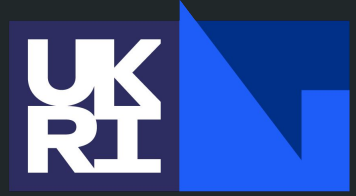

Science and Technology Facilities Council

Introduction to Trigger Systems 16th June 2026

Tim Martin (RAL Particle Physics Department)

Introduction

Acknowledgements and Disclaimer

- Lecture contains inputs from prior presenters, **Julie Kirk** (RAL PPD) & **Will Panduro Vazquez** (RAL PPD).
- ***I work on ATLAS – much of this talk will be collider-based and ATLAS-centric, but the concepts are general.***

What is a trigger system, and why do we need it?

- The problem:
  - Particle physicists are typically searching for **rare processes**.
  - Particle physics experiments, in particular at colliders, generate **enormous amounts of data**.

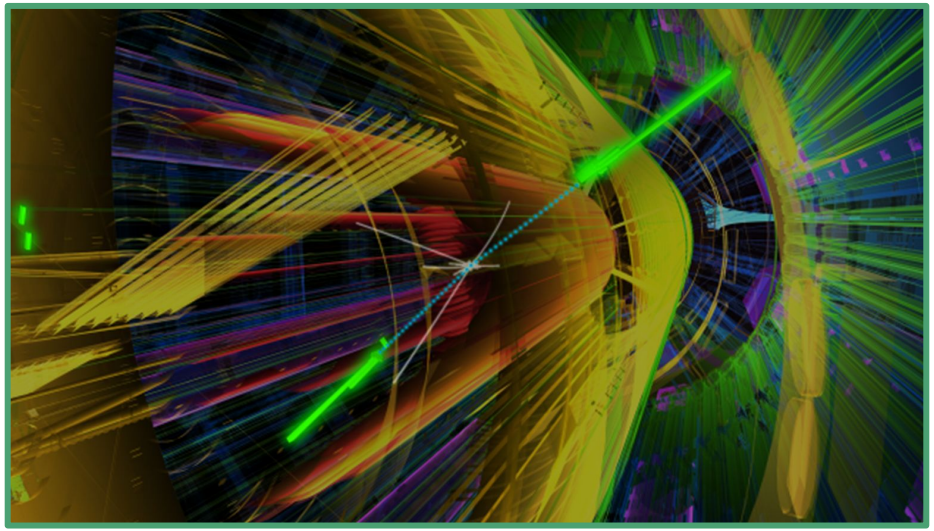

One barn, **10^{-24} cm^2**, is approximately the cross sectional area of a Uranium nucleus.

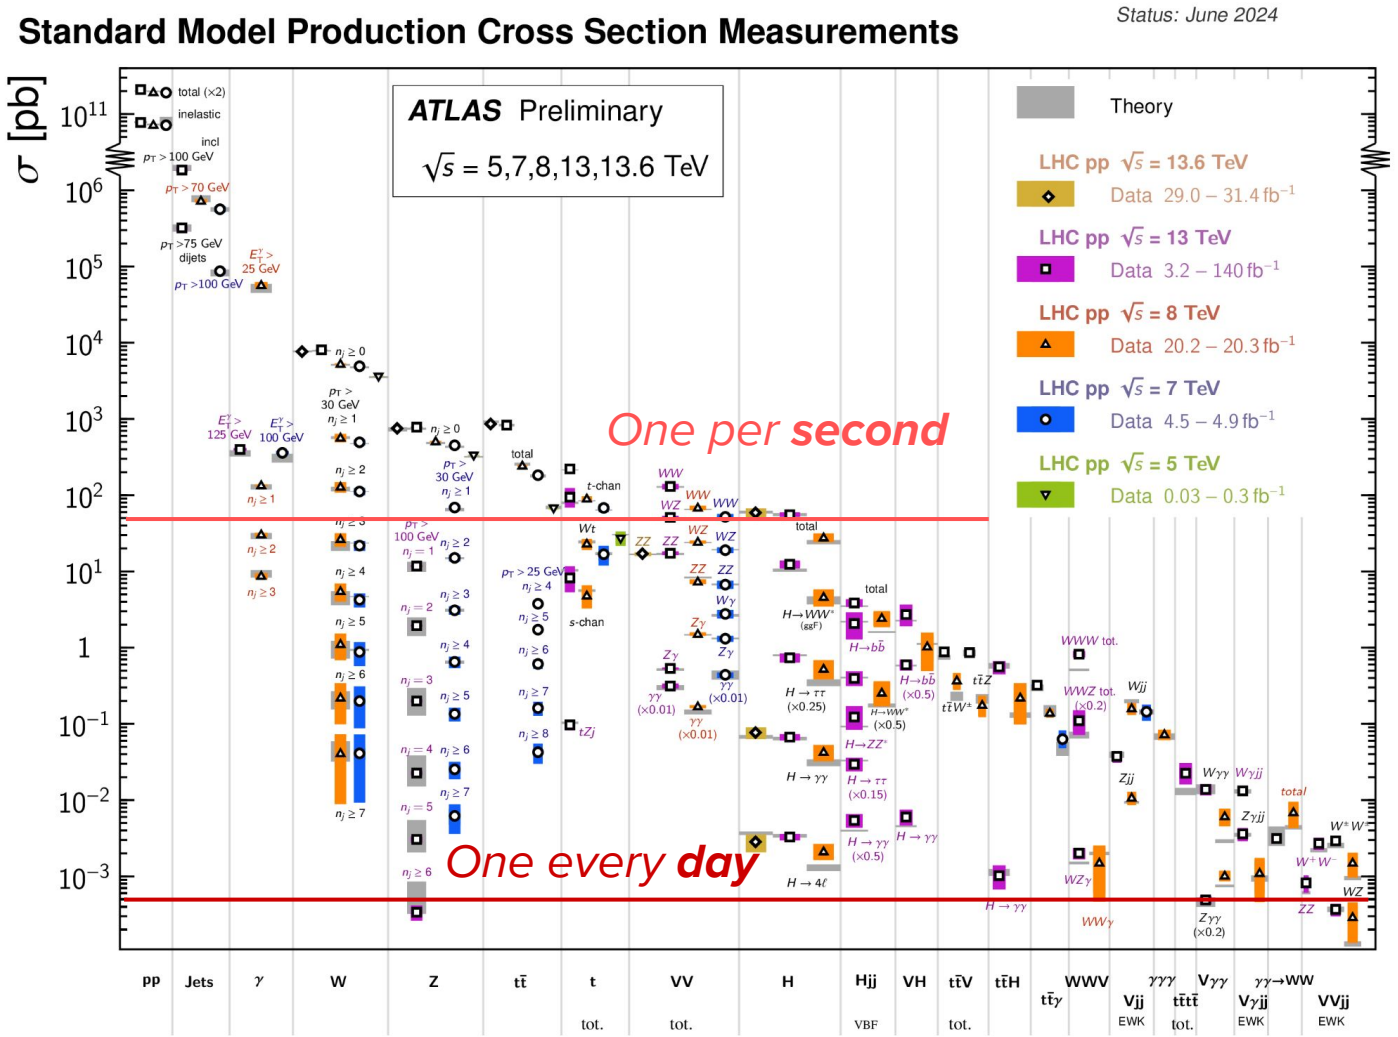

The LHC currently delivers an interaction flux (*luminosity*) of **2×10^{34} cm^{-2} s^{-1}**

One barn, **10^{-24} cm^2**, is approximately the cross sectional area of a Uranium nucleus.

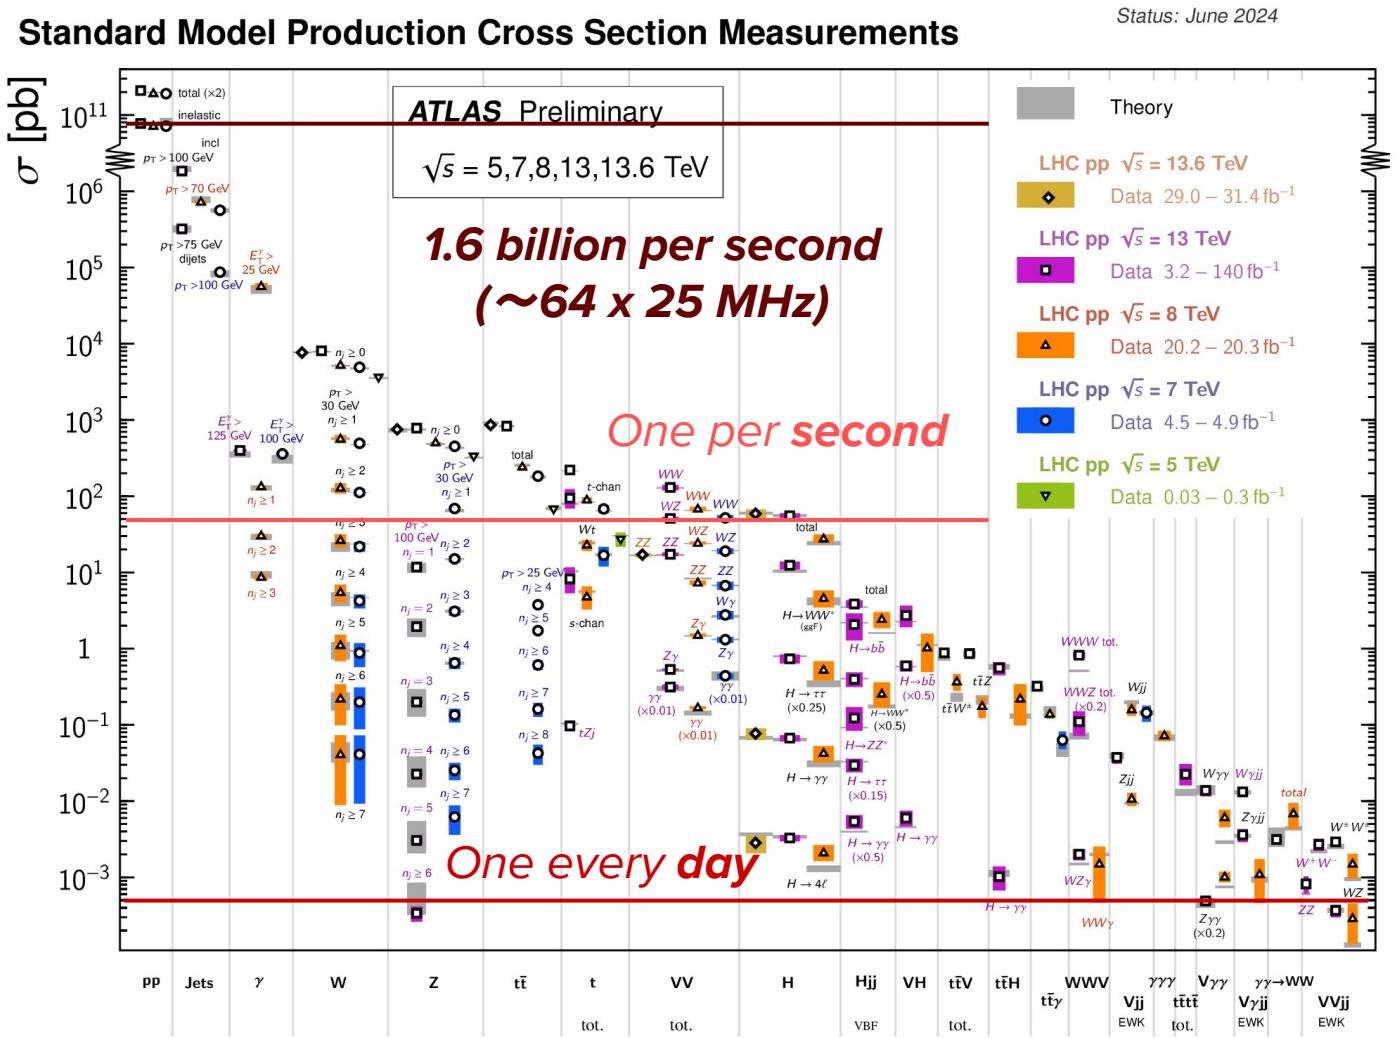

The LHC currently delivers an interaction flux (*luminosity*) of **2×10^{34} cm^{-2} s^{-1}**

What is a trigger system, and why do we need it?

