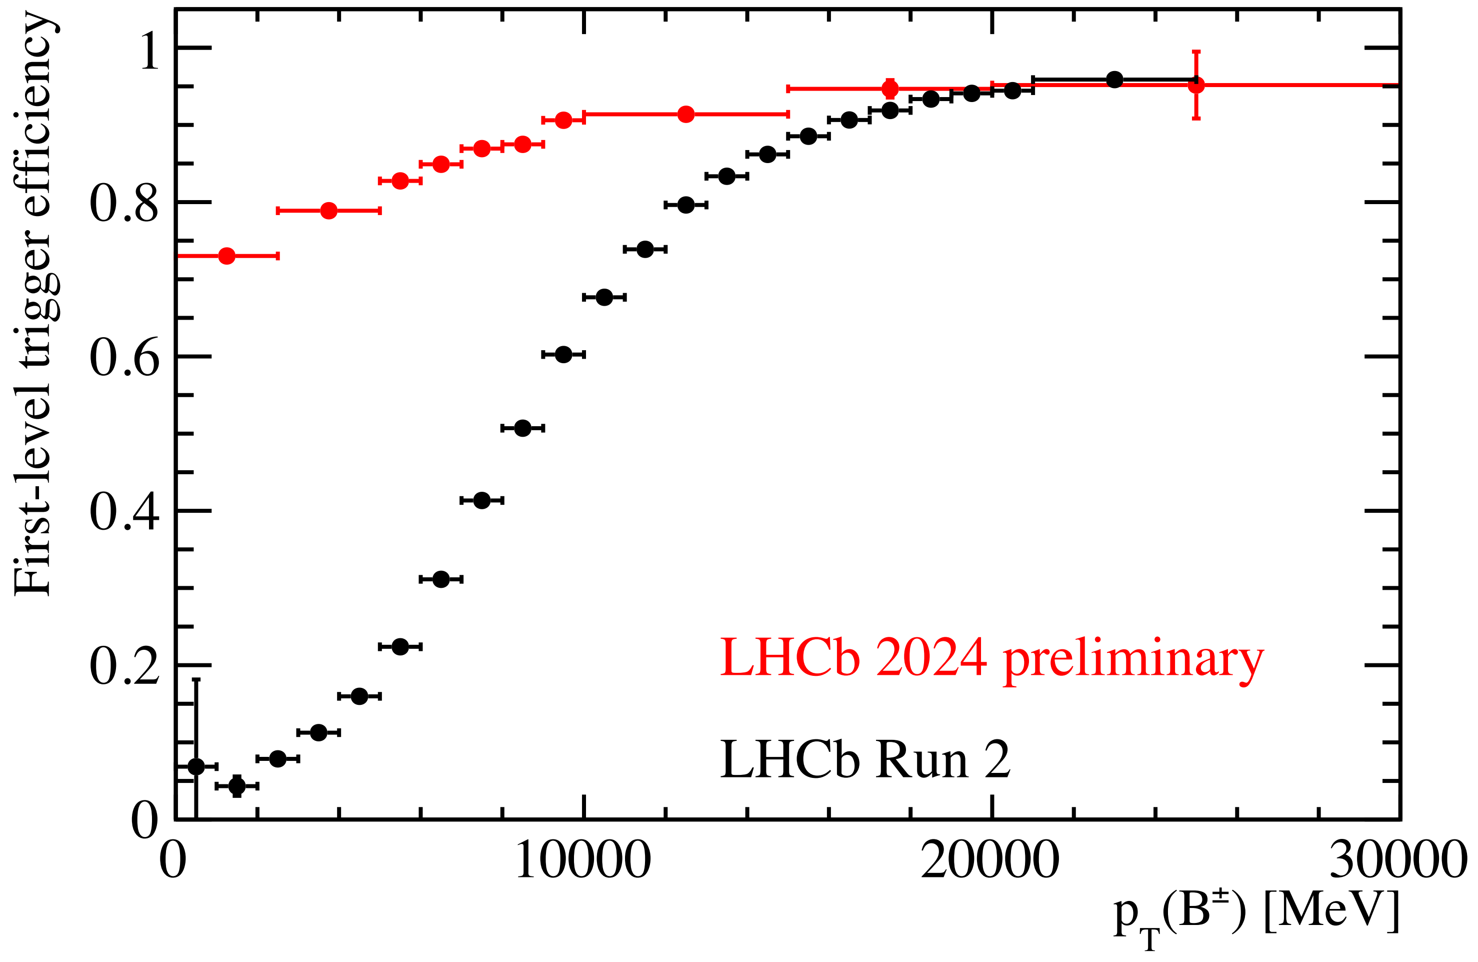

$D^0 \to K\pi$

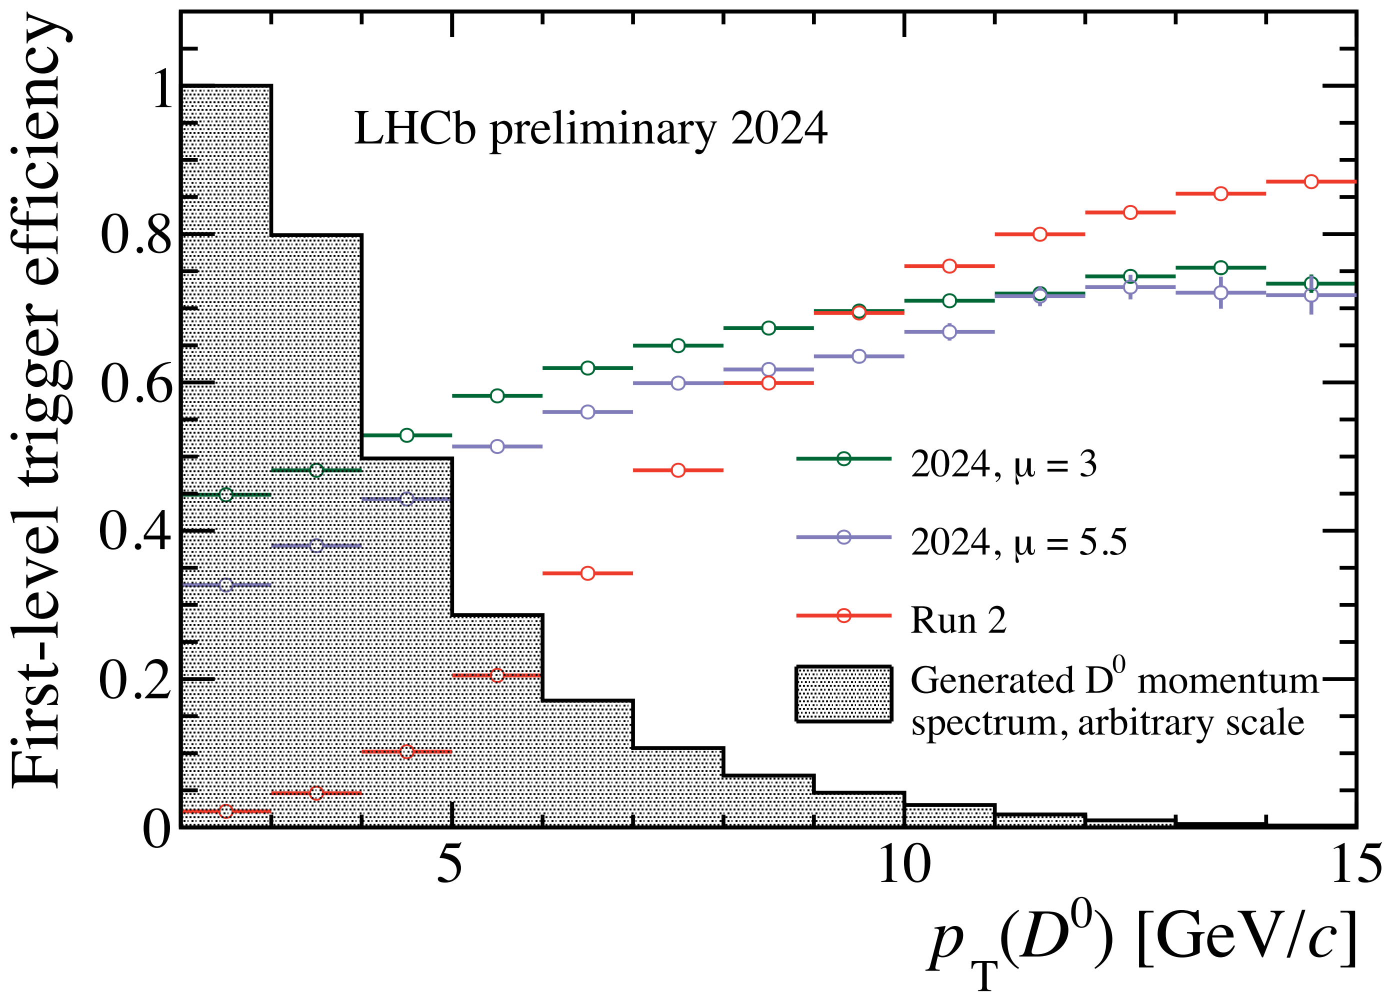

Software trigger gains more than x2 yields per unit luminosity for hadronic final states. Ability to reconstruct the detector at all occupancies also means significant gains for many muon channels

Other benefits of COTS solutions

Although the LHCb DAQ and trigger design is fundamentally driven by physics, the choice to use COTS elements wherever possible comes with benefits

- Relatively low learning curve for newcomers
- Ease of maintenance
- Possiblity to reuse code and even computing architectures (if desirable) for online, offline, and physics analysis data processing steps
- Can upgrade without redesigning, so continuously benefits from external technology progress

Considering the pace of computing technology development outside HEP, the third advantage is a particularly attractive one compared to locking ourselves into bespoke solutions many years before the experiments actually have to take data.

