ACAT 2024, 14 March, Stony Brook University (NY)

ATLAS TDAQ upgrades for Phase-2

F.Pastore (RHUL) on behalf of the ATLAS Collaboration

The High-Luminosity LHC

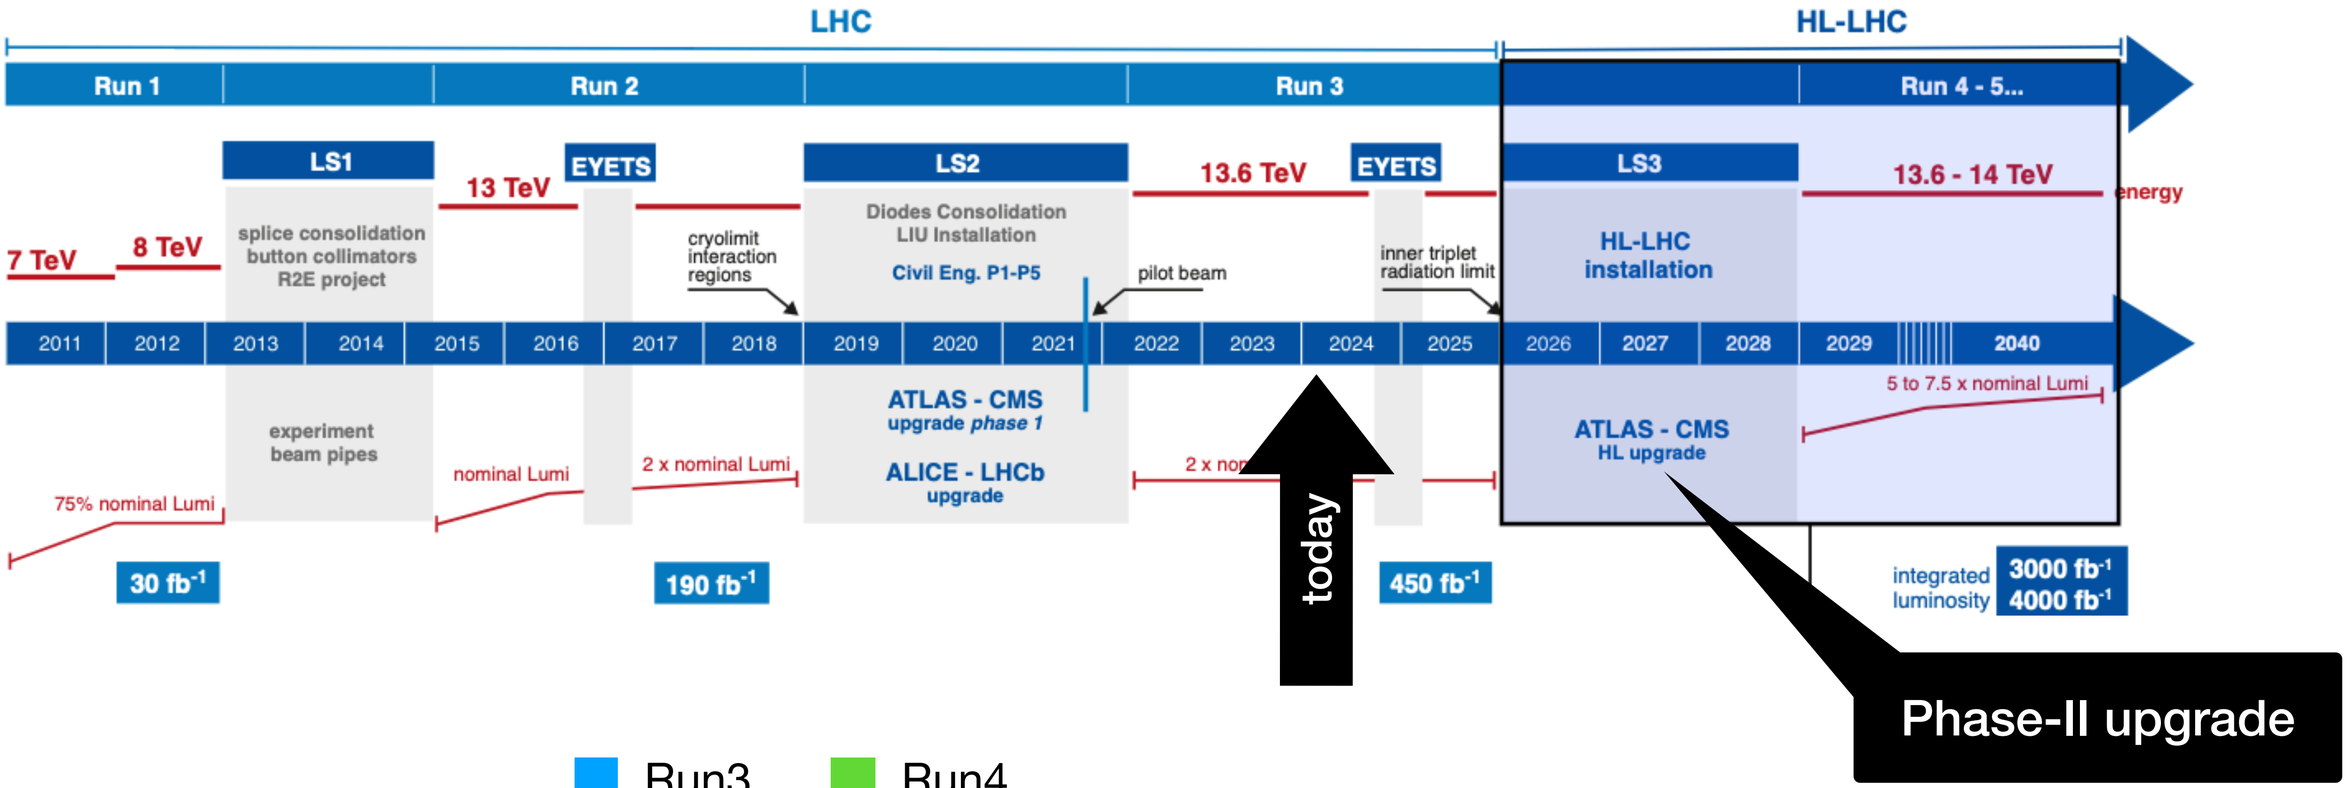

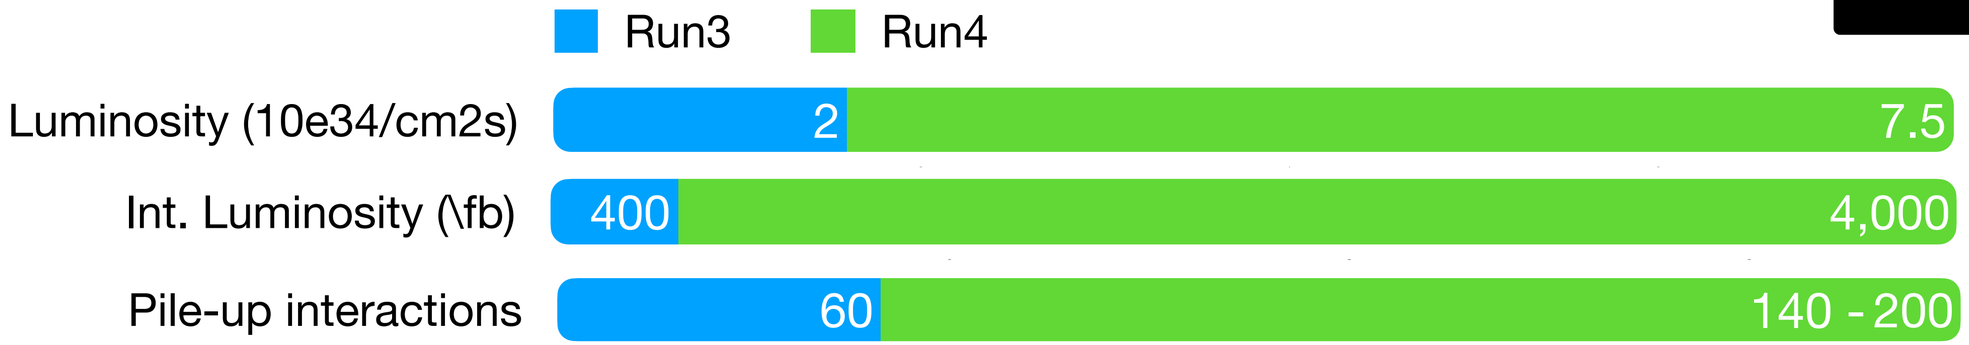

ATLAS at High-Luminosity L = 7.5×10^{34} /cm^2/s

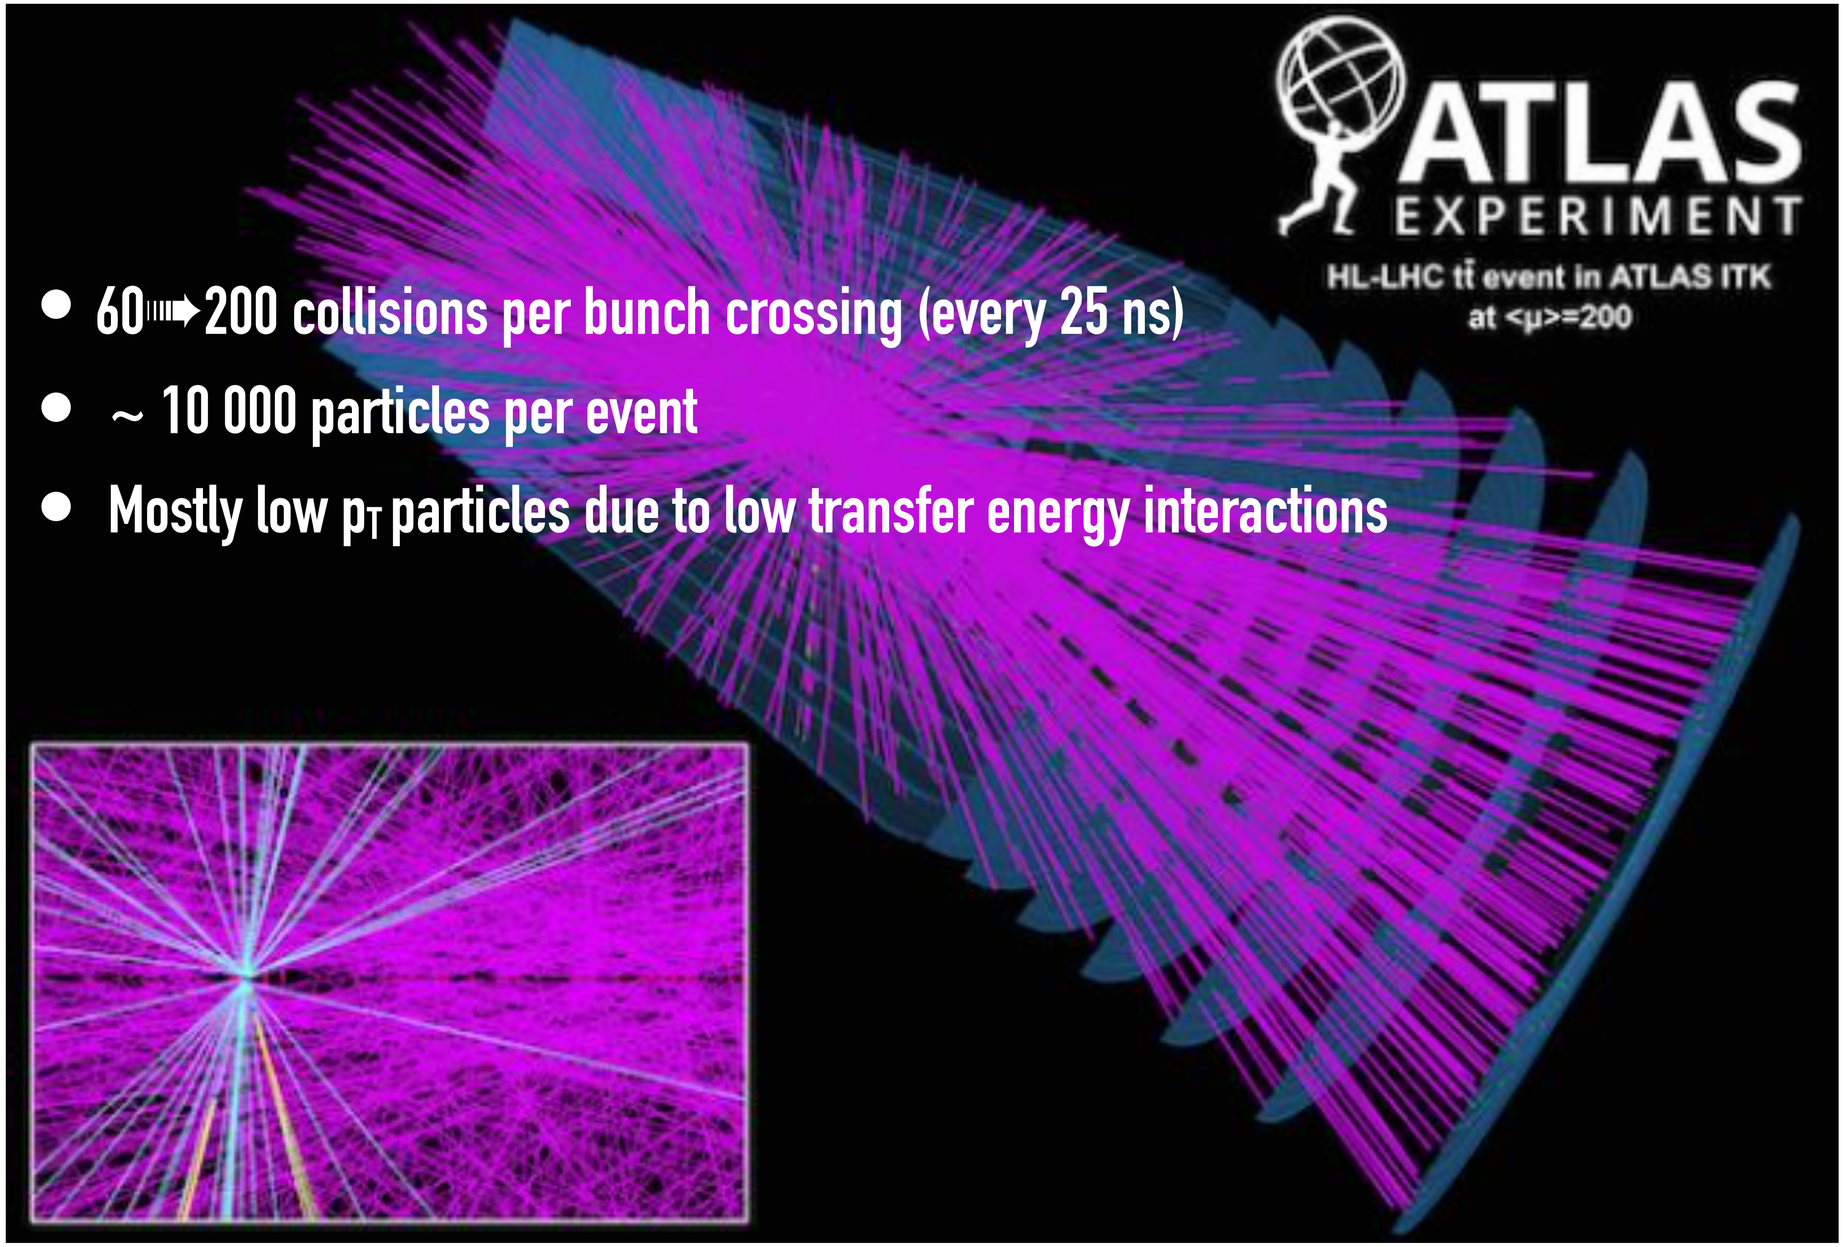

ATLAS for Phase-II

New Front End and Back End electronics for calorimeter and trigger

New full silicon Inner Tracker (ITk) - sensors/mechanics/Front End

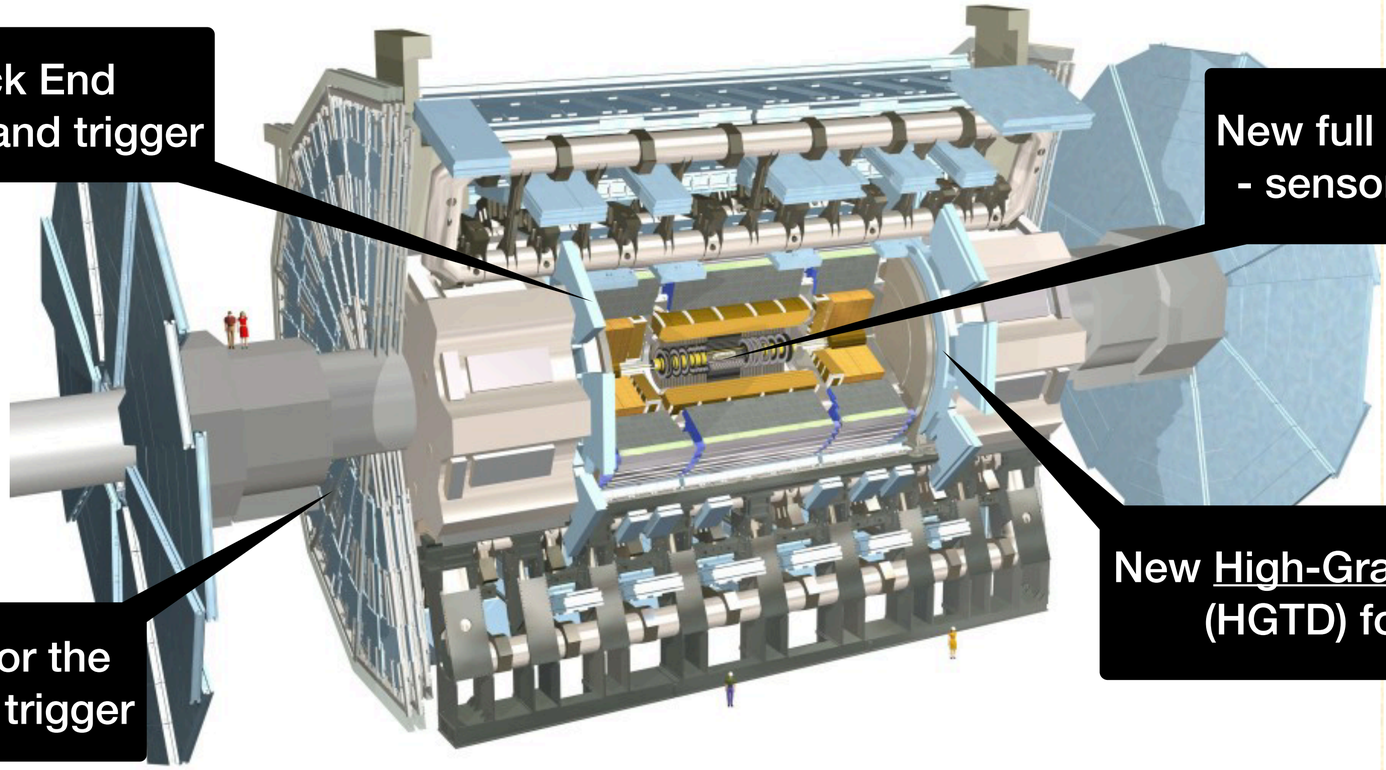

New High-Granularity Timing detector (HGTD) for pile-up mitigation

New Front End electronics for the muon spectrometer and new trigger

