CMS Trigger and DAQ upgrades for Run-3

D. Rabady (CERN), on behalf of the CMS collaboration

EP Detector Seminar

3 May 2024

CMS triggers in two stages

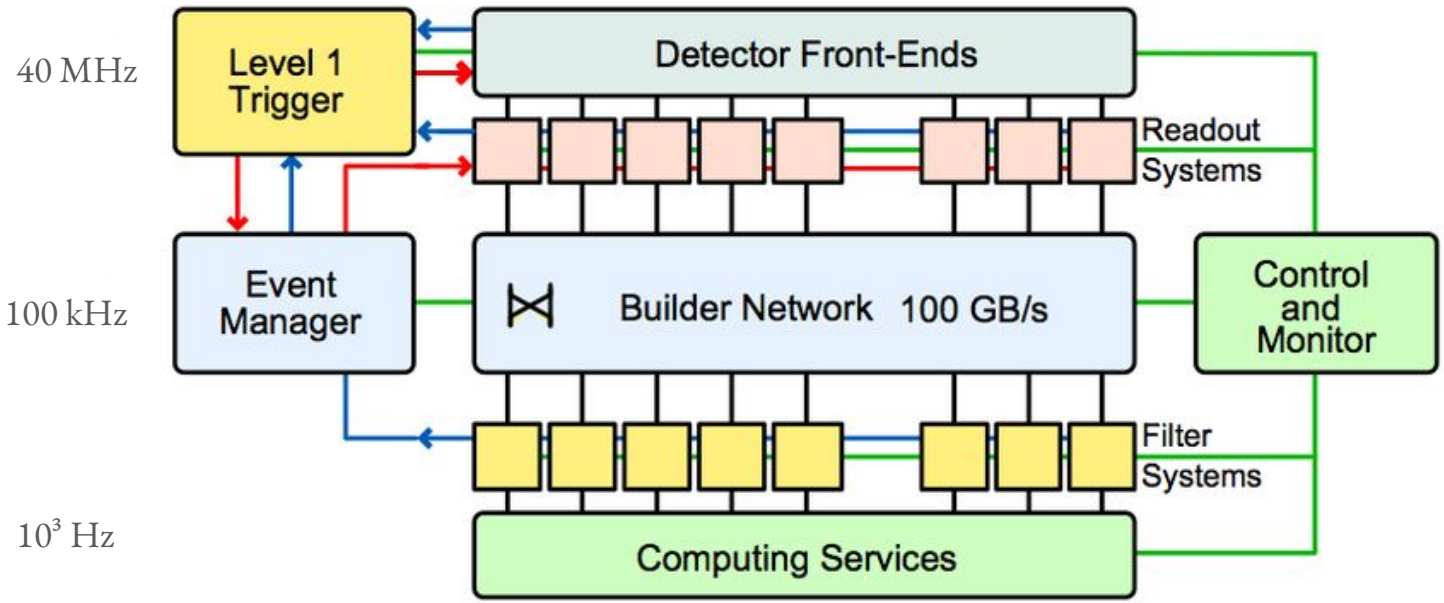

Level-1 trigger system design

- Muon trigger combining all three muon detectors
  - Regional track-finders
  - Global sorting and cancel-out layer also extrapolates coordinates back to interaction region
- Two-layer calorimeter trigger
  - Tower-level calibrations
  - Pileup subtraction
  - Independent calibrations for jets, taus, e/gamma
- Global Trigger with 512 possible algorithms
  - Single, multi- and cross-object
  - Topological information