- The problem:
  - Particle physicists are typically searching for **rare processes**.
  - Particle physics experiments, in particular at colliders, generate **enormous amounts of data**.
- For example:
  - **CMS** and **ATLAS** unfiltered off-detector data rate of order **60 TB/s**
  - Instantaneous global internet traffic in 2025: **395 TB/s** [AppLogic/Sandvine GIPR 2025]
  - ***No storage system we could possibly afford could cope with this volume of data!***

What is a trigger system, and why do we need it?

- The solution:
  - **Trigger systems!** Broadly -
    - *A system that uses simple criteria to rapidly decide which events in a particle detector to retain, when only a small fraction of the total can be recorded.*
- Example from Colliders:
  - ***Interesting events usually have high-p_T particles.***

$pp \rightarrow W(\rightarrow e\nu)\gamma\gamma$

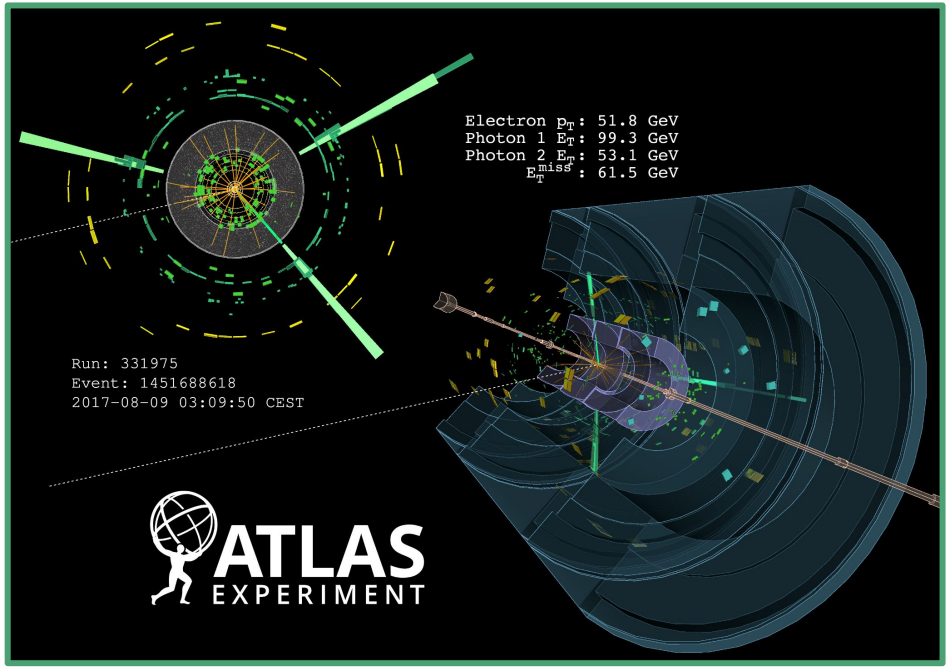

Phys. Lett. B 848 (2024) 138400

What is a trigger system, and why do we need it?

- The solution:
  - **Trigger systems!** Broadly -
    - *A system that uses simple criteria to rapidly decide which events in a particle detector to retain, when only a small fraction of the total can be recorded.*
- Example from Colliders:
  - ***Interesting events usually have high-p_T particles.***
  - ***... hidden in a mass of low-p_T background.***

2025 high-intensity test
⟨μ⟩= 150, 92 PVs

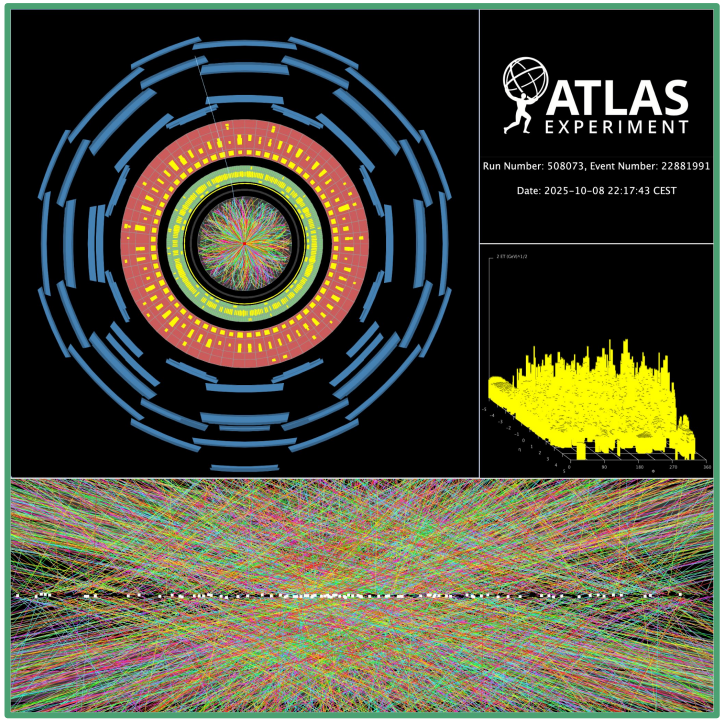

What is a trigger system, and why do we need it?

- The solution:
  - **Trigger syste** …adly -
    - *A system* …
      criteria …
      which …
      detect …
      only a small …
      total can be recorded.
- Example from Colliders:
  - ***Interesting events usually have high-p_T particles.***
  - ***… hidden in a mass of low-p_T background.***

No second chances - one the event is discarded, it's gone forever!

2025 high-intensity test
⟨μ⟩= 150, 92 PVs

From Gas to Higgs Bosons

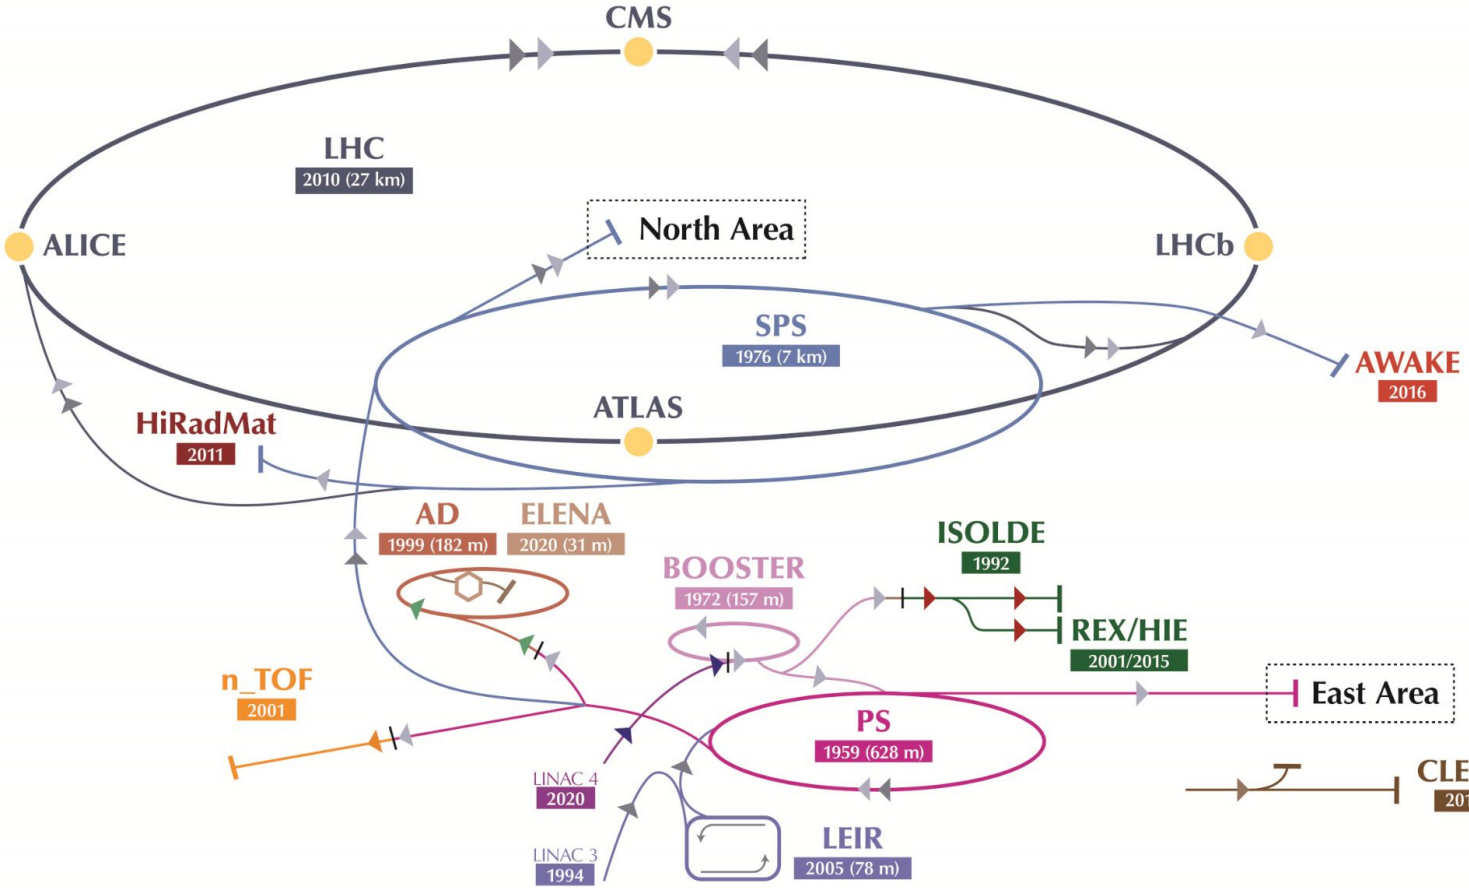
CMS
LHC
2010 (27 km)
ALICE
North Area
LHCb
SPS
1976 (7 km)
AWAKE
2016
ATLAS
HiRadMat
2011
AD
1999 (182 m)
ELENA
2020 (31 m)
ISOLDE
1992
BOOSTER
1972 (157 m)
REX/HIE
2001/2015
n_TOF
2001
East Area
PS
1959 (628 m)
LINAC 4
2020
CLEAR
2017
LINAC 3
1994
LEIR
2005 (78 m)


PREINJECTEUR
LINAC 2

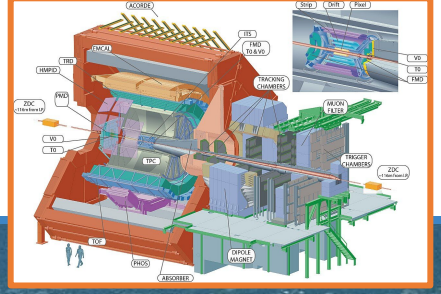
ACORDE
TRD
HMPID
PMD
ZDC
PHOS
TOF
TPC
TRACKING CHAMBERS
MUON FILTER
TRIGGER CHAMBERS
DIPOLE MAGNET
ABSORBER
Strip
Drift
Pixel

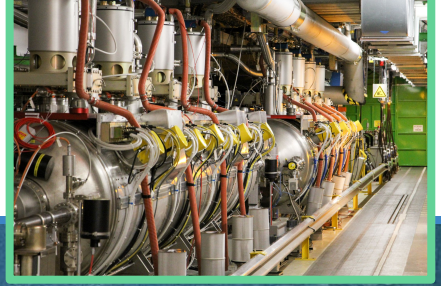

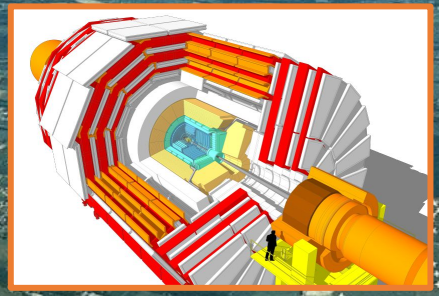

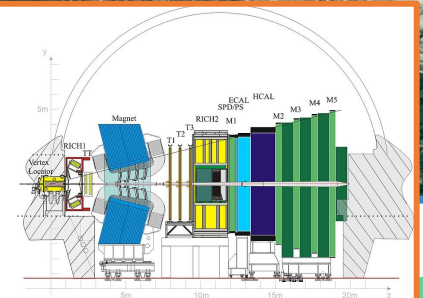
Magnet
RICH1
RICH2
ECAL
HCAL
SPD/PS
M1
M2
M3
M4
M5
T1
T2
T3
Vertex Locator

- LHC Radio Frequency is **400.8 MHz**
- LHC Circumference is **26,659 m**
- Ring divides into **35,640** buckets
  - Buckets separated by around **0.75 m**
- Proton **bunches** from the PS separated by **25 ns**
  - Buckets which may have protons are separated by around **7.5 m**
- The ring is never *entirely filled*.
- With 2,340 filled ***bunches***, bunch-crossings occur at **26,314,470 Hz**