- Upgrades of detectors to be more robust against pileup
- Upgrades on Front End (FE) electronics towards new hardware triggers
- Completely new Trigger/DAQ architecture (this talk)

Changes in T/DAQ

- **More data: physics gain exploiting higher LHC Luminosity through increasing the trigger rates**
  - The hardware trigger rate is a limit already in Run 3
  - DAQ throughput increase, due also to larger event size (~2 to 5 MB/event)
- **Better data: increase trigger selectivity, to set lower thresholds than today, with longer processing time, due to event complexity**
  - Longer hardware trigger latency (⟹ new readout/DAQ architecture)
  - Software trigger more robust against pile-up

ATL-COM-DAQ-2016-059

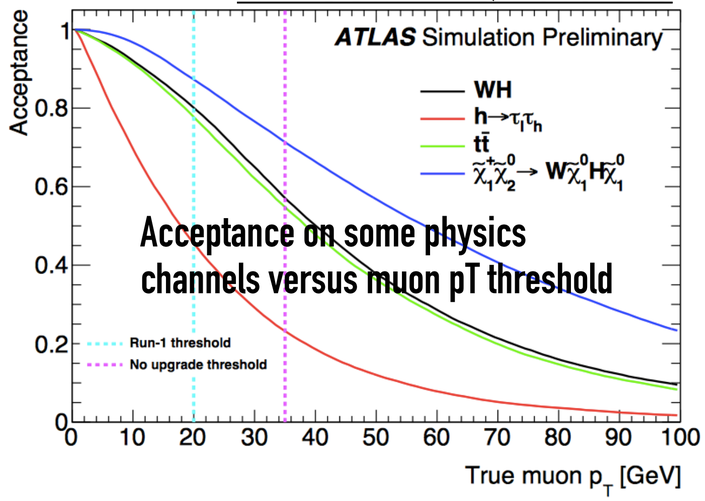

ATL-COM-DAQ-2015-148

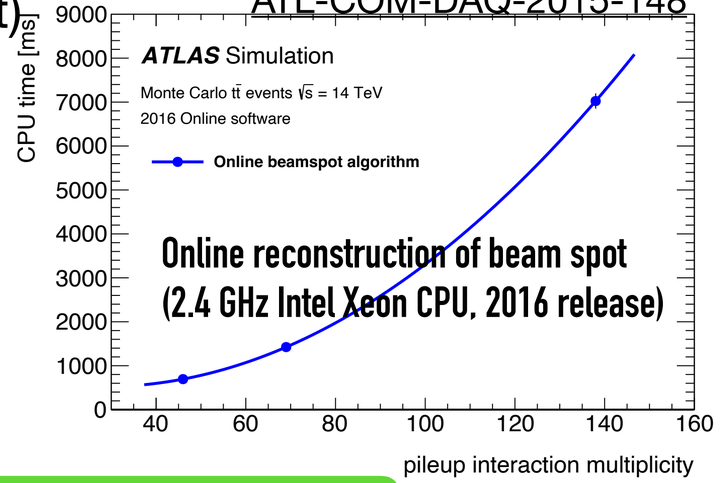

ATLAS TDAQ architecture for Phase-II

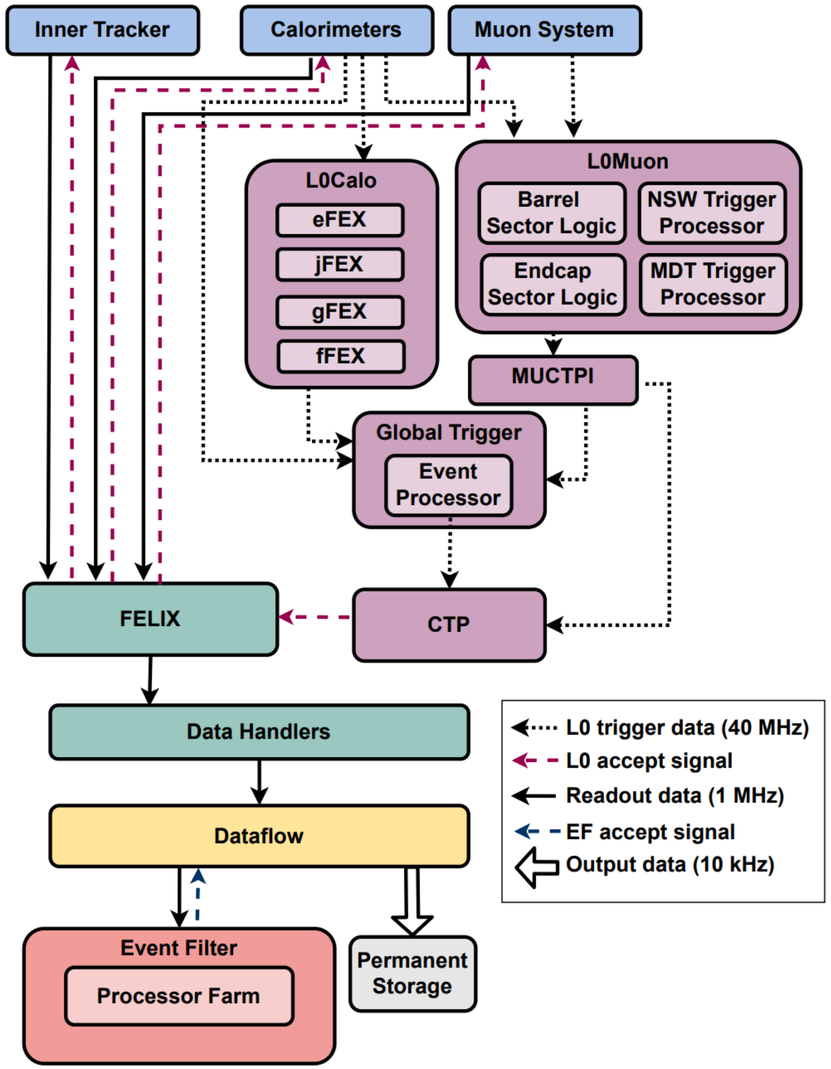

- **Level-0 (L0) trigger (hardware)**
  - 40 MHz ⟹ 1 MHz
  - 10 μs latency
- **Event Filter trigger (software)**
  - 1 MHz ⟹ 10 kHz
- **DAQ: Readout and Dataflow**
  - 5 TB/s data throughput

- **Technical Design Report (TDR) in 2017 [ATLAS-TDR-029]**
- **TDR Amendment in 2022 [ATLAS-TDR-029-ADD-1]**
  - Re-design of the Event Filter architecture based on COTs