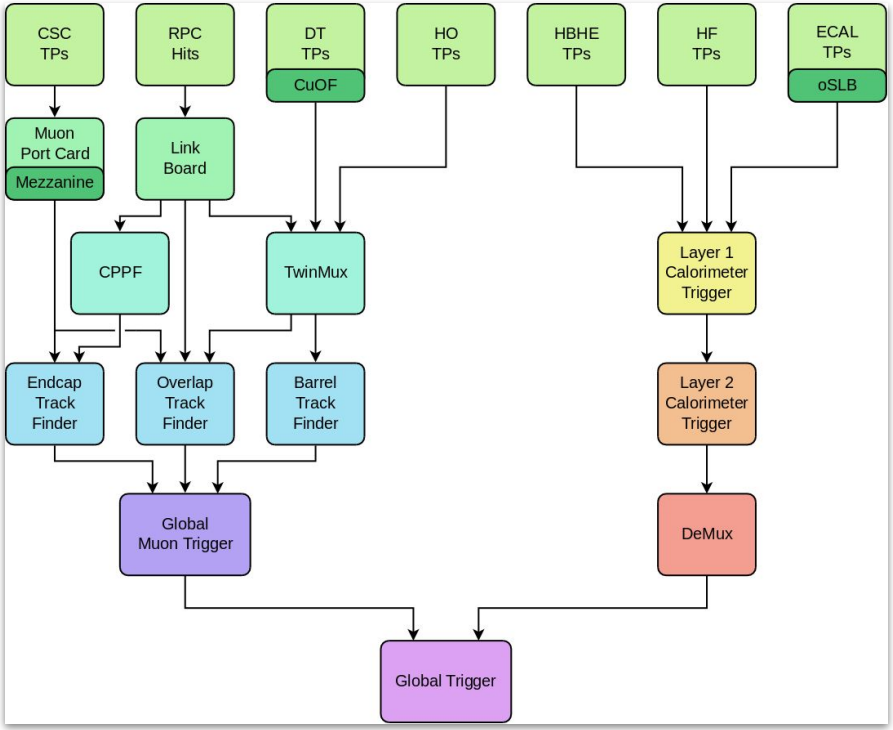

Level-1 trigger system improvements

- Dedicated algorithms for long-lived particle searches
- Improved pile-up mitigation
- Integration of GEM and ZDC detectors
- Demonstrating data scouting at the bunch crossing frequency

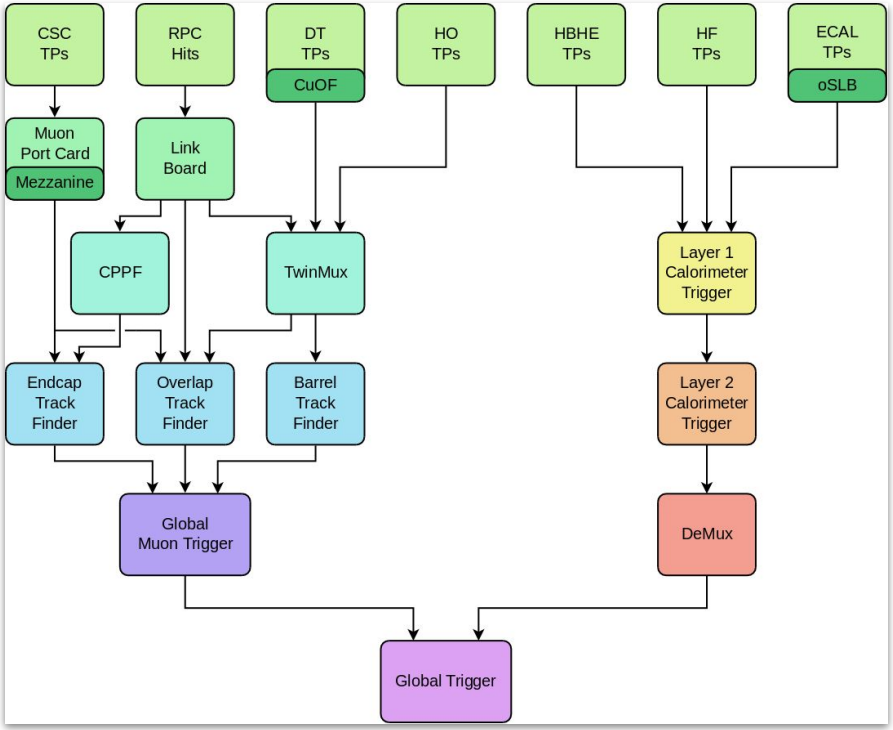

Displaced muons for Run-3

- **Barrel:**
  - Profit from Phase-2 developed Kalman filter-based algorithm implemented using high-level synthesis (HLS) in firmware
  - Applying KF without constraint to beamspot allows to compute impact parameter and p_T of displaced track ("unconstrained p_T")

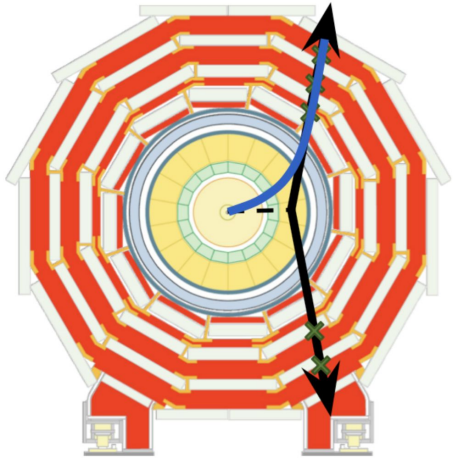

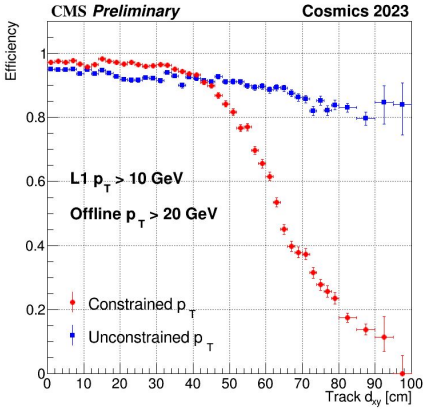

- **Endcap:**
  - Simplified Phase-2 NN used to assign impact parameter and p_T without beamspot constraint
- **Overlap:**
  - Naive Bayes classifier estimates the p_T of prompt muons based on a set of precomputed patterns
  - For displaced tracks, the same approach is used, but track extrapolation using bending angle from DT reference chambers is compared to measured location in order to obtain p_T without beamspot constraint