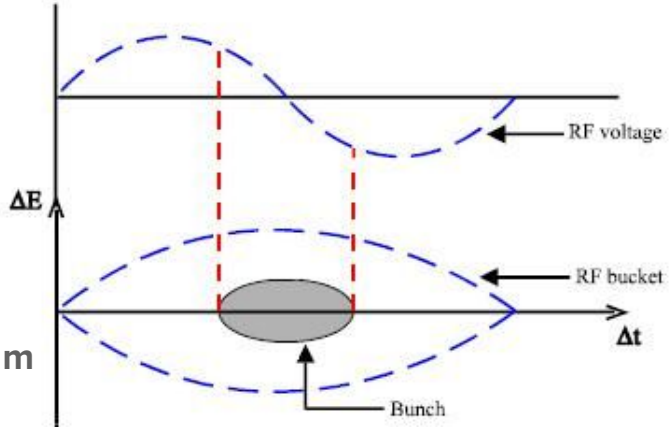

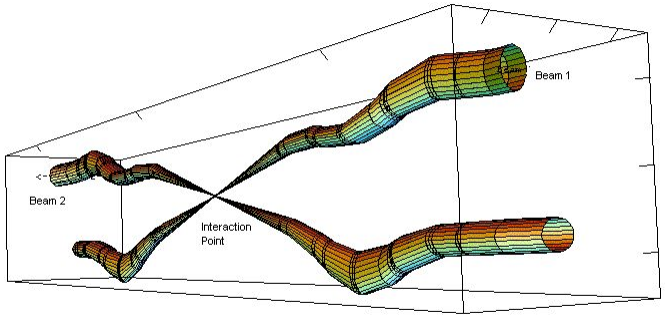

Relative beam sizes around IP1 (Atlas) in collision

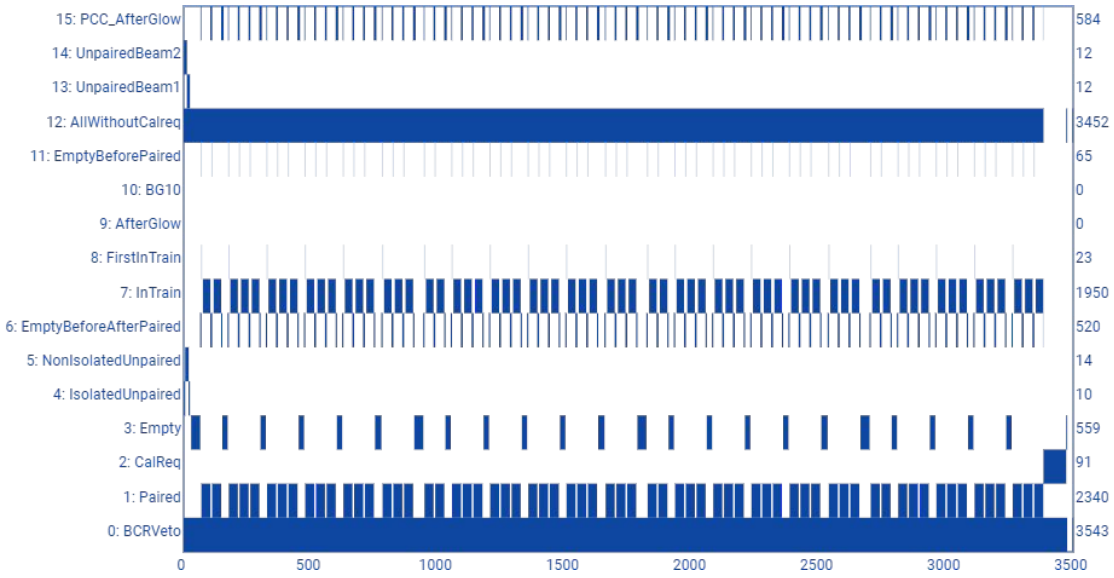

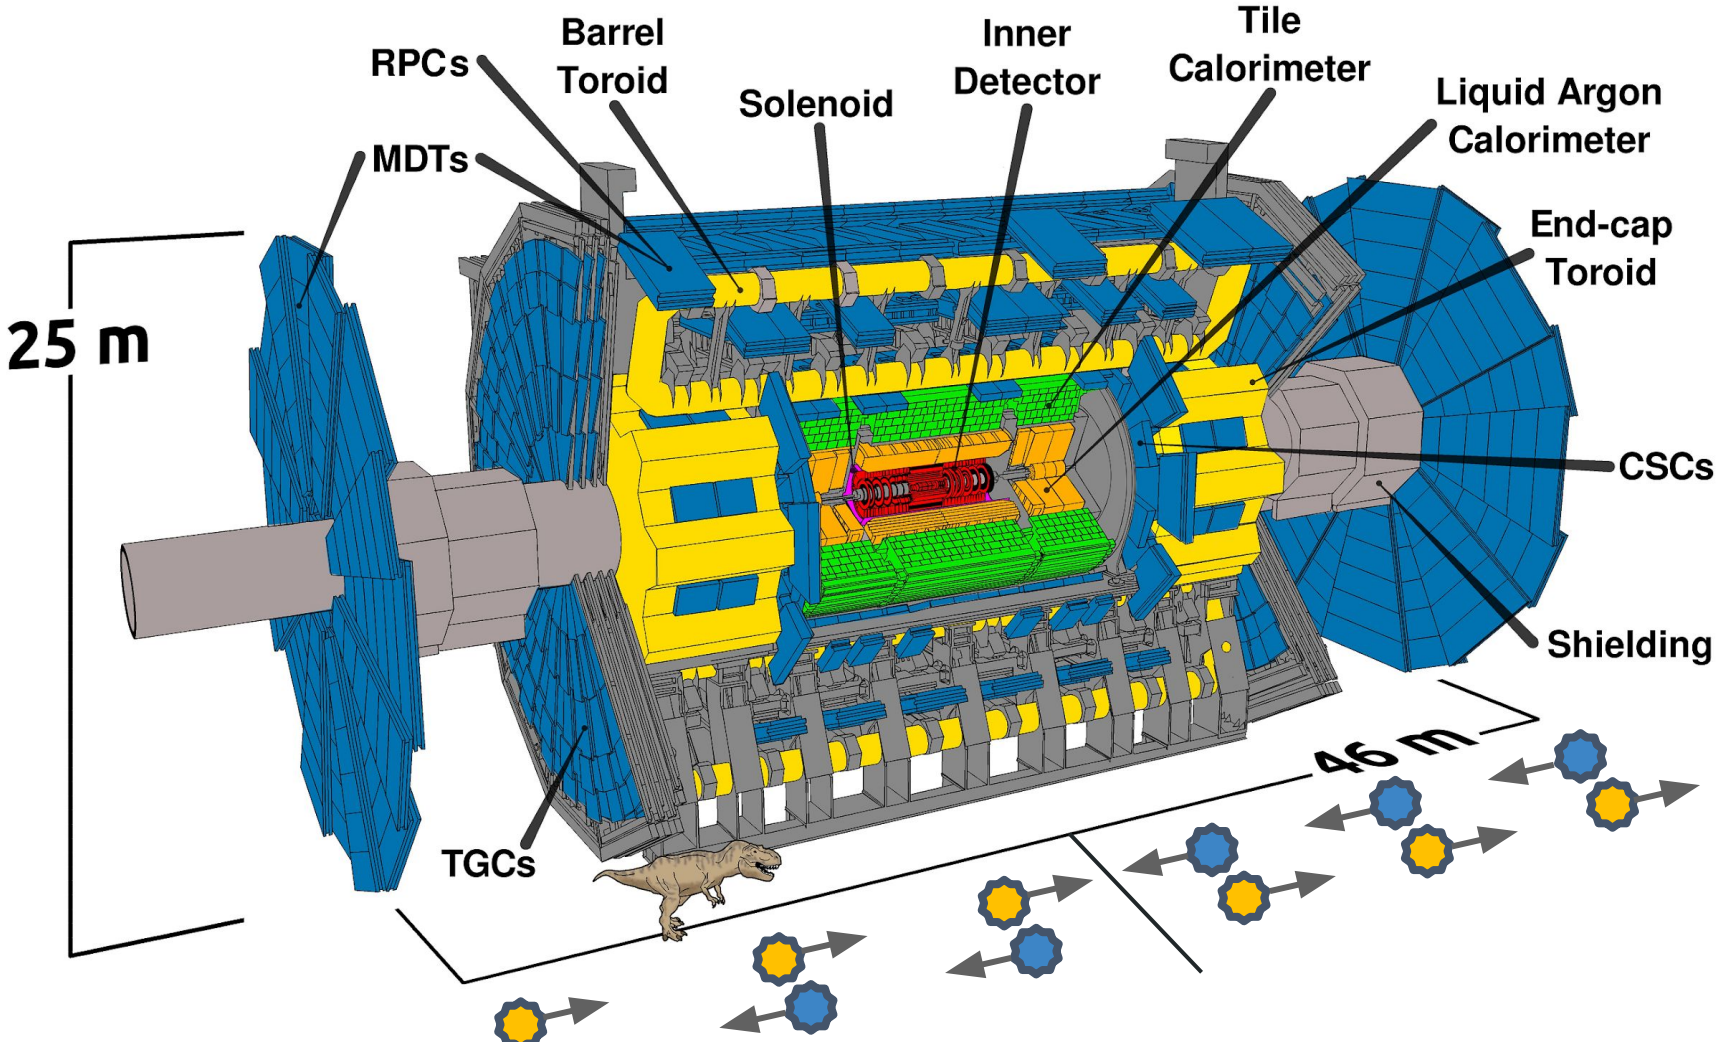
RPCs
Barrel Toroid
Solenoid
Inner Detector
Tile Calorimeter
Liquid Argon Calorimeter
MDTs
End-cap Toroid
25 m
CSCs
Shielding
46 m
TGCs