L0 trigger for Phase-II

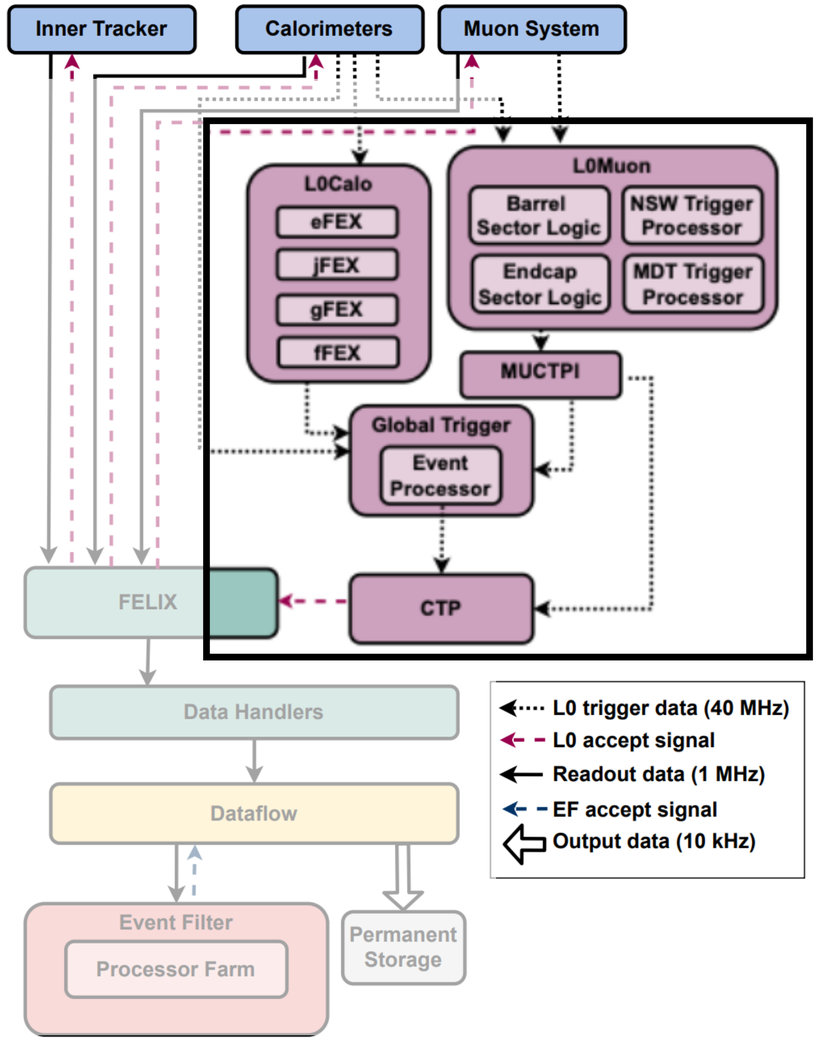

- Identify **muon** and **calorimeter** L0 trigger objects (**TOBs**) in local regions (Region-of-interest)
- TOBs are collected into the **Global Trigger** which implements full granularity offline-like reconstruction @40MHz
- **Central Trigger Processor (CTP)** provides hardware trigger decision, sent to the Readout and the Event Filter
- Technology for custom build components
  - ATCA-based architecture
  - FPGA processing
  - data I/O with optical links 2.5-25 Gb/s

L0 Calorimeter trigger

- **Dedicated boards to identify calorimeter objects, improved resolution thanks to full granularity calorimeter information (for both EM and HAD)**
  - additionally improved by use of Global Trigger board (see next)
- **Four boards for Feature Extraction**
  - e(lectron)FEX, j(et)FEX, g(lobal)FEX already installed in Run 3
    - Retaining hardware with significantly updated firmware
  - f(orward)FEX: new board for forward electrons |η|>2.5 and jets |η|>3.3
    - Schematic capture & layout completed, prototype soon ready

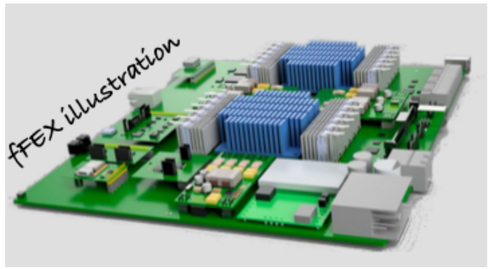

- **New optical connections to Global Trigger (fFOX, gFOX, FOX++)**
  - Custom designed, final design approval this year

Single-electron trigger rates TDAQ-TDR

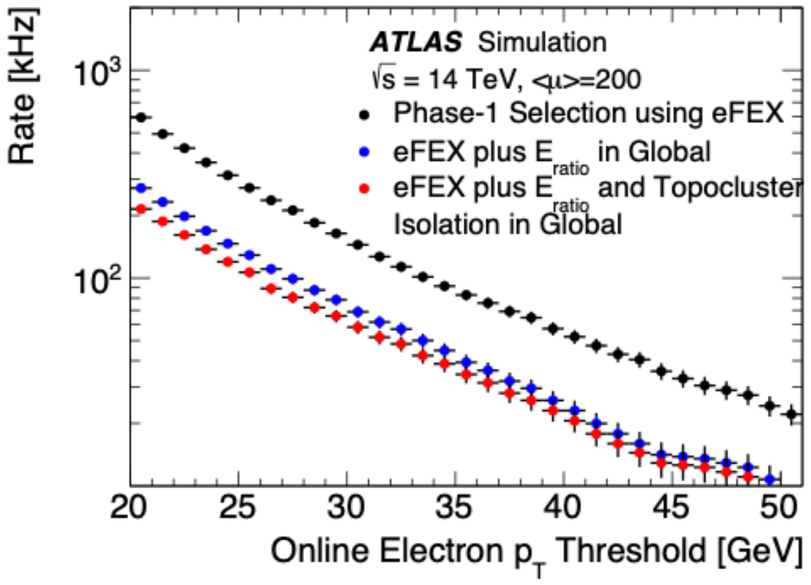

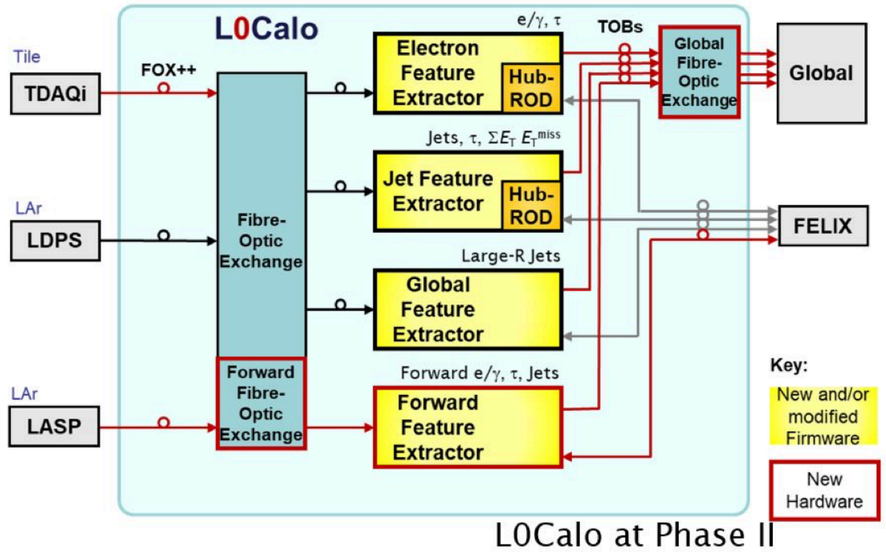

L0Calo at Phase II

L0 Muon trigger

- **Extended to include all muon detectors (in addition to legacy RPC & TGC & NSW)**
  - New trigger processor for endcap New Small Wheel (NSW-TP)
  - New trigger processor for MDT precision chambers (MDT-TP) with improved p_T resolution on muons identified by RPC and TGC
    - Final Design Review in 2023, pre-production has started
- **Data is streamed out and the entire trigger logic is moved off-detector (Sector Logic board)**
  - New common board for all technologies
  - 2nd prototype layout started
- **Integration tests in full swing (more connectivity to test)**
  - Including on-detector boards, MDT-TP, Sector Logic and readout (FELIX)