Showers in the muon detectors

- Muon detector showers are a powerful tool for LLP searches
  - Probing new LLP parameter space: long lifetimes and low masses
    - Decay after calorimeters causes showers in muon detectors
  - Traditionally triggering for these signatures using missing energy, but low efficiency
- Implementation:
  - Use spare bandwidth from cathode strip chambers to endcap muon track finders (EMTF) to transmit few bits to indicate when chambers detect more than a certain number of local charged tracks (LCTs)
  - EMTF and Global Muon Trigger apply logic to create
    - Single shower trigger with high threshold
      - Triggers on any LLP energetic enough to create shower in CSCs
    - Two showers in different sectors trigger with lower thresholds
      - Use to trigger on pair produced LLPs

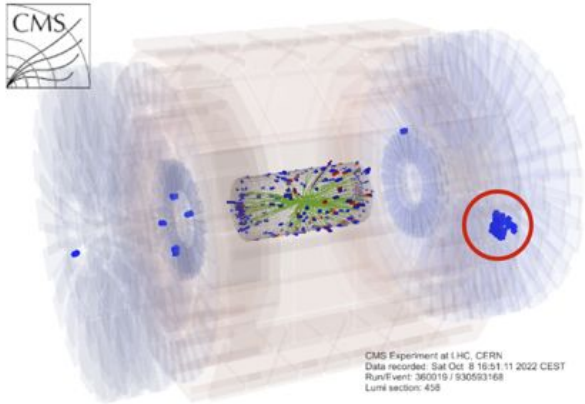

Long-lived particles in the calorimeters

- HCAL transmits depth and timing information to trigger
  - Six bits available in HCAL → Calorimeter Layer-1 trigger link
  - Used to flag signals characteristic of exotic long-lived particle decays
    - TDC timing used to mark hits with late arrival times
    - Shower profile used to mark decays within HCAL volume
  - Only one bit available on Calorimeter Layer-1 → Layer-2 link
    - Can use timing & shower profile to create single "LLP flag"
- Can select hadronic signatures from LLPs with decay lengths of 1–2 m which decay prior to or within HCAL
  - Require two trigger towers with LLP flag set in jet algorithm

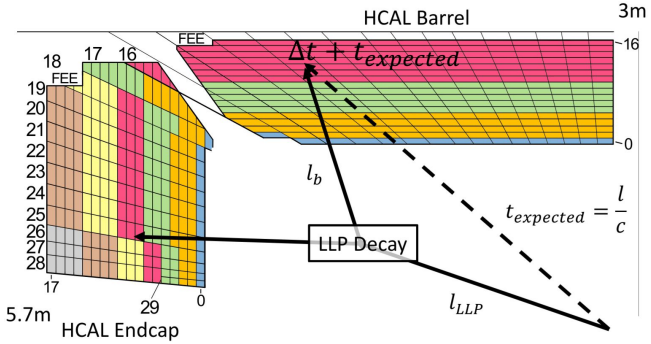

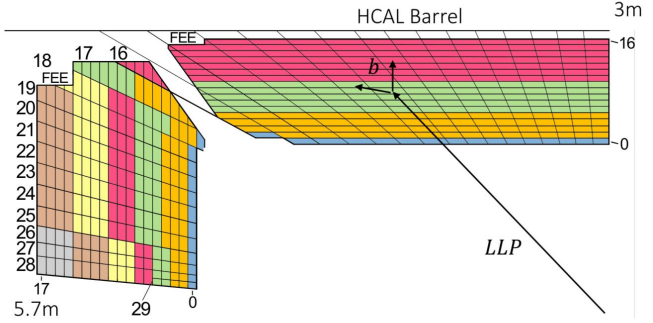

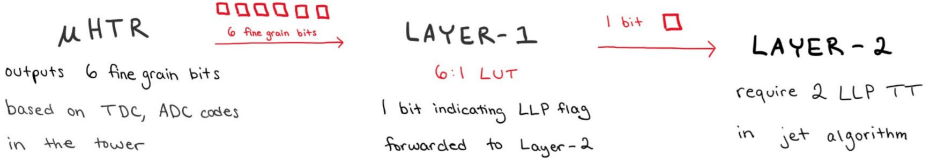

40 MHz data scouting system

- Use objects reconstructed by L1 trigger to do physics
  - Semi real-time analysis and/or storing of tiny event record
  - Enables study of exotic signatures that cannot be fit into the trigger budget
    - Targeted at Phase-2
    - Demonstrator taking data now
- L1 scouting system included in CMS runs regularly
  - Data taken so far primarily used for data quality monitoring
    - Pre-firing, bunch-to-bunch correlations, luminosity monitoring etc.
  - First online analyses currently under commissioning