Conclusion

The upgraded LHCb detector has successfully implemented a nearly triggerless readout in which all detector information is processed by a data centre consisting of heterogeneous processing units

The system eliminates latency as a consideration and within the limit of using the same architectures online and offline (currently not the case, but a matter of choice) enables maximum reuse of processing code and overlap between real-time and offline data processing

This is an inherently scalable solution for any experiment which does not physically require a hardware trigger (e.g. as the HL-LHC experiments require because among other things of material budget considerations) which we may evolve towards a truly streaming/triggerless readout in the future if that proves to be the best way forward.

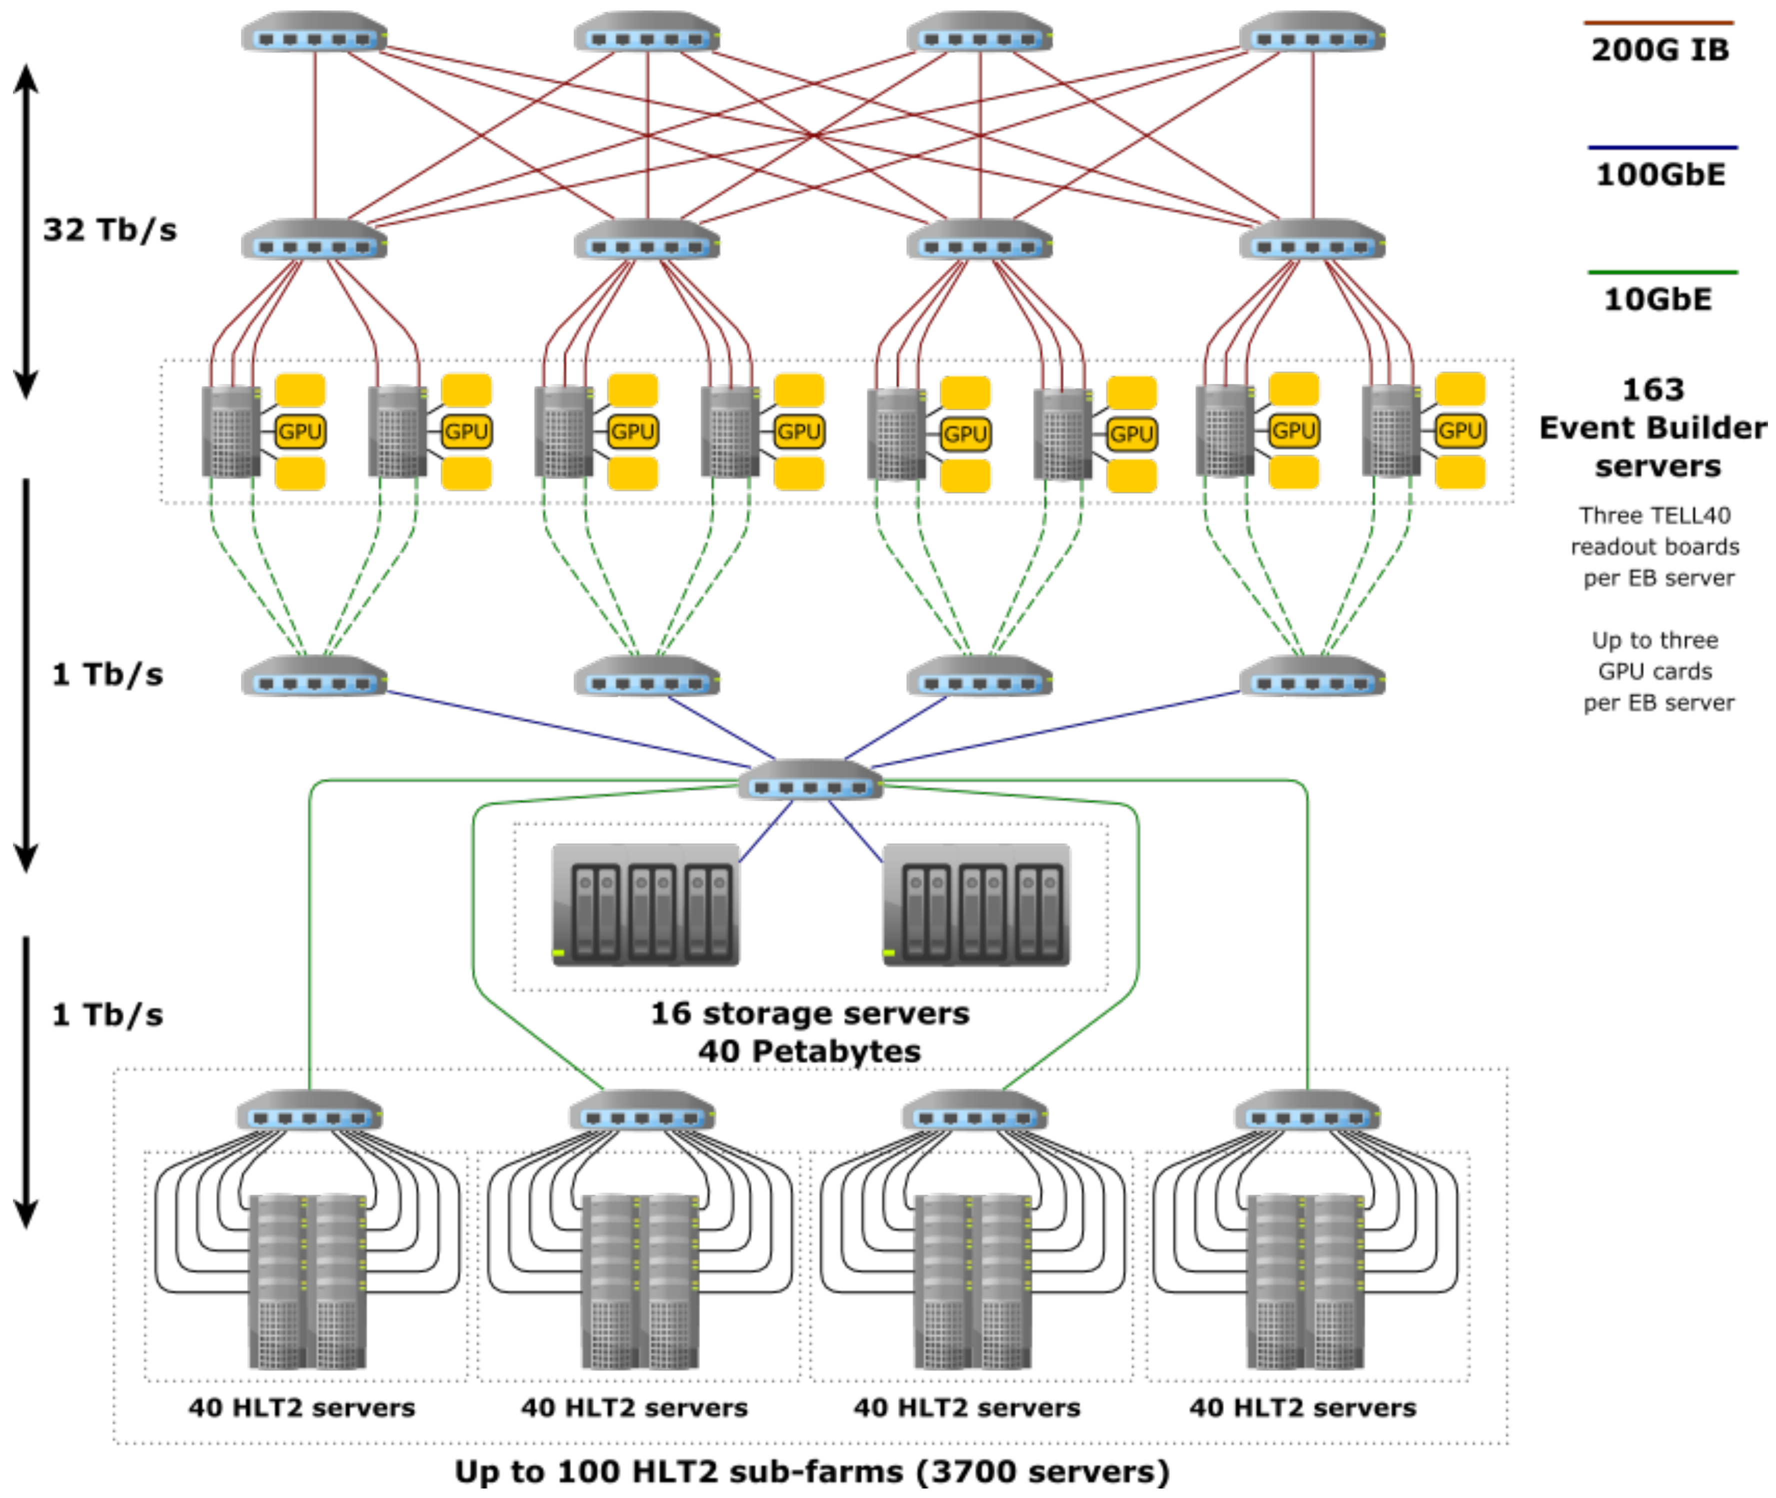

LHCb DAQ & Trigger could be an interesting model for FCCee experiments

Backup