Single muon 20 GeV turn-on TDAQ-TDR

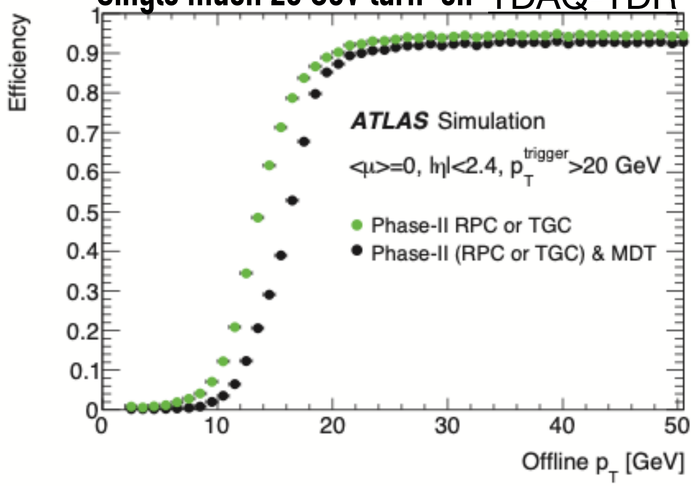

MDT-TP Command Module

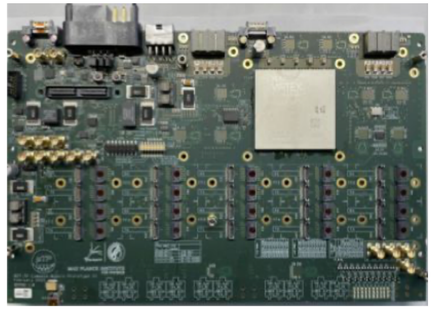

Sector-Logic prototype

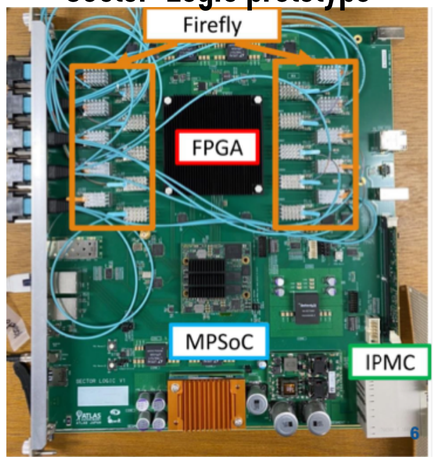

L0 Global Trigger

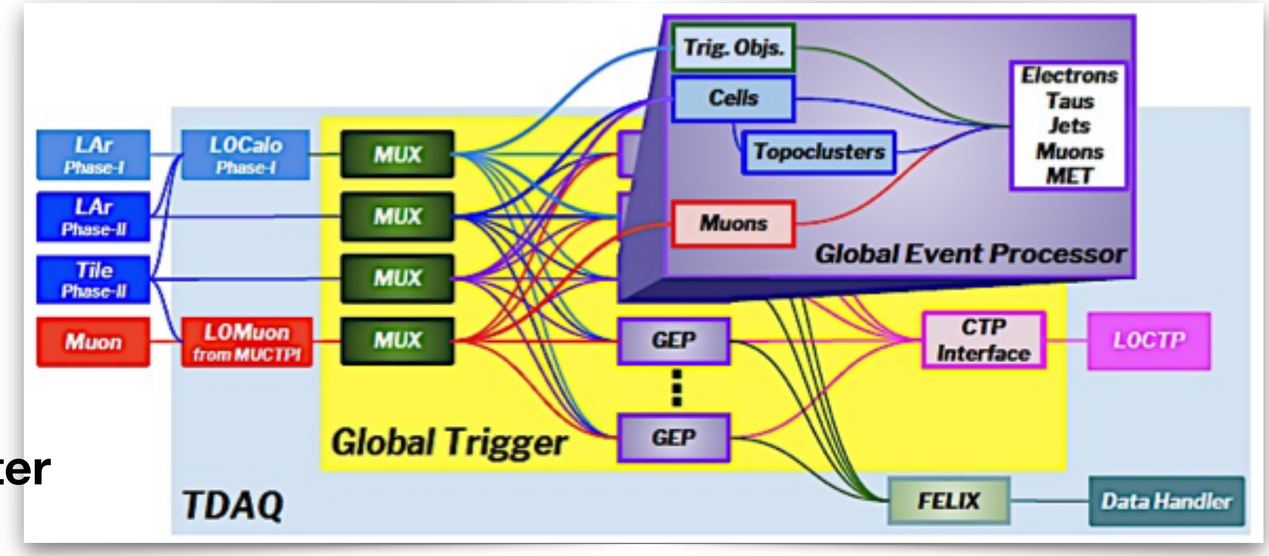

- **Process L0 TOBs + more high-granularity calorimeter**
  - Offline-like algorithms to refine identification of muons, calorimeters (topological-clusters), jets (anti-k_T) and pile-up subtraction
  - With topological functionality as in Run 3
- **Farm of boards with same hardware platform (Global Common Module) for different functionalities:**
  - Data aggregation and time-multiplexing per bunch-crossing (MUX)
  - Processing algorithms for each event in Global Event Processors (GEP)
  - Send GEP outputs to the Central Trigger Processor (gCTPi)
- **Status**
  - Preliminary design review of the board in 2023, early tests on prototype (2 Xilinx VP1802)
  - Firmware progressing, review passed for critical algorithms (tau and pile-up suppression)
  - Slice test (including MUX, GEP and gCTPi functionality) in fall 2024

L0 Global v3 prototype

L0 Central Trigger

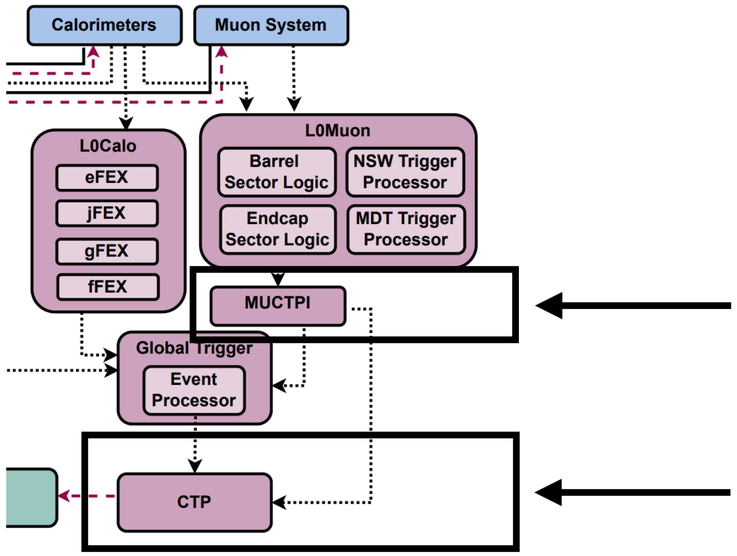

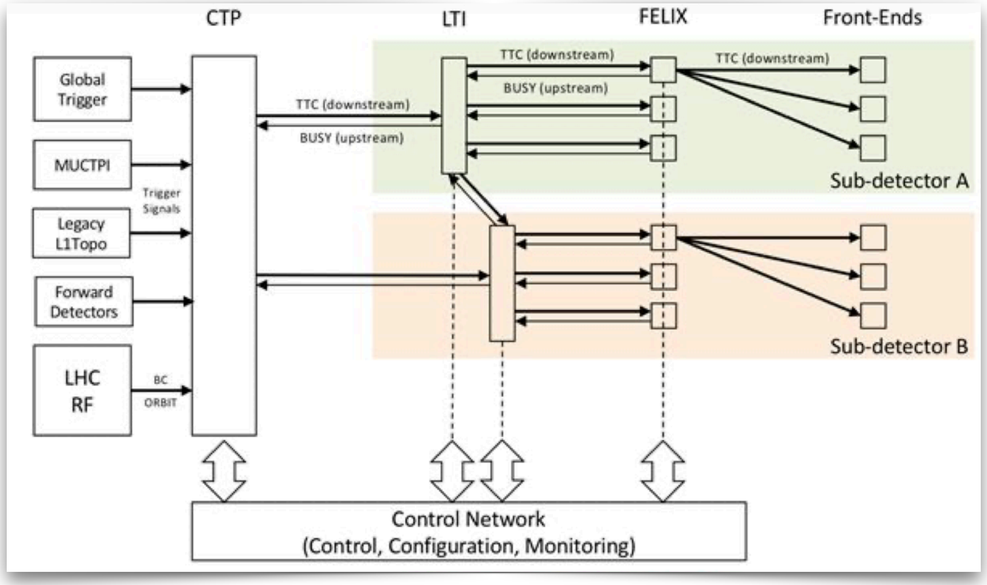

- **Takes final L0 decision, applying vetos and pre-scales, and drives the timing signals**
- **New Central Trigger Processor (CTP) board for**
  - More trigger inputs: 512 to 1024
  - More bandwidth, more complex criteria
  - Preliminary design expected this spring
- **Muon-to-CTP-Interface MuCTPI to remove overlaps between muons and calculate multiplicities**
  - Reuse Run 3 board with upgraded firmware
- **Trigger, Timing and Control (TTC) system network, distributed via new Local Trigger Interface (LTI) modules**
  - Preliminary design in 2023, new prototype underway
  - Tests on going of phase stability for UltraScale+ GTHe4