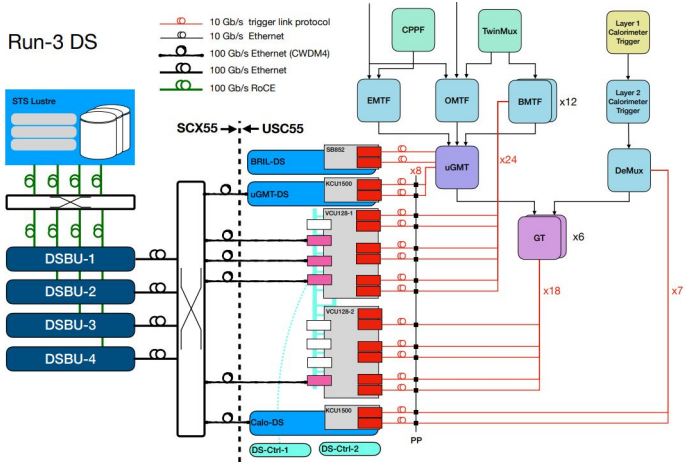

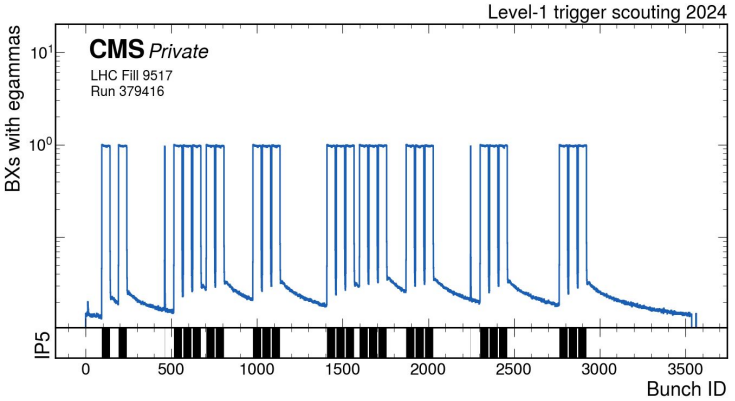

The data acquisition system

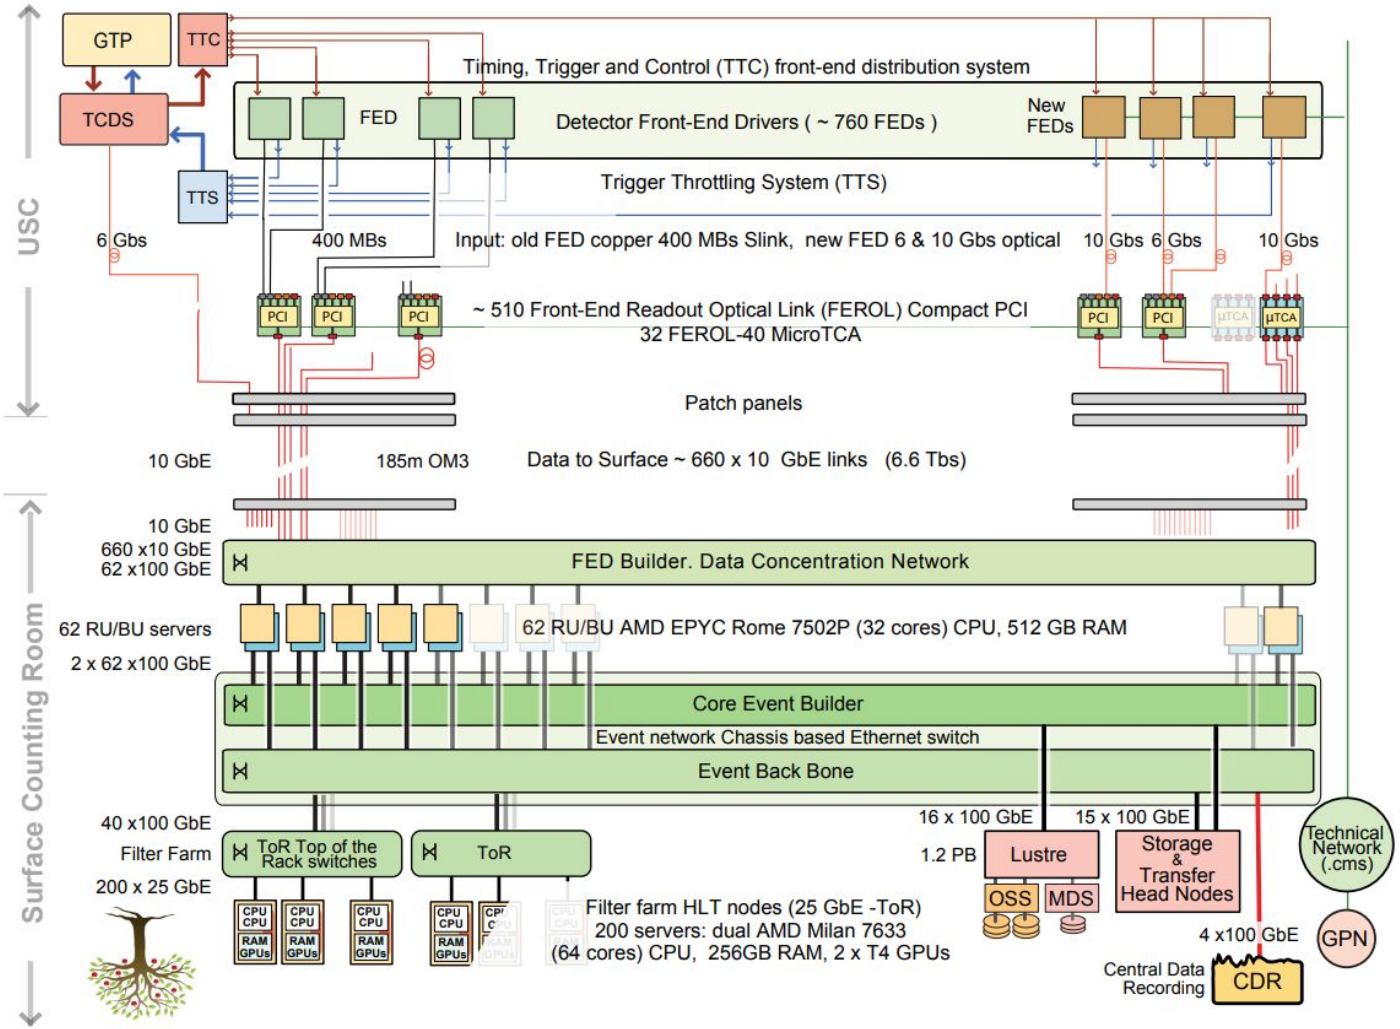

The data acquisition system

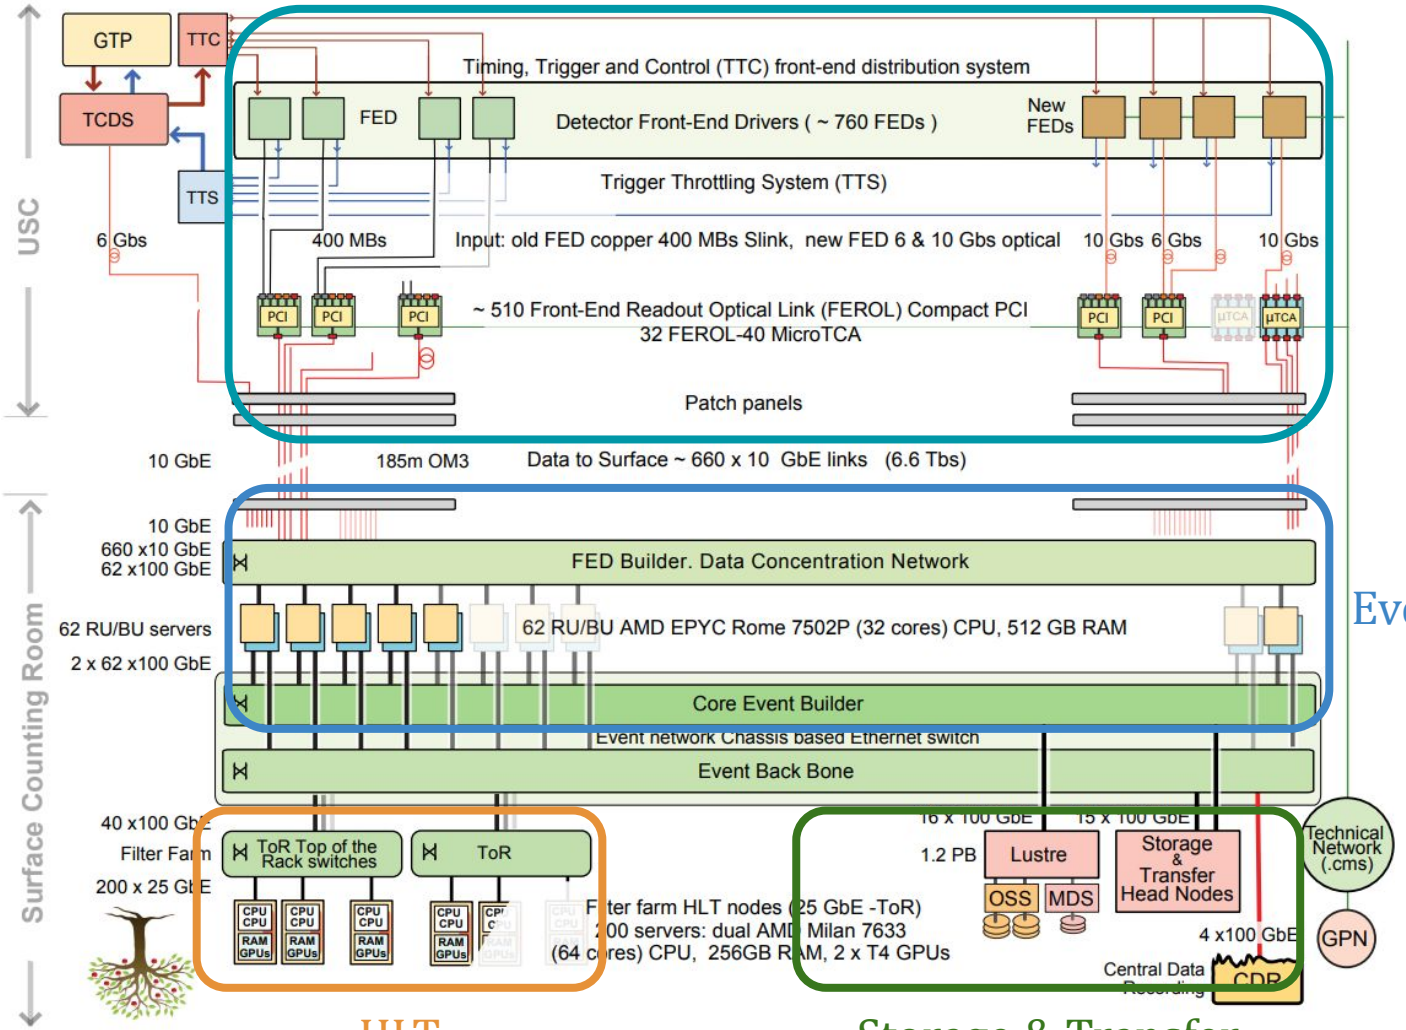

Event building in Run-2

- Event fragments → Readout units
  - Receiver units build super fragments