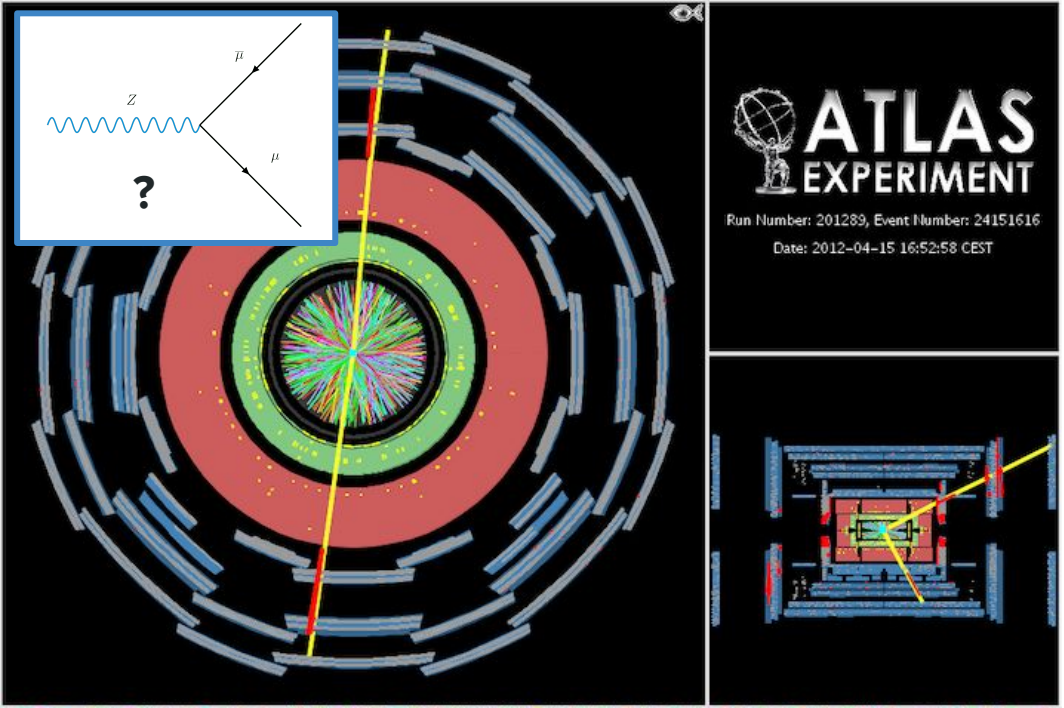

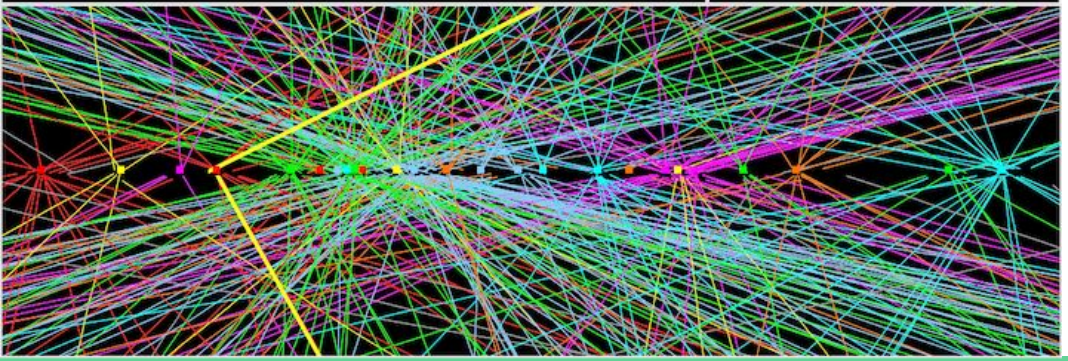

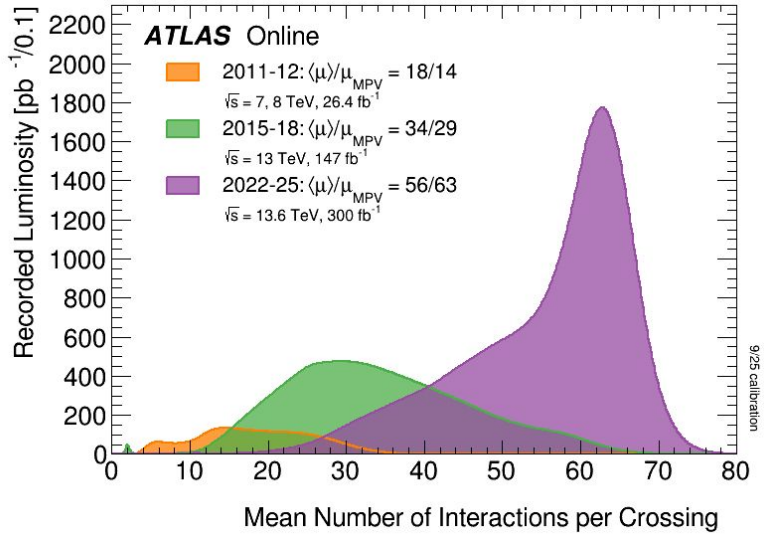

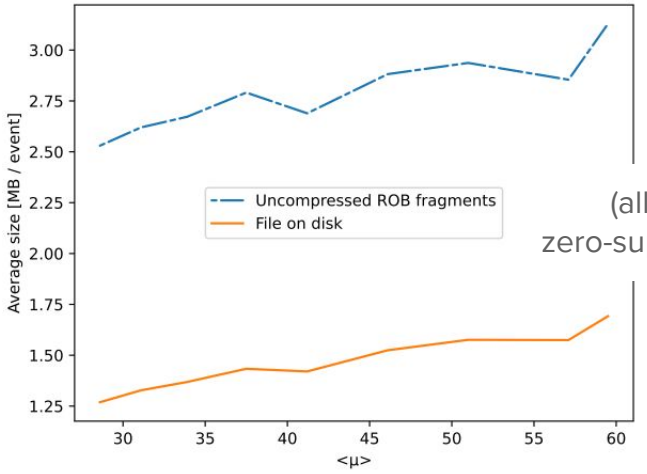

(all after zero-suppression)

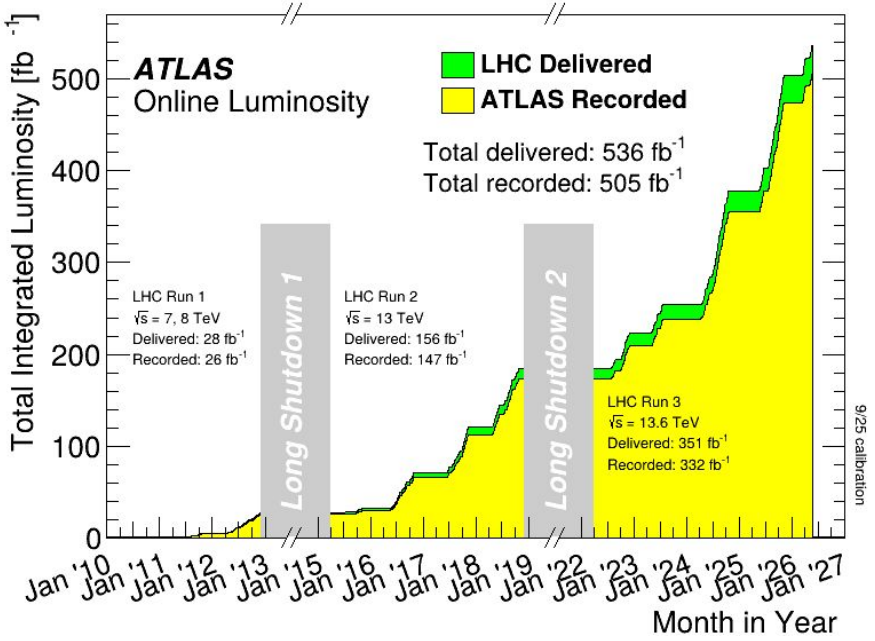
ATLAS
Online Luminosity
LHC Delivered
ATLAS Recorded
Total delivered: 536 fb-1
Total recorded: 505 fb-1
LHC Run 1
√s = 7, 8 TeV
Delivered: 28 fb-1
Recorded: 26 fb-1
Long Shutdown 1
LHC Run 2
√s = 13 TeV
Delivered: 156 fb-1
Recorded: 147 fb-1
Long Shutdown 2
LHC Run 3
√s = 13.6 TeV
Delivered: 351 fb-1
Recorded: 332 fb-1
Total Integrated Luminosity [fb-1]
Month in Year
9/25 calibration

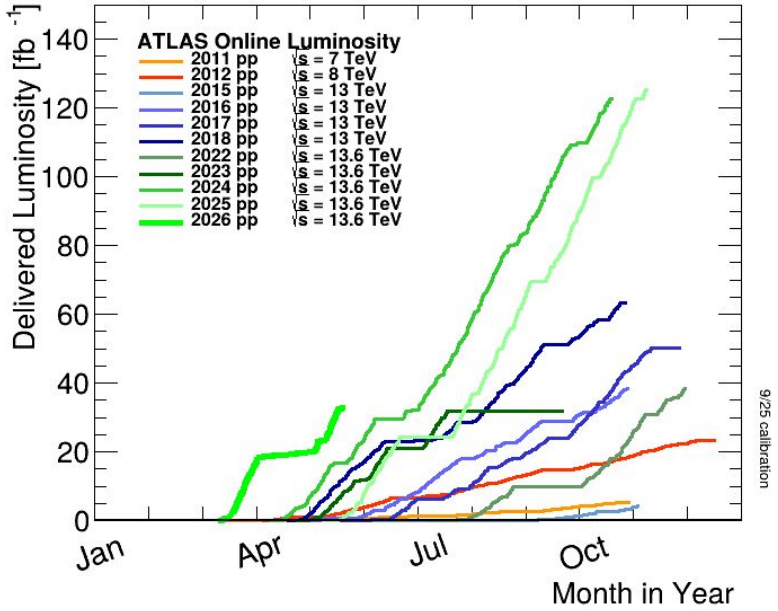
ATLAS Online Luminosity
2011 pp √s = 7 TeV
2012 pp √s = 8 TeV
2015 pp √s = 13 TeV
2016 pp √s = 13 TeV
2017 pp √s = 13 TeV
2018 pp √s = 13 TeV
2022 pp √s = 13.6 TeV
2023 pp √s = 13.6 TeV
2024 pp √s = 13.6 TeV
2025 pp √s = 13.6 TeV
2026 pp √s = 13.6 TeV
Delivered Luminosity [fb-1]
Month in Year
9/25 calibration

General Overview

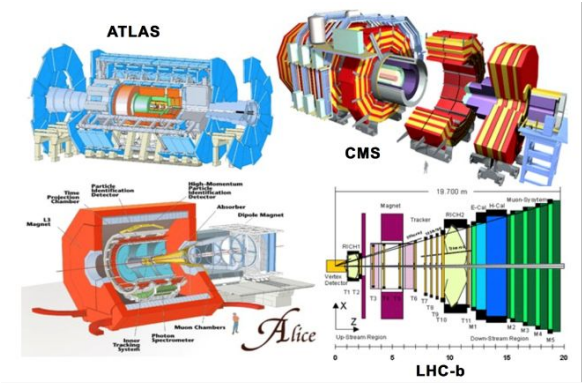

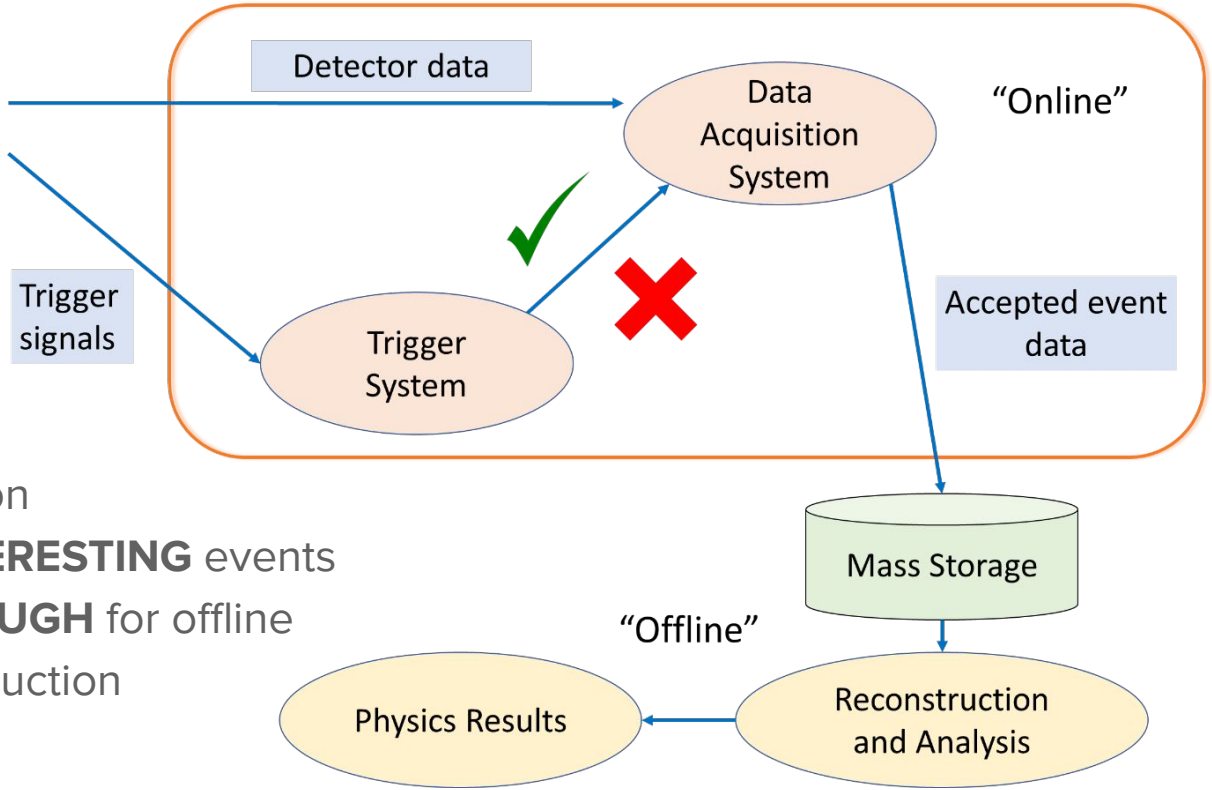

- Trigger must:
  - Make a **FAST** decision
  - Select the most **INTERESTING** events
  - Keep rate **LOW ENOUGH** for offline storage and reconstruction