More details in A. Koulouris' talk

Readout

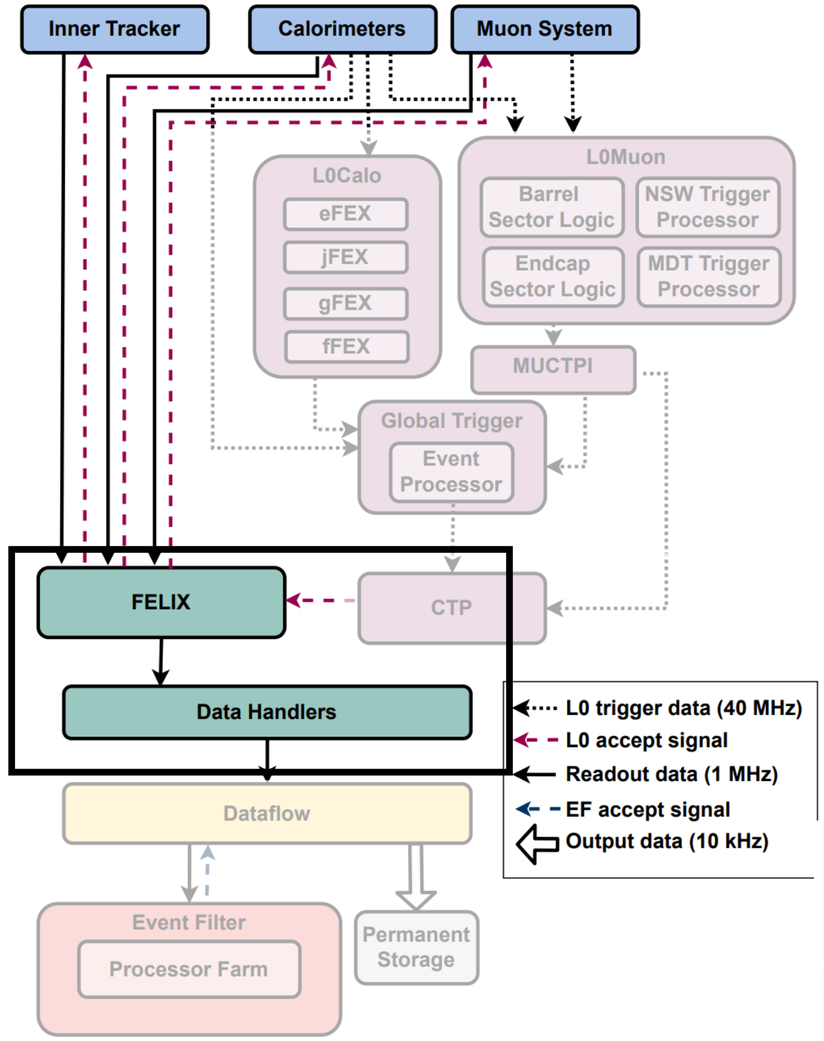

- **New common interface between detectors and DAQ, with commodity servers and switched network**
  - 4.6 TB/s aggregate throughput across whole system
- **FELIX (Front-End LInk eXchange)**
  - PCIe cards with single FPGA, custom optical links for front-end
  - Already running in Run 3 for new detectors and trigger components
- **Data Handler collects data fragments for detector-specific preparation**
- **FELIX prototypes upgrade for Run 4**
  - **2nd prototype (FLX-182)**: Xilinx VM1802, PCIe Gen4, 24 links up to 25 Gb/s + TTC interface - 50 boards production for integration tests
  - **3rd prototype (FLX-155)**: Xilinx VP1552, PCIe Gen5, up to 48 links. Prototype design approved for production
  - Firmware is mature and expandable
  - Final Design Review by the end of 2024

Dataflow and network

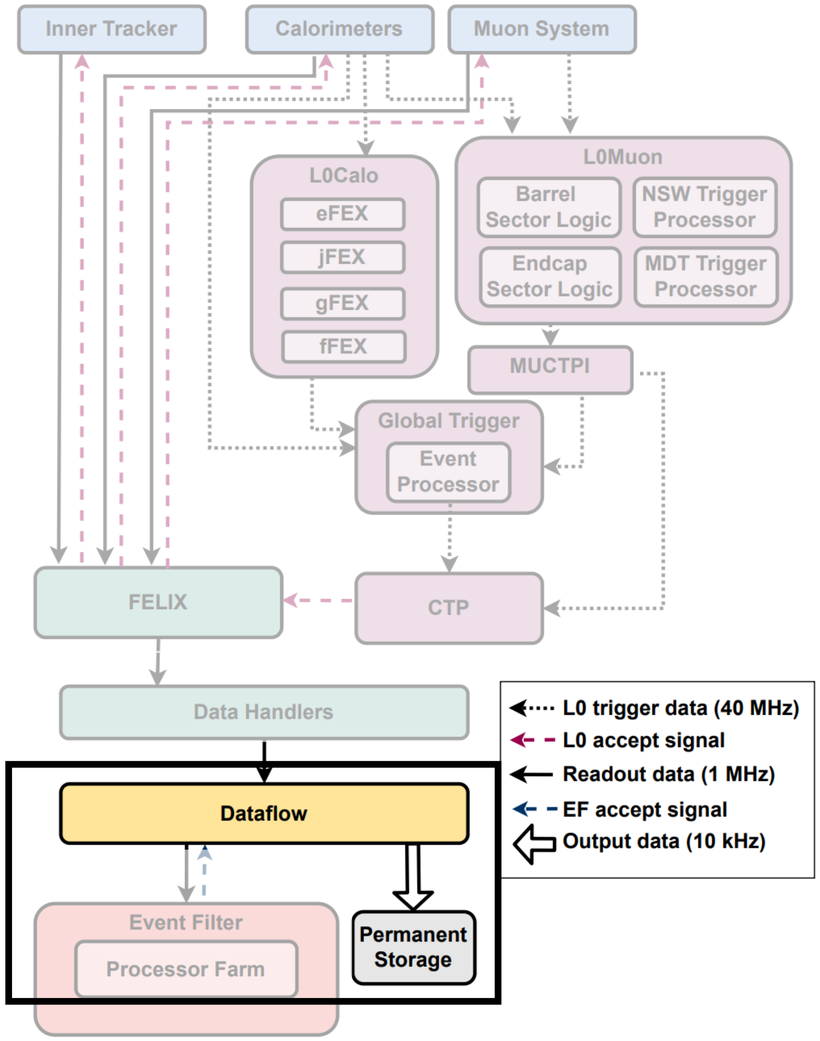

- **Dataflow aggregates data (full-event-building), buffered for Event Filter and sent to permanent storage**
  - Software prototype, based on Run 3, plus optimisations to support required 12 kHz event rate per Event Filter rack
  - Preliminary design in 2023, large scale tests in Run 3 system
- **Networking capabilities expanded, new simulation models [link]**
  - Controlled with traffic-shaping and sufficient buffers
- **Persistent-storage buffer prototype using DAOS (open-source high-performance distributed storage system)**
  - Alternative to Run 3 application-based design