- All super fragments for a given event → Builder unit
- Drawback: unidirectional use of network links

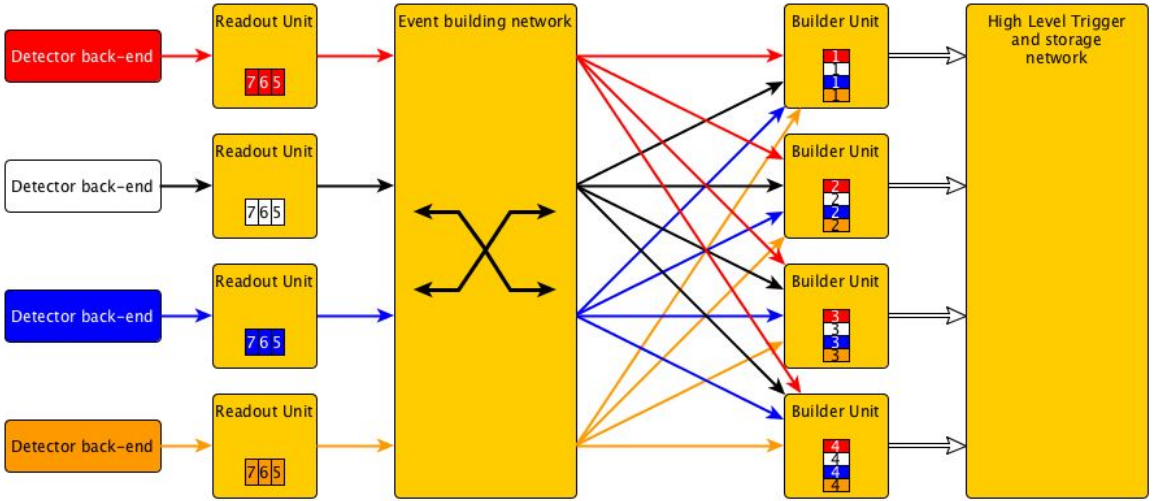

Folded event building in Run-3

- Readout and builder unit appliances housed in the same physical machine
- Allows bi-directional use of network links
  - Reduction of almost 50% of RU/BU nodes and network bandwidth
- Drawback: significant demand on I/O and memory performance of individual nodes
  - Required to receive and merge event fragments, exchange super fragments, build and serve events, pass HLT output to storage and transfer service