Hardware Triggering

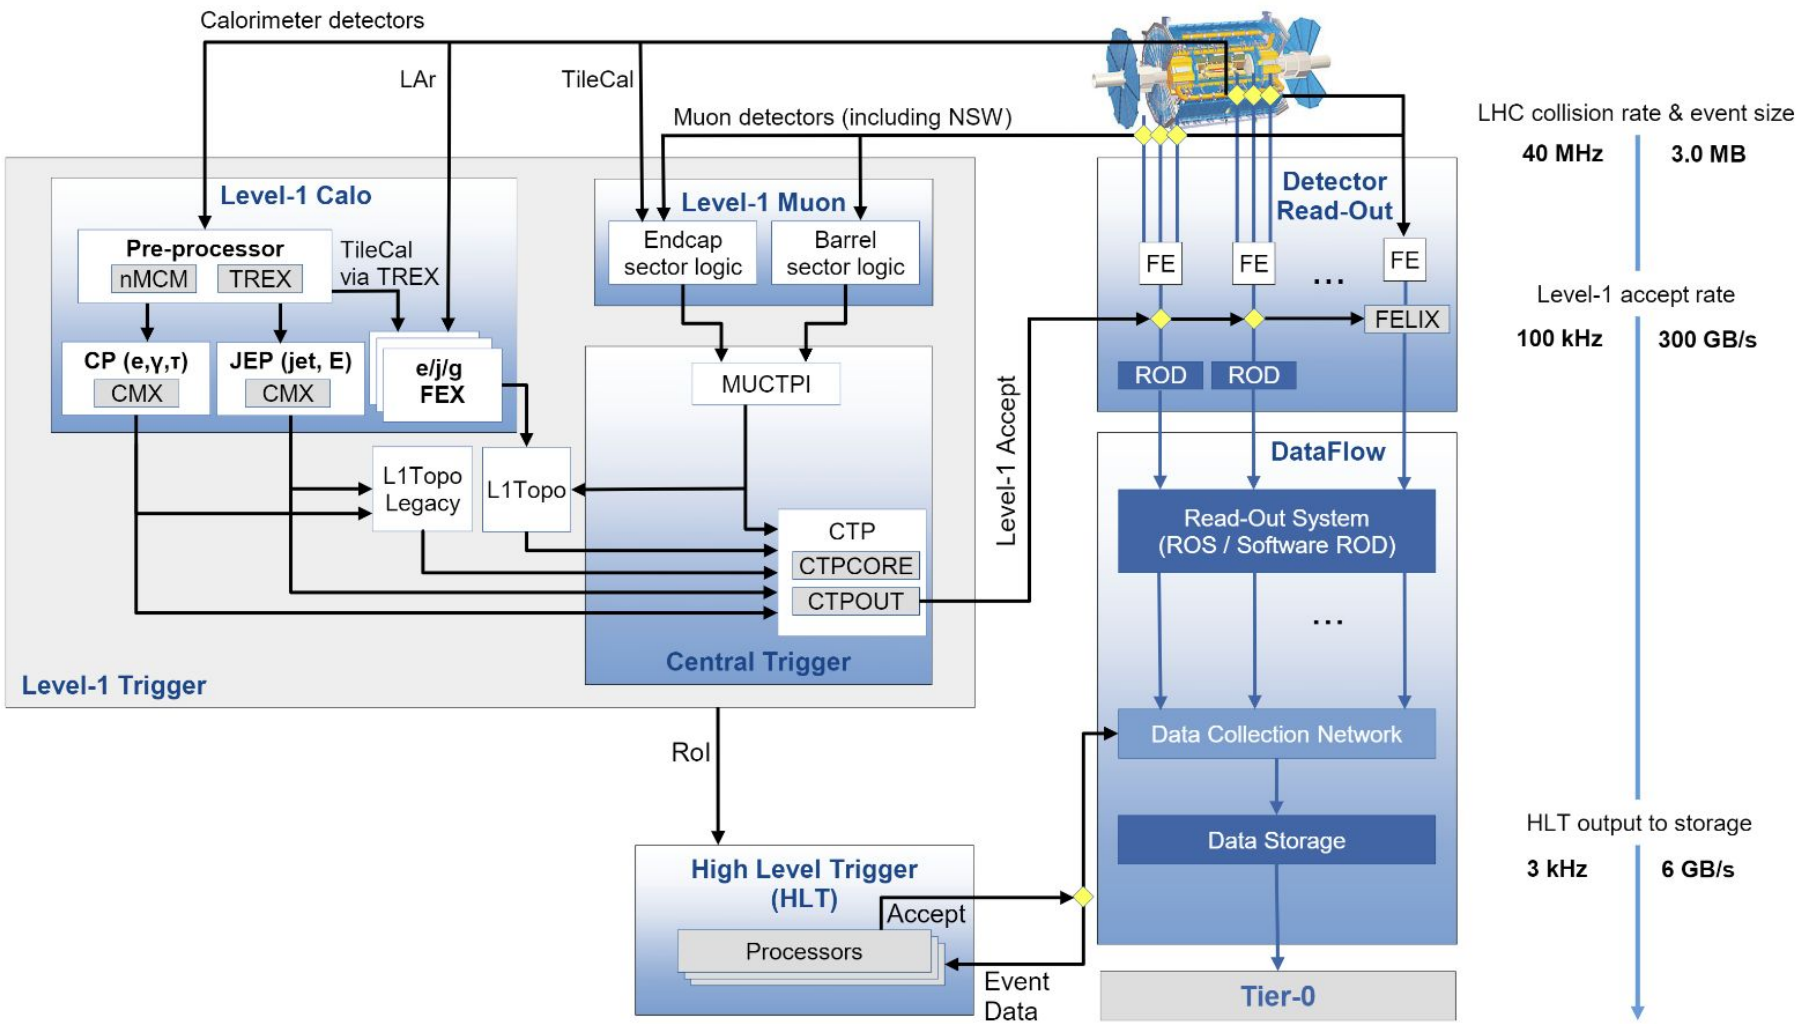
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 Accept
Detector Read-Out
FE
FE
...
FE
FELIX
ROD
ROD
DataFlow
Read-Out System (ROS / Software ROD)
...
Data Collection Network
Data Storage
Tier-0
RoI
High Level Trigger (HLT)
Processors
Accept
Event Data
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 accept rate
100 kHz
300 GB/s
HLT output to storage
3 kHz
6 GB/s

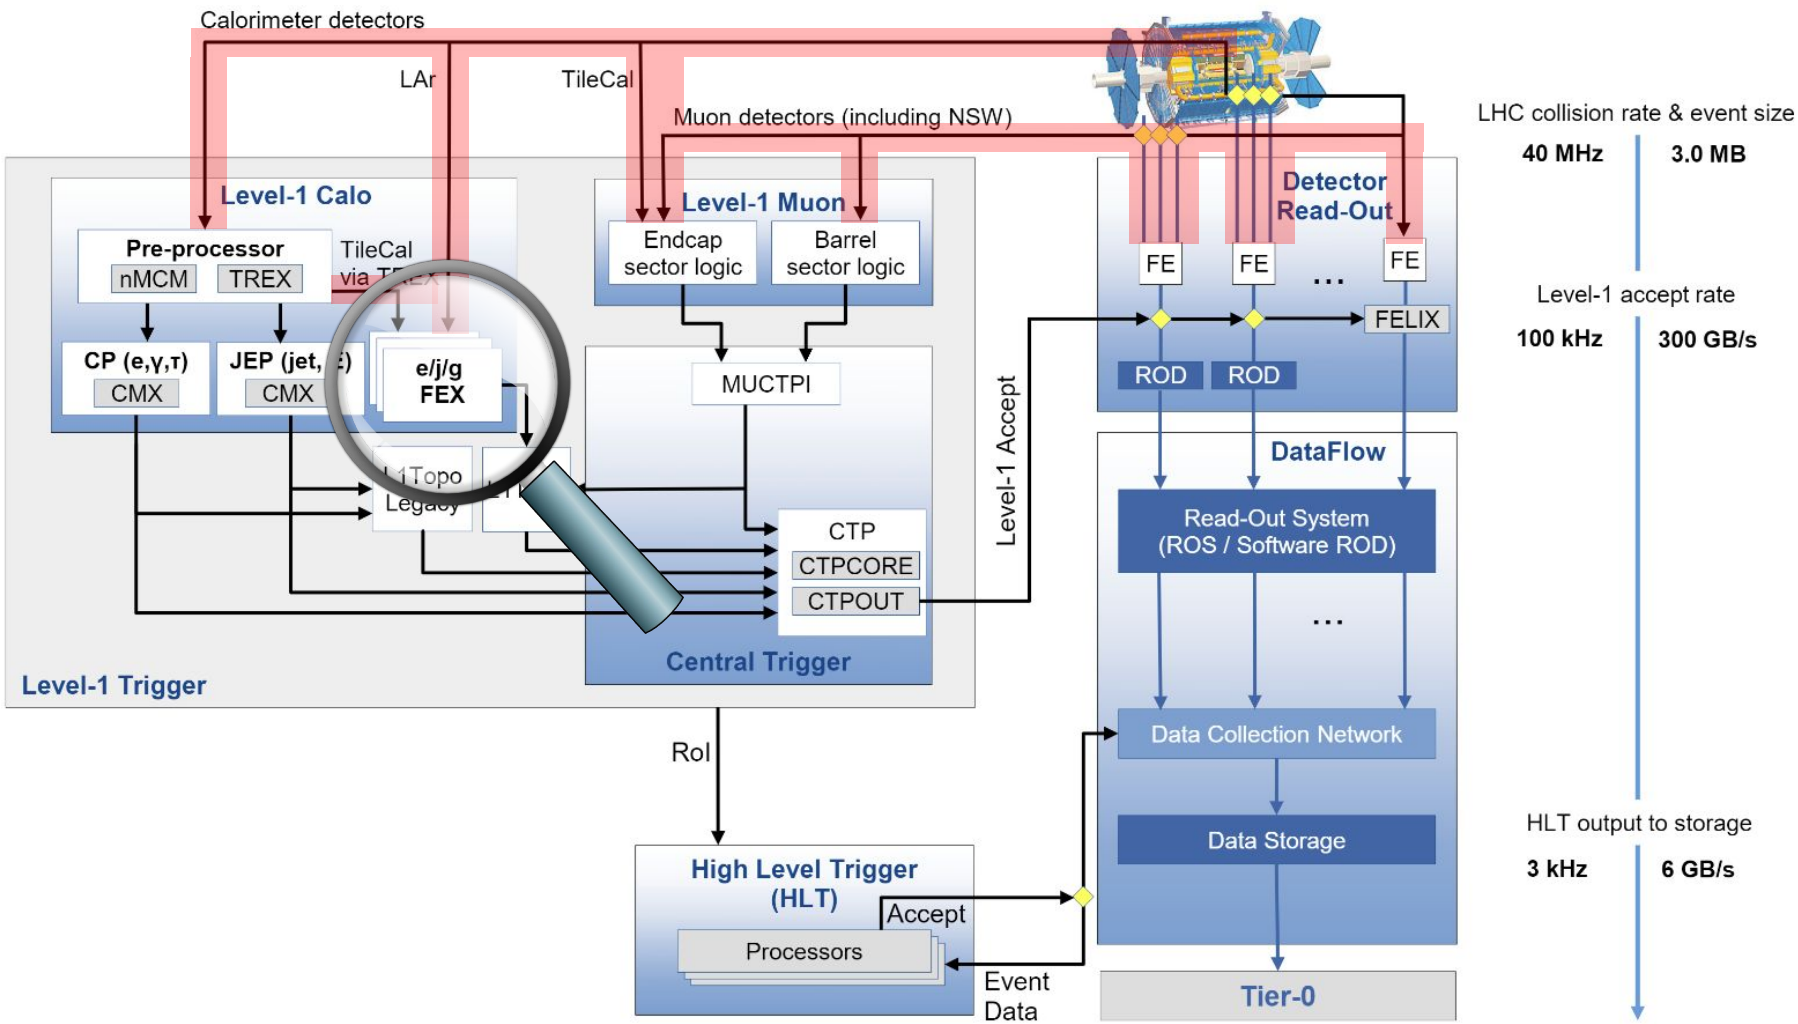
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal
via TREX
CP (e,γ,τ)
CMX
JEP (jet,
CMX
e/j/g
FEX
L1Topo
Legacy
Level-1 Muon
Endcap
sector logic
Barrel
sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 Accept
Detector
Read-Out
FE
FE
...
FE
FELIX
ROD
ROD
DataFlow
Read-Out System
(ROS / Software ROD)
...
Data Collection Network
Data Storage
Tier-0
RoI
High Level Trigger
(HLT)
Processors
Accept
Event
Data
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 accept rate
100 kHz
300 GB/s
HLT output to storage
3 kHz
6 GB/s

- The signals in the detector are **digitized** and **buffered** in a **front-end circular pipeline.**
- The race is now on to decide if we want to *keep the event* **before** it gets overwritten in the buffer.
- Need to make and return our decision within **2.5 µs**.
- *That's only 750m travel time at the speed of light…*
- The L1 trigger is released primarily via **pipelined algorithms** implemented on **FPGAs**.

Make a FAST decision

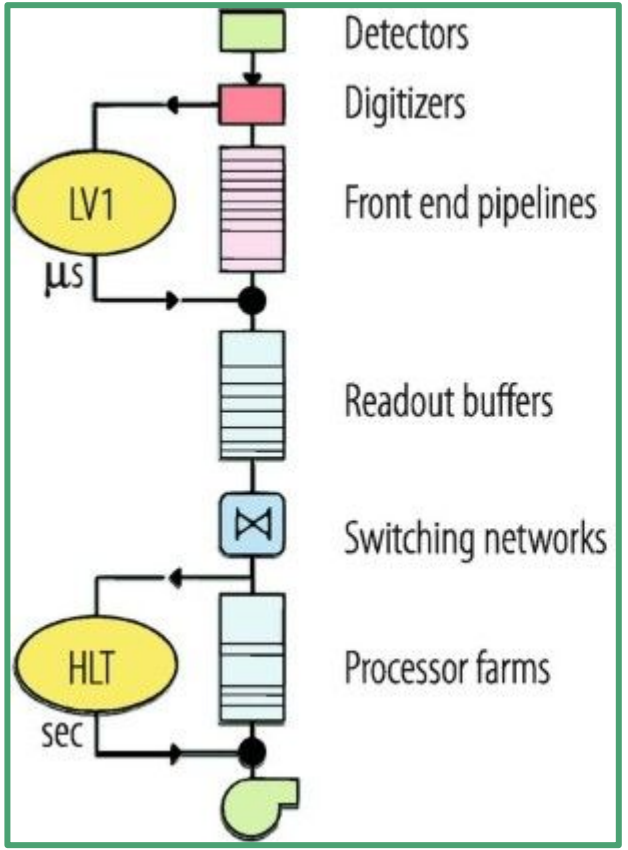

Electromagnetic Feature EXtractor (eFEX)

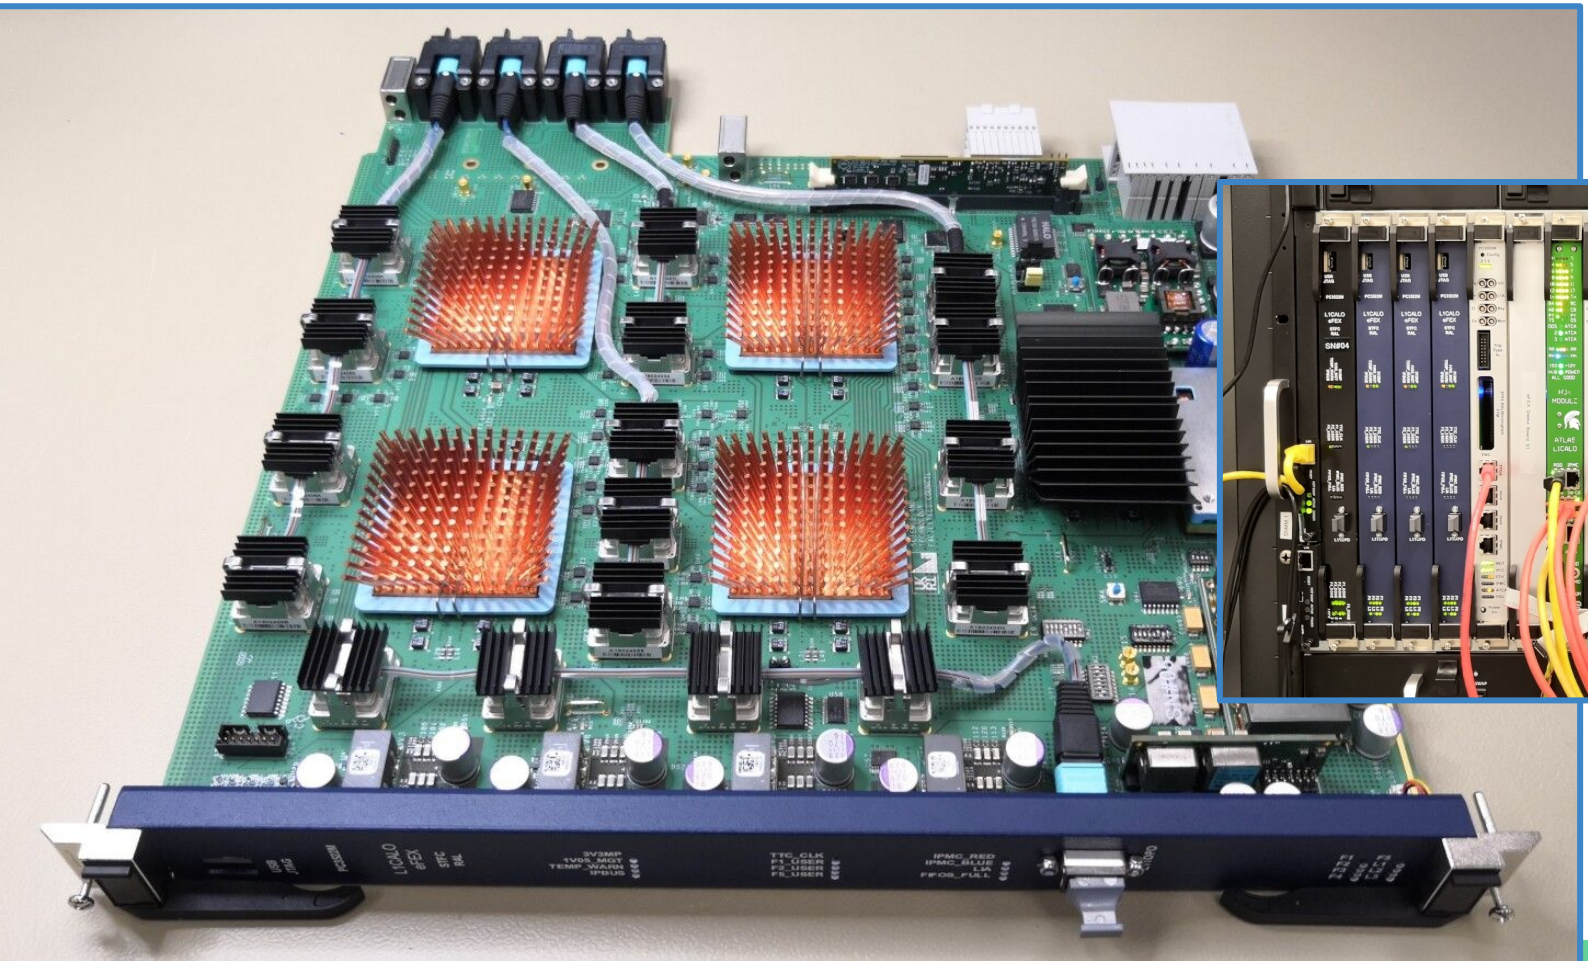

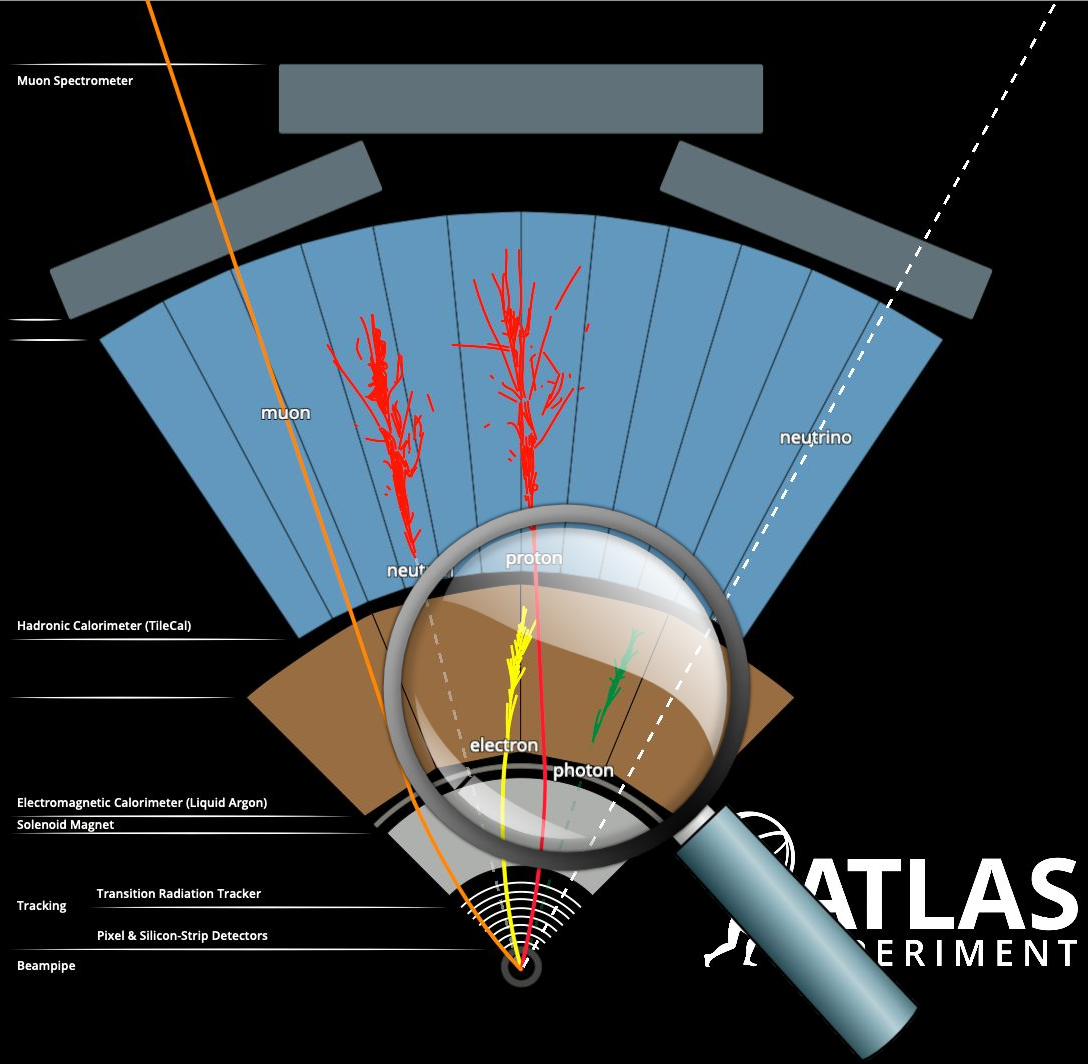
Muon Spectrometer
muon
neutrino
proton
Hadronic Calorimeter (TileCal)
electron
photon
Electromagnetic Calorimeter (Liquid Argon)
Solenoid Magnet
Tracking
Transition Radiation Tracker
Pixel & Silicon-Strip Detectors
Beampipe
ATLAS

- Prior to 2022, the L1 trigger only had a single EM energy reading per 0.1 x 0.1 region.
- In our upgraded system, eFEX gets **10 EM energy readings** from the same region.
- The eFEX system cuts on **shower shape variables**.
  - **Shower width** in the **2nd electromagnetic calorimeter layer**.
  - **Fraction of energy in the hadronic calorimeter**.
  - **Shower width** in the **1st electromagnetic layer**.

$$R_\eta = 1 - \frac{E_{\mathrm{T},3\times2}}{E_{\mathrm{T},7\times3}}$$

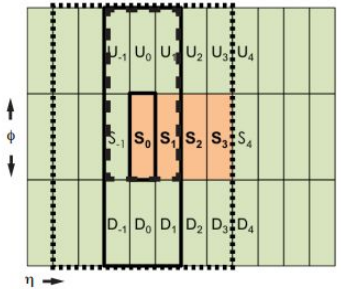

$$R_{\mathrm{had}} = \frac{E_{\mathrm{T,had}}}{E_{\mathrm{T,EM}} + E_{\mathrm{T,had}}}$$

$$w_{\mathrm{s,tot}} = \sqrt{\frac{\sum E_{\mathrm{T},i} \times (i - i_{\max})^2}{\sum E_{\mathrm{T},i}}}$$

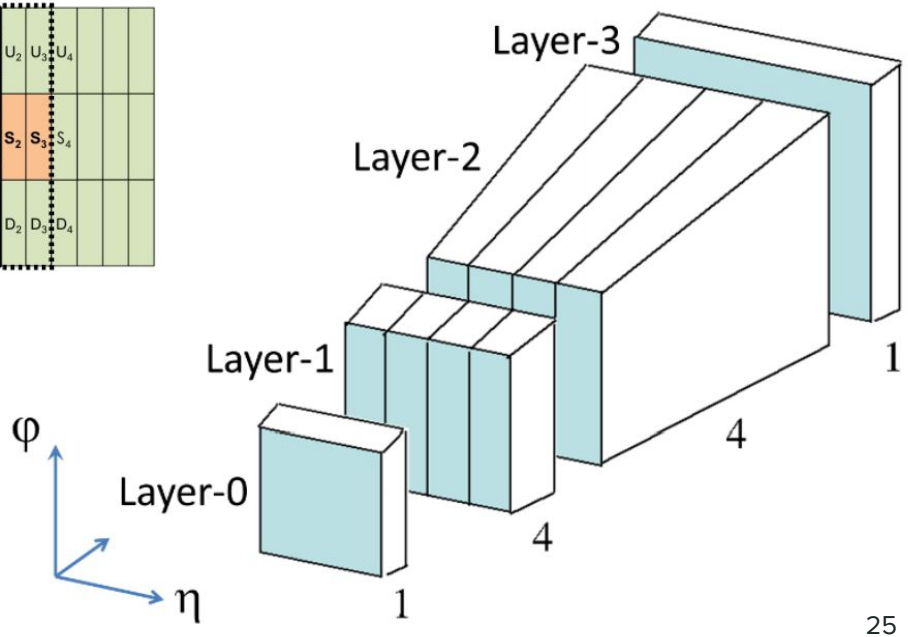

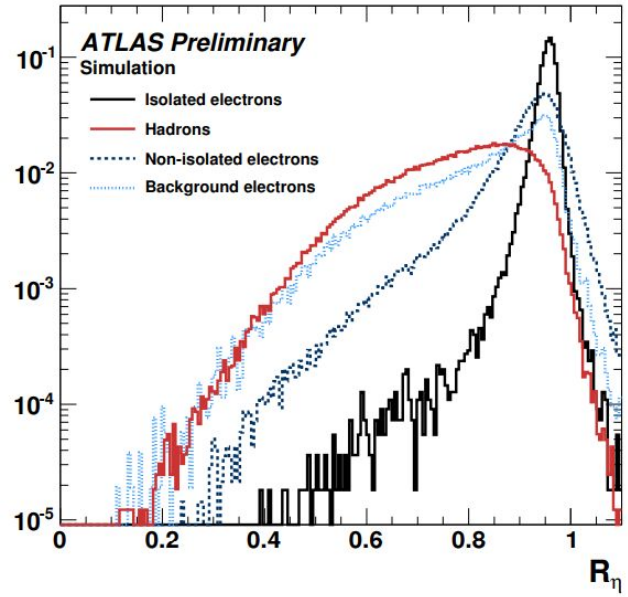

“Good” isolated electrons are (probably) from a W or Z

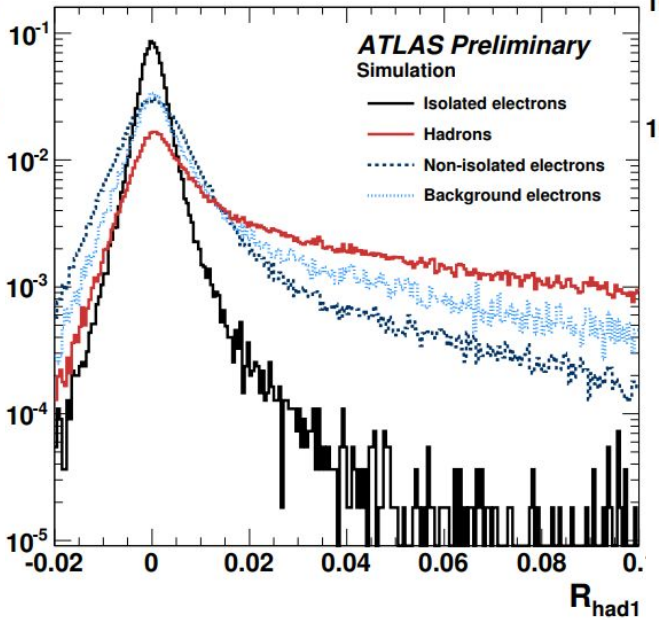

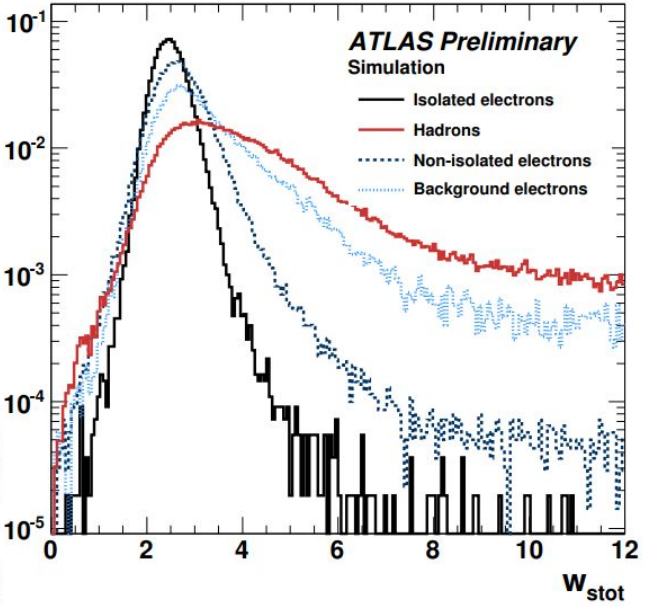

Select the most INTERESTING events

Keep the rate LOW ENOUGH

- eFEX: Both **purer** (*greater fraction of triggered candidates are actually electrons*) …
- … and more **efficient** (*greater fraction of electrons are triggered*) than the prior system.
- *(Nomenclature: EM22VHI ~= eEM26M)*

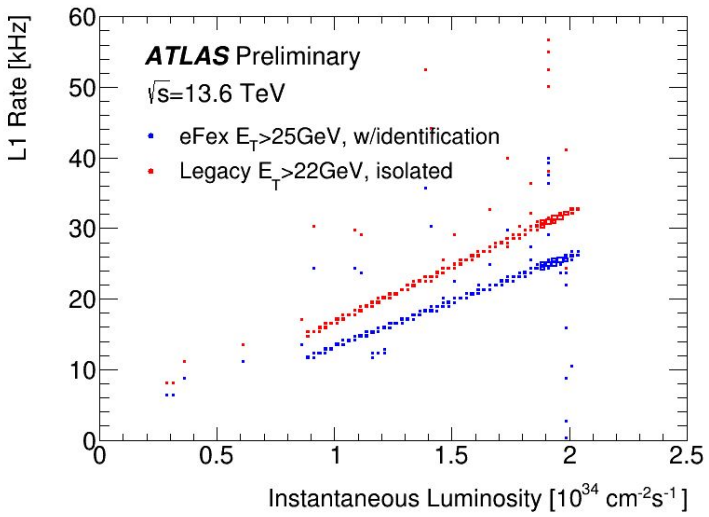

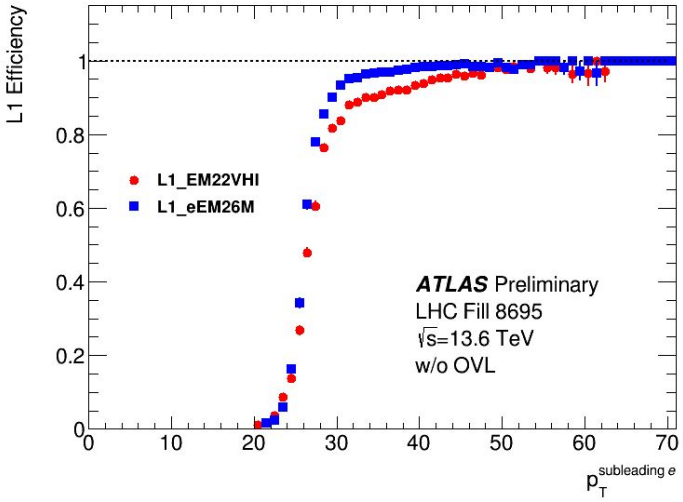

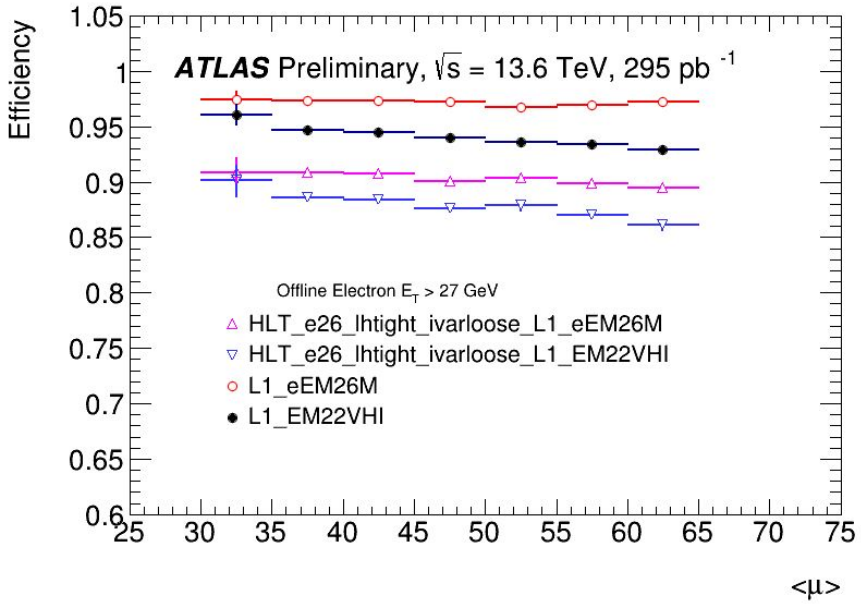

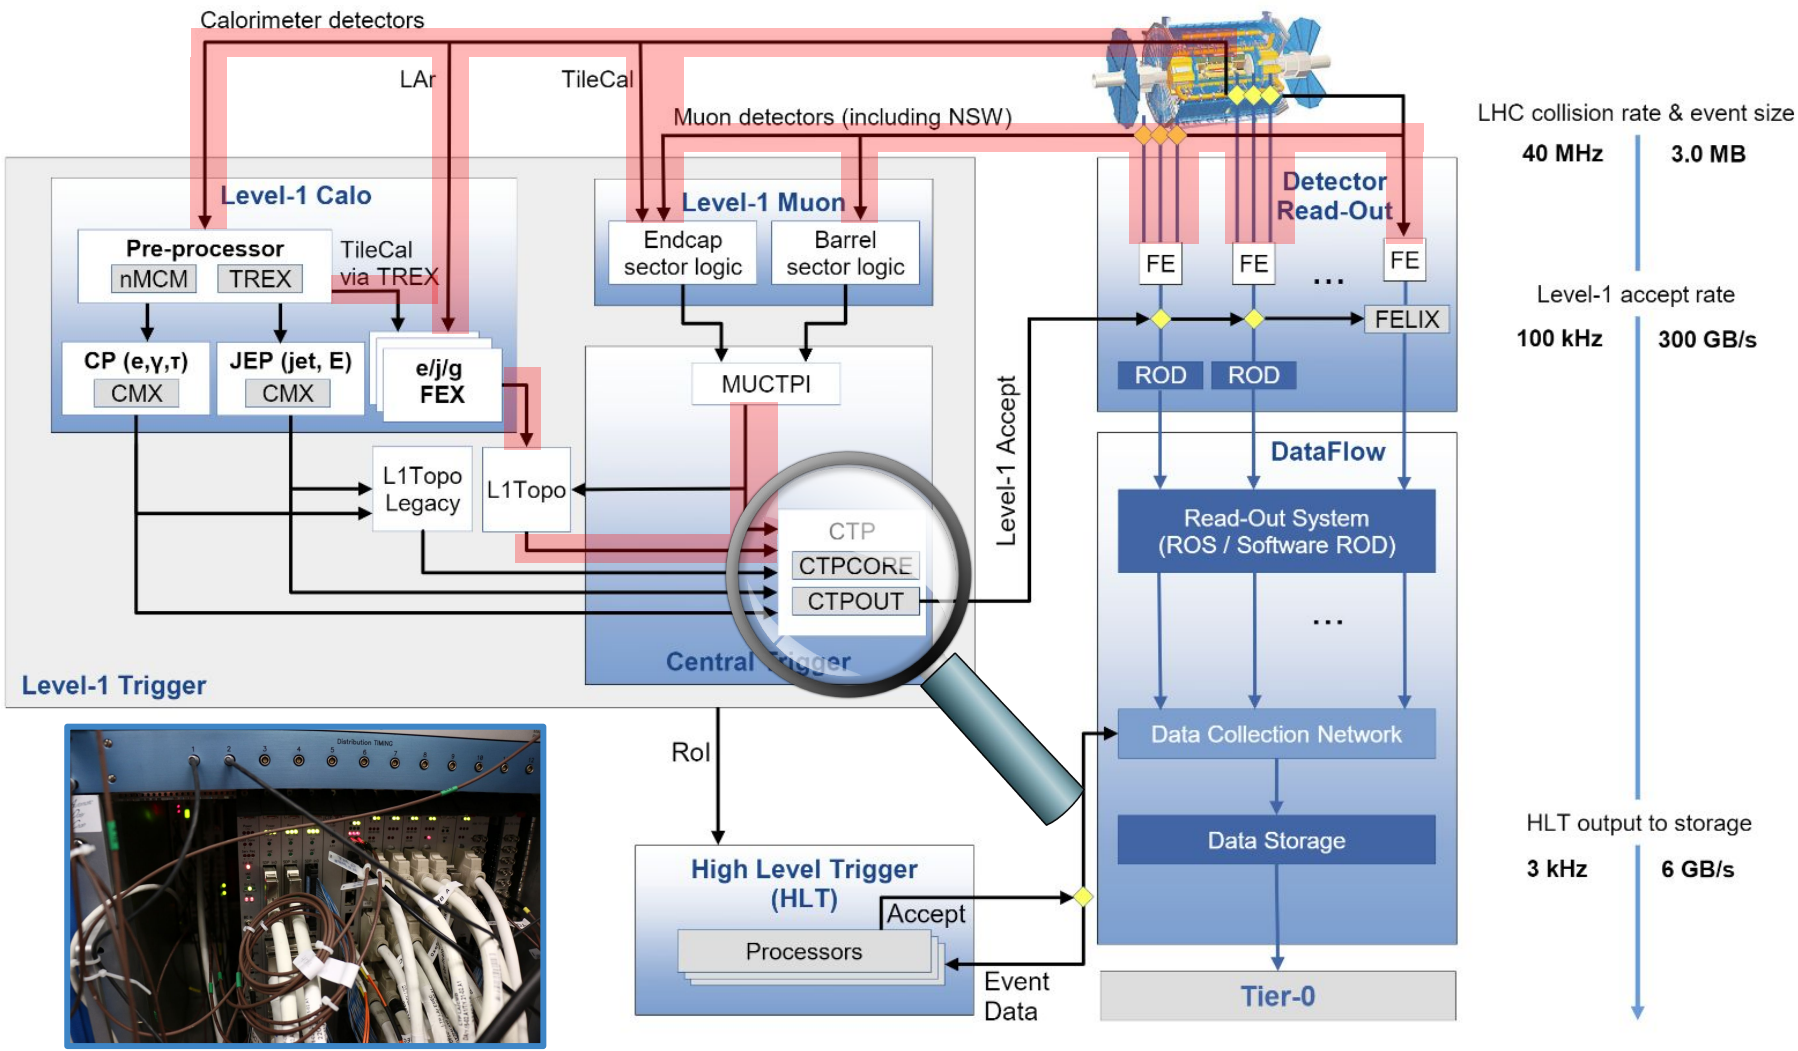
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 Accept
Detector Read-Out
FE
FE
...
FE
FELIX
ROD
ROD
DataFlow
Read-Out System (ROS / Software ROD)
...
Data Collection Network
Data Storage
Tier-0
RoI
High Level Trigger (HLT)
Processors
Accept
Event Data
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 accept rate
100 kHz
300 GB/s
HLT output to storage
3 kHz
6 GB/s

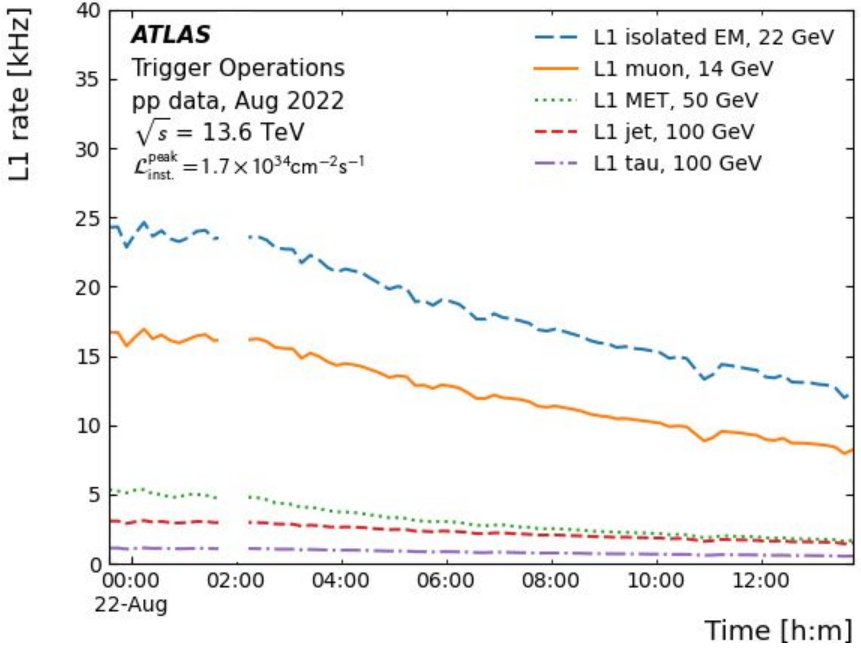

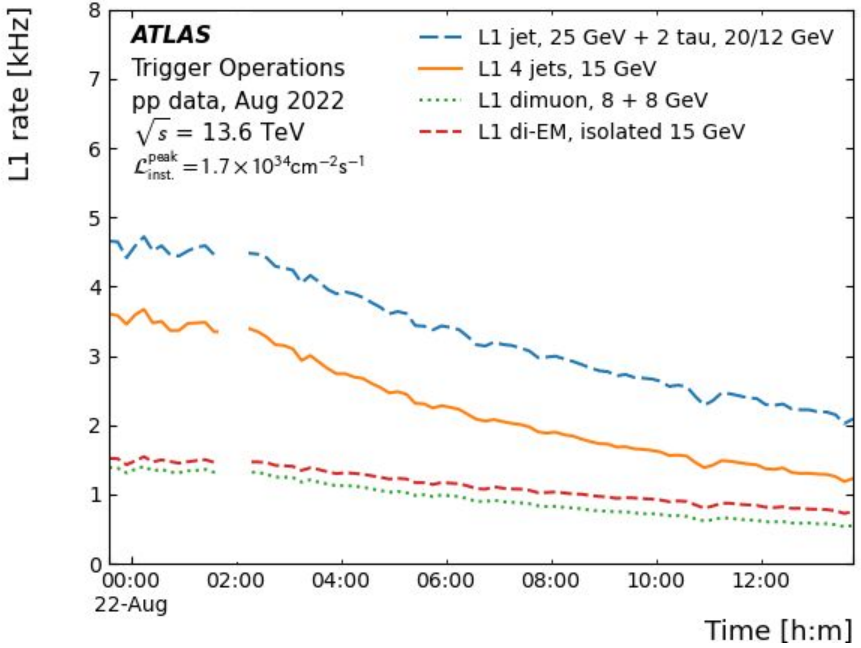

- A ***Trigger Menu*** converts the collaboration's physics goals into **momentum & identification requirements** on objects identified by the L1 trigger.
- ATLAS' L1 trigger limited to a maximum of **100 kHz** (around 250-400 rejection).
- The **Central Trigger Processor** emits **Accept** signals for events above threshold.

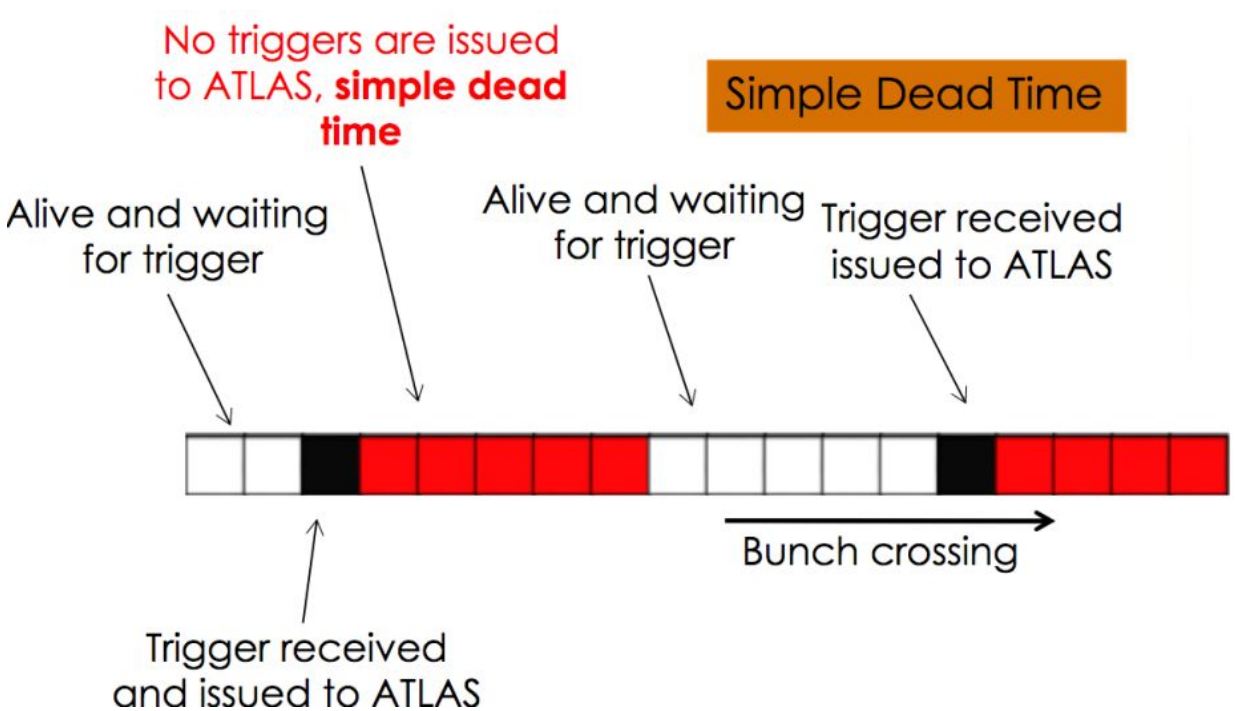

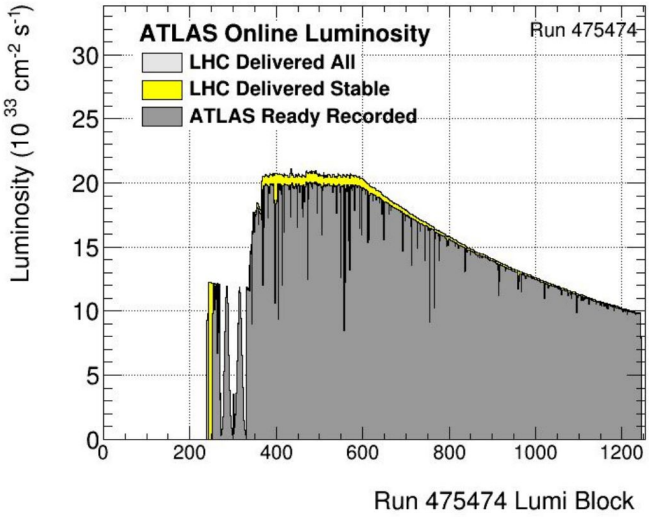

- Once a trigger signal issued, the data must be **copied off of the front end pipelines**.
- **No triggers can be issued in the following bunch crossings**.
- The amount of luminosity **recorded by the experiment** is less than the amount of data **delivered to the experiment** by the LHC.