Dataflow and network

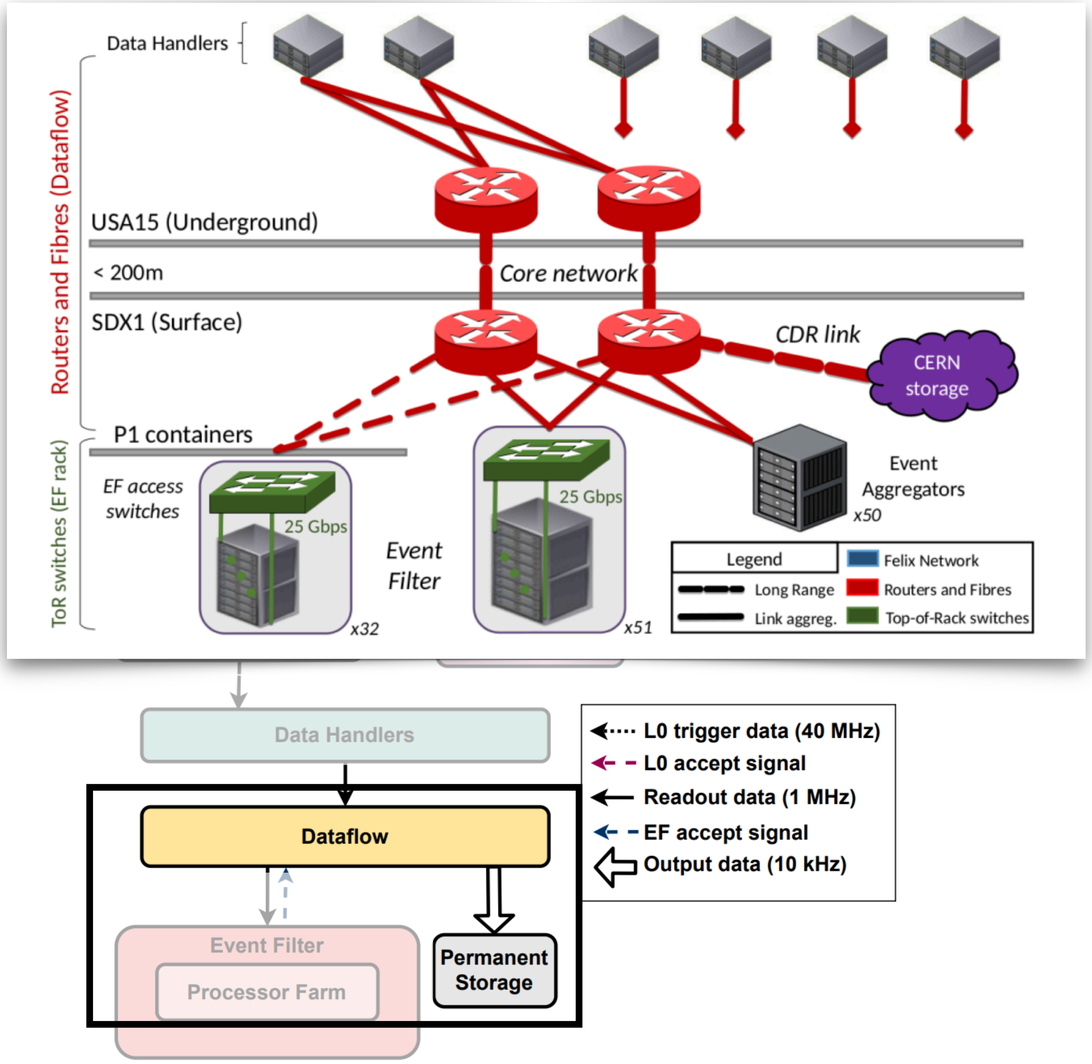

- **Dataflow aggregates data (full-event-building), buffered for Event Filter and sent to permanent storage**
  - Software prototype, based on Run 3, plus optimisations to support required 12 kHz event rate per Event Filter rack
  - Preliminary design in 2023, large scale tests in Run 3 system
- **Networking capabilities expanded, new simulation models [link]**
  - Controlled with traffic-shaping and sufficient buffers
- **Persistent-storage buffer prototype using DAOS (open-source high-performance distributed storage system)**
  - Alternative to Run 3 application-based design

Event Filter

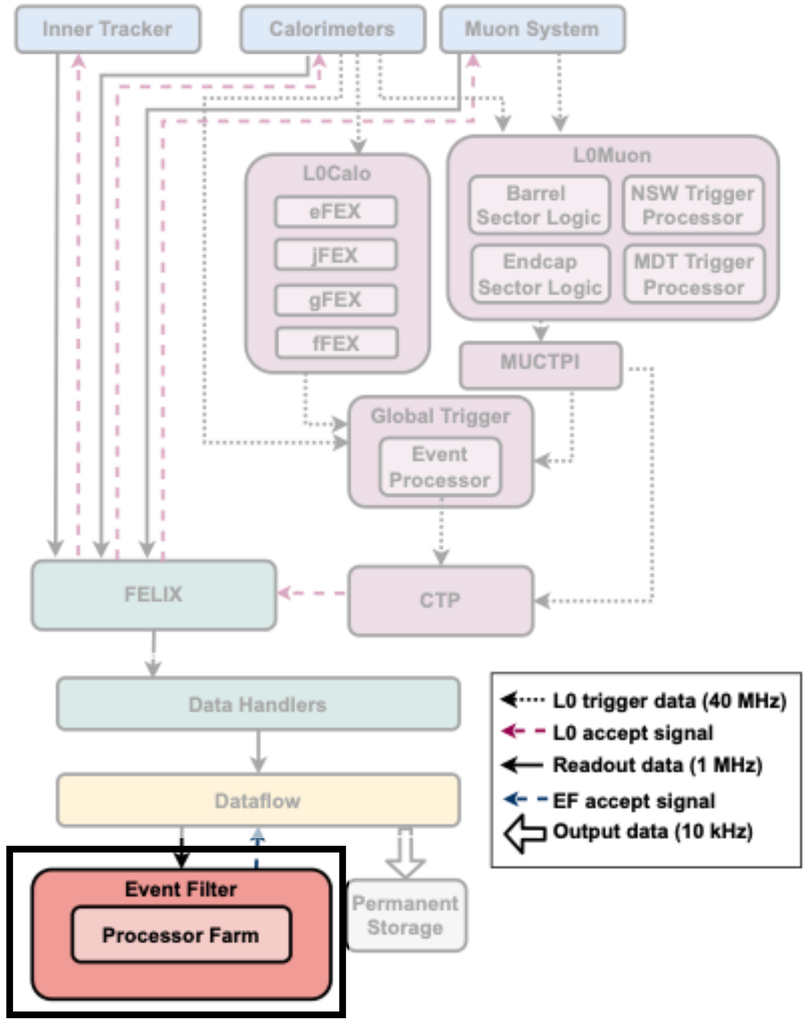

- **Software trigger with full event-building @1MHz**
- **Farm of heterogeneous commodity processors, w/offline-like algorithms, possibly with accelerators (CPU + GPUs and/or FPGAs)**
- **Feasibility demonstrated with algorithms on different platforms**
  - CPU: x8 speed-up with fast tracking
  - GPU: x12 speed-up on topological calorimeter cell clustering
  - Many ML approaches (GNN, CNN, RNN) look promising
    - Forum dedicated to ML on FPGA
- **Building demonstrators, to investigate use of accelerators, with final technology choice in 2025. Mainly driven by tracking needs**

ATLAS-TDR-029-ADD

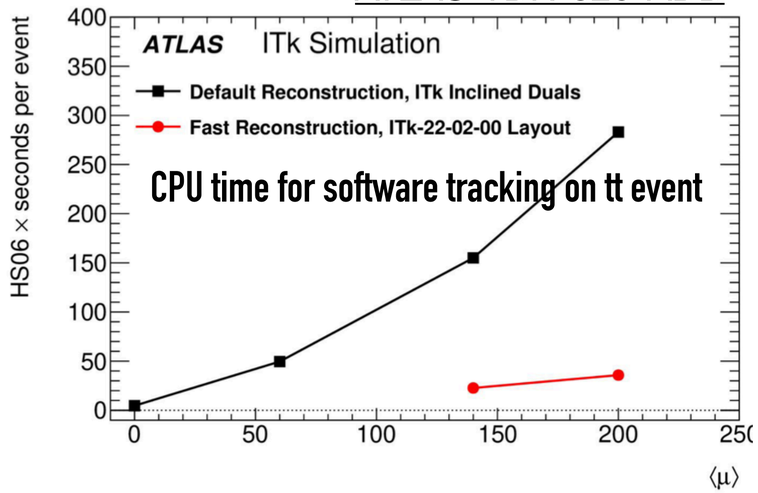

ATL-DAQ-PROC-2022-002

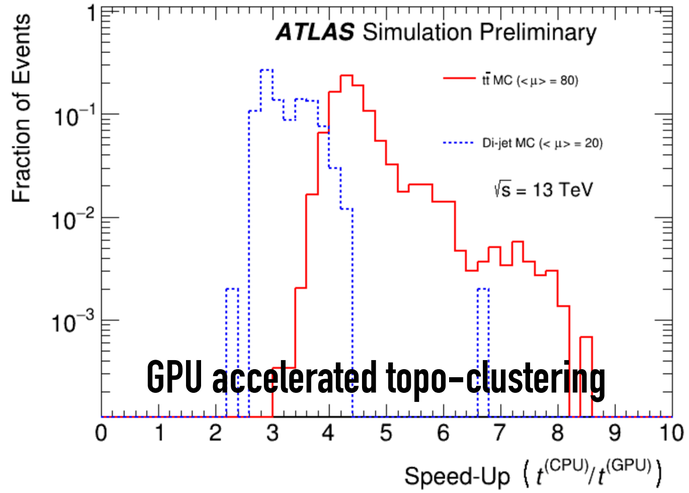

EF Tracking

- **Project dedicated to assess the best technology choice for Event Filter**
- **Ongoing R&D on accelerators, comparing CPU, GPU and FPGA**
  - Support for common GPU language/API to all ATLAS experiment
  - Settled FPGA family on AMD/Xilinx
- **Demonstrators for test slices of tracking algorithms on multiple technologies (CPU+GPU+FPGA) for commodity boards**
  - Split into track seeding, pattern recognition, track fitting, ambiguity removal
  - Including NN options (eg GNN)
  - Exploring use of High Level Synthesis (HLS)
- **Interfaces to ATLAS software via A Common Tracking Software (ACTS)**
  - Experiment independent toolkit for tracking

Accelerator based | CPU based

?

Online software

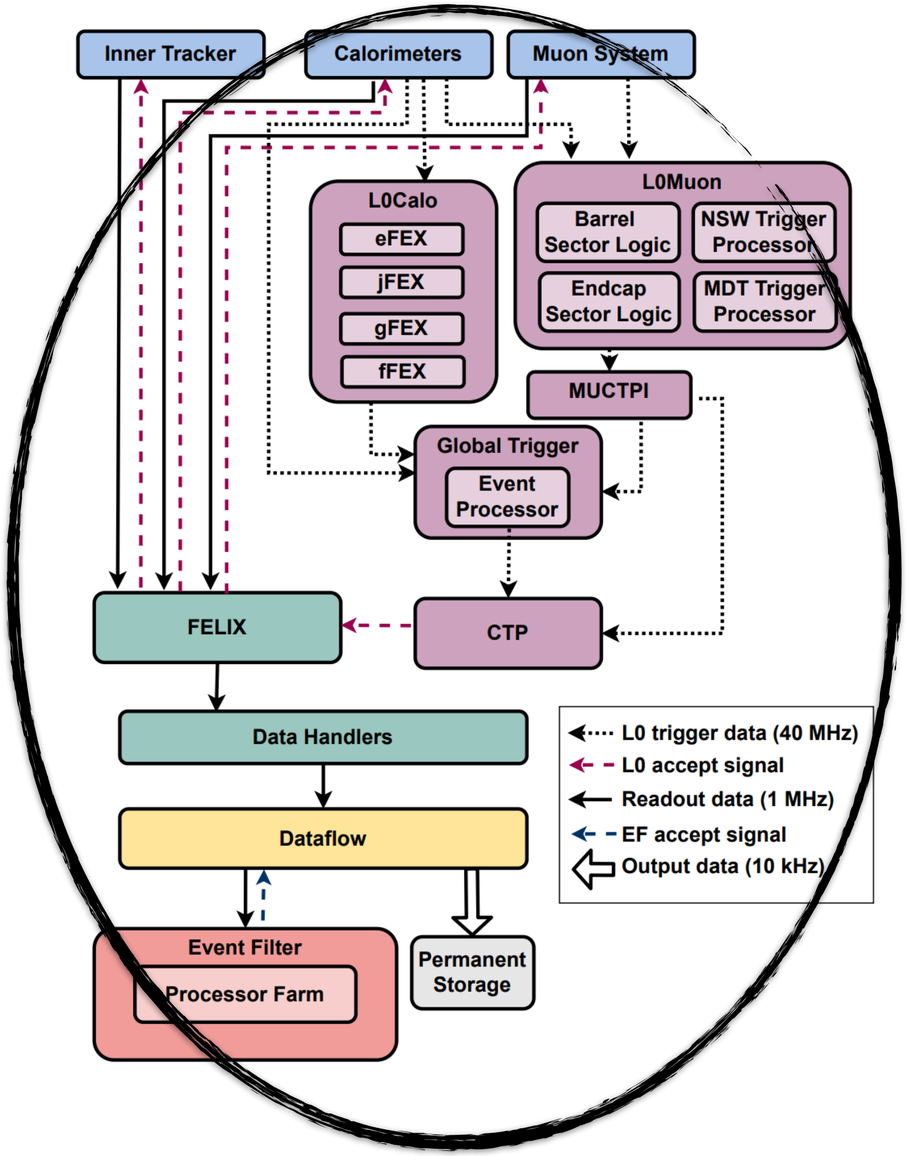

kubernetes

- **Gluing whole T/DAQ system together for common configuration, control and monitoring**
- **Prototype based on Kubernetes as farm orchestrator**
  - Open-source platform to automate deployment, management and scaling with containerised applications
- **Tested successfully on Run 3 farm, with 2600+ nodes**
  - Scaling the cluster size is still a research topic [OpenAI]
  - Very few in the world working with a cluster of this size
- **Being integrated within ATLAS Finite State Machine (Run Control)**

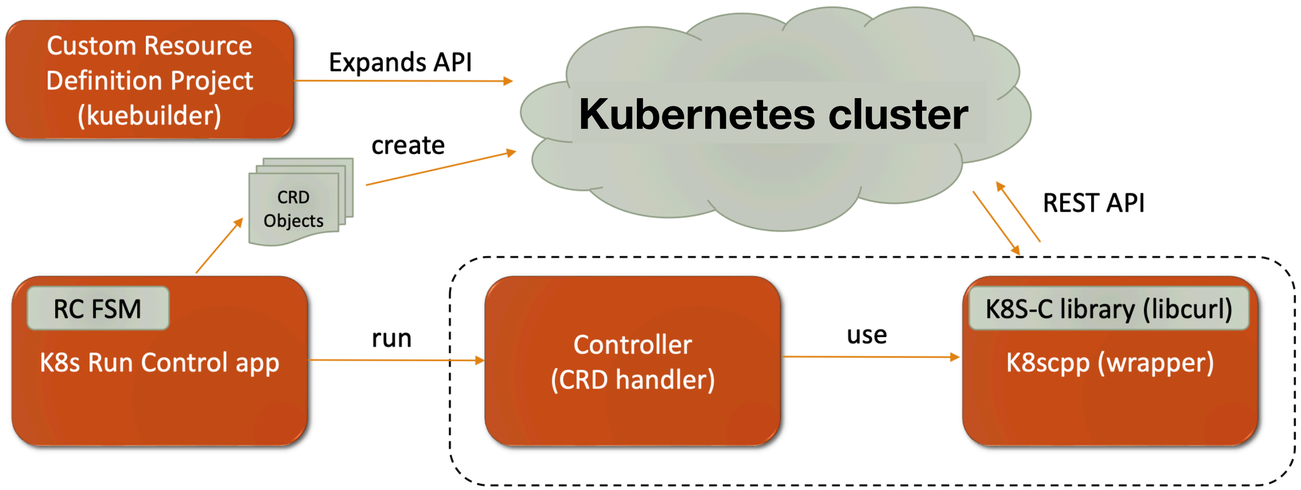

Red boxes -> In-house Software

Summary

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**
- **Level-0 systems focusing on finalising design and progressing to integration with prototypes that already exist**
  - With major effort in firmware development

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**
- **Level-0 systems focusing on finalising design and progressing to integration with prototypes that already exist**
  - With major effort in firmware development
- **Readout and DAQ have developed prototypes, now under testing**
  - Adopting latest products available on the market

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**
- **Level-0 systems focusing on finalising design and progressing to integration with prototypes that already exist**
  - With major effort in firmware development
- **Readout and DAQ have developed prototypes, now under testing**
  - Adopting latest products available on the market
- **Event Filter software trigger studying best options for use of accelerators**
  - Extensive R&D for both algorithms and hardware choices

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**
- **Level-0 systems focusing on finalising design and progressing to integration with prototypes that already exist**
  - With major effort in firmware development
- **Readout and DAQ have developed prototypes, now under testing**
  - Adopting latest products available on the market
- **Event Filter software trigger studying best options for use of accelerators**
  - Extensive R&D for both algorithms and hardware choices
- **Physics and performance studies ongoing to ensure performance goals met**
  - Enabling new trigger algorithms and developing updated Run 4 trigger menu

Summary

- **Very active area of Phase-2 upgrades in the ATLAS Trigger and DAQ systems**
- **Level-0 systems focusing on finalising design and progressing to integration with prototypes that already exist**
  - With major effort in firmware development
- **Readout and DAQ have developed prototypes, now under testing**
  - Adopting latest products available on the market
- **Event Filter software trigger studying best options for use of accelerators**
  - Extensive R&D for both algorithms and hardware choices
- **Physics and performance studies ongoing to ensure performance goals met**
  - Enabling new trigger algorithms and developing updated Run 4 trigger menu
- **Detailed integration plan, from now through to installation and commissioning**


ACAT 2024
22nd International Workshop on Advanced Computing Analysis Techniques in Physics Research
11 – 15 March 2024 | Stony Brook, New York
Charles B Wang Center – Stony Brook University
Abstract Submission by January 20th, 2024
Website: https://acat2024.sdcc.bnl.gov/
Email: acat-loc2024@cern.ch
International Advisory Committee: David Britton (Glasgow Univ, UK) Chair, Gang Chen (iHEP, China), Latifa Elouadrhiri (Jefferson Lab, USA), Ian Fisk (FII, USA), Maria Girone (CERN, Switzerland), Gudrun Heinrich (KIT, Germany), Gregor Kasieczka (UHH, Germany), Doris Kim (Soongsil Univ, Korea), Jerome Lauret (BNL, USA), Daniel Maitre (Durham, UK), Axel Naumann (CERN, Switzerland), Jennifer Ngadiuba (FNAL, USA), Fons Rademakers (CERN, Switzerland), Liz Sexton-Kennedy (Fermilab, USA), Lucia Silvestris (INFN Bari, Italy), Liliana Teodorescu, (Brunel Univ, UK), Claudio Torres (USM, Chile), Gordon T. Watts (Univ. Washington, USA), Monique Werlen (LAPTH, France)
Local Organizing Committee: Christopher Bee (SBU), Alan Calder (SBU) Co-chair, Ian Fisk (FII), Ann DeGregorio Gardner (SBU), Robert Harrison (SBU), John Hobbs (SBU), Jerome Lauret (BNL) Chair, Meifeng Lin (BNL), Liz Mogavero (BNL), Karin Muscarella (SBU), Michael Poat (BNL), Danielle Zimmerman (SBU), Michael Zingale (SBU)