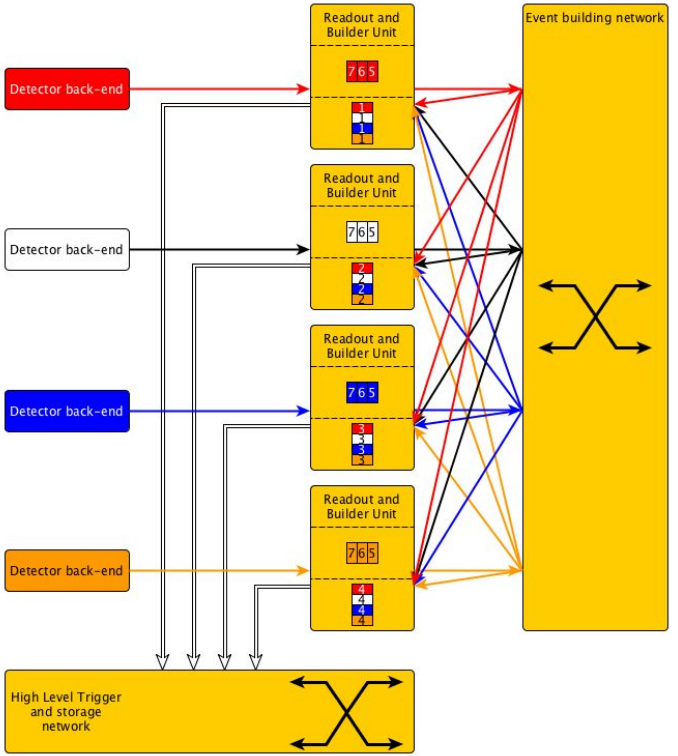

Event building tuning and performance

- Selected AMD EPYC Rome single socket CPUs to host RU/BU appliances
  - Only mild tuning needed when run in single NUMA domain mode
    - Groups of threads running similar tasks pinned to same or adjacent cores and closer to corresponding network interface PCIe lanes

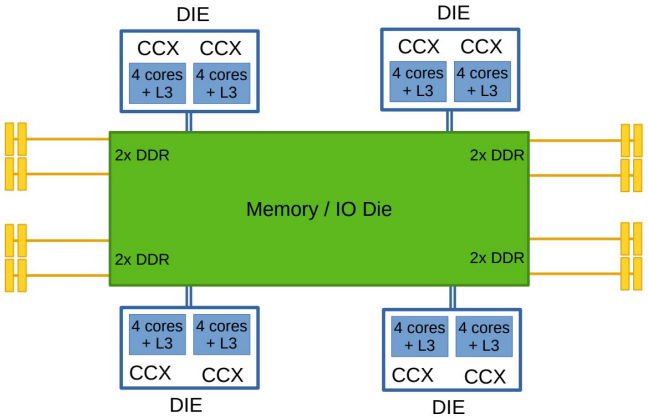

- Performance evaluated and deemed to be capable of fulfilling the Run-3 requirements
  - Measurements taken with 50 RU/BU nodes and shown to allow event building at 100 kHz up to event sizes of 2.5 MB

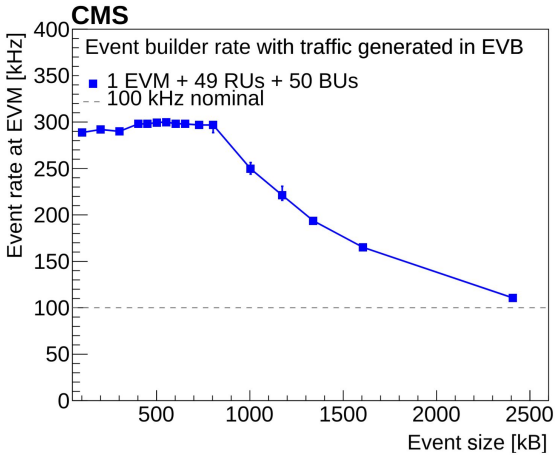

The CMS file-based filter farm

- HLT decoupled from central DAQ via file system
  - Builder unit application writes events in custom binary format to ramdisk that is exported via NFS
  - Files are distributed to HLT processes running on filter units (FUs) assigned to a particular builder unit
- HLT is implemented in offline analysis framework (CMSSW)
  - Allows to reuse data structures, unified use of conditions, and faster deployment of algorithms
- Since beginning of Run-2 CMSSW uses task-based multi-threaded event processing
  - Reduced memory footprint due to sharing of detector conditions and calibrations within a process
  - Allowed to exploit all logical CPU cores

Filter units with accelerator support

- For Run-3 support for offloading to accelerators was added
  - GPUs in HLT reduce average event processing time by 40%
    - Cost of the filter farm could be reduced by 15%, power consumption by 30% vs. farm without GPUs
  - Using Alpaka abstraction library to allow vendor-agnostic development of algorithms
- Run-3 filter farm implemented with AMD EPYC Milan dual socket machines with two Nvidia T4 GPUs
  - CMSSW processes split into two groups with each group pinned to one single-socket NUMA domain as well as to GPU that is attached to PCIe lanes of associated socket

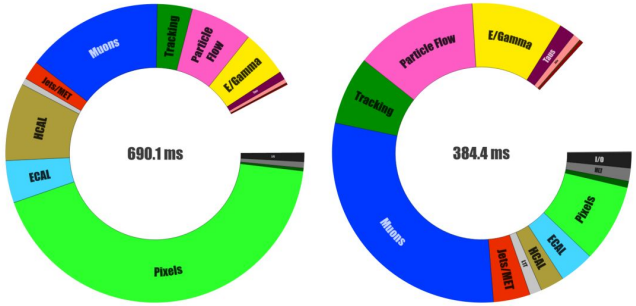

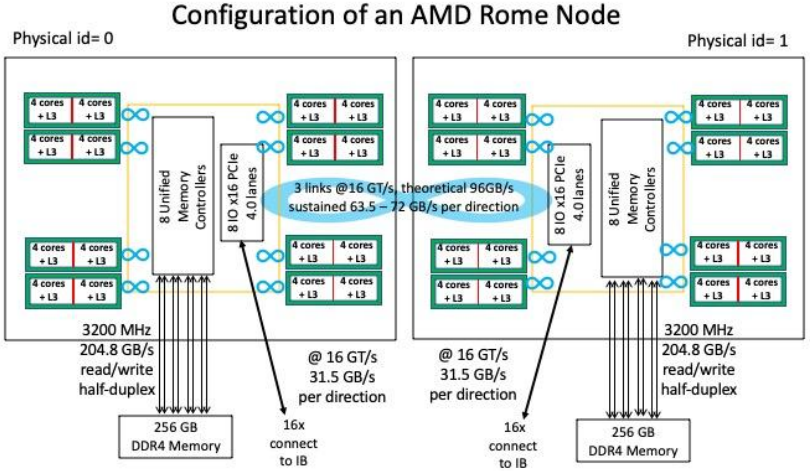

HLT Run-3 algorithm improvements

- Tracking
  - For **Run-2:** three/four iterations of a combinatorial Kalman filter
  - **Run-3:** Single global iteration, seeded by loose selection of pixel tracks reconstructed by parallelised algorithm developed for use in GPUs
    - Fewer iterations and less CPU time than Run-2 algorithm, but better performance
- Tau lepton reconstruction in two steps
  - First pass reconstruction and isolation criterion computed using a CNN with pixel tracks and calorimeter candidates as input
  - Full reconstruction improved vs. Run 2 using DeepTau NN
    - For 2024 will be moving to ParticleNet
- B jet tagging
  - Improved light flavour jet misidentification rates vs. Run 2 (DeepCSV) by using two new NN-based jet taggers, DeepJet and ParticleNet

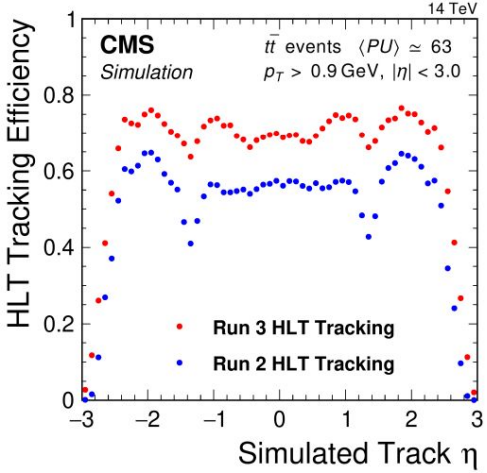

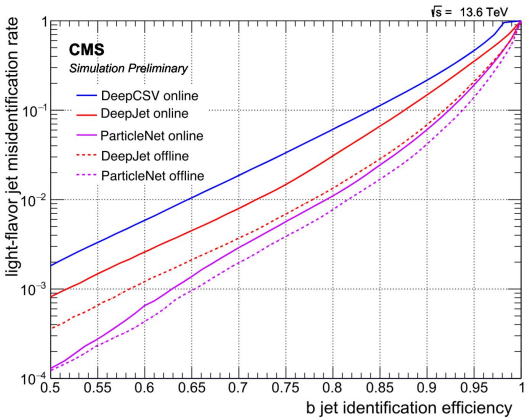

HLT Run-3 triggers

- Special paths for long-lived particles
  - Displaced jets, displaced leptons and photons, and delayed jets
    - Using timing information from ECAL and HCAL, as well as HCAL depth information
    - Partly seeded by newly developed L1 seeds
  - Three new high-multiplicity trigger (HMT) paths that target hadronic showers in the muon system
    - 2x seeded by L1 HMT paths, reconstructing clusters in the CSC or in the CSC and DT systems
    - 1x seeded by L1 MET, reconstructing cluster in the DT system
- Data scouting
  - Stores only the most relevant physics information, as reconstructed by the HLT
    - Significantly smaller event sizes: 11 kB vs. 1 MB for standard events
  - For Run 3 pixel tracks from parallelised algorithm used to directly seed particle flow reconstruction
    - Since 2024 running full HLT reconstruction for scouting stream
  - In addition to muons, jets, particle flow candidates stored during Run 2, also electrons, photons, and tracks are now stored for Run 3

Summary

- Level-1 trigger introduced significant improvements for LLP searches
  - Displaced muons, displaced/delayed jets, muon detector showers
- Demonstrator for scouting at the bunch crossing frequency has begun data taking
  - Currently commissioning the online analysis feature
- Using more efficient event building architecture
  - Bi-directional use of links reduced event builder size by almost 50%
- GPU-accelerated filter farm reduces average HLT event processing time by 40%
  - Efforts made to avoid vendor lock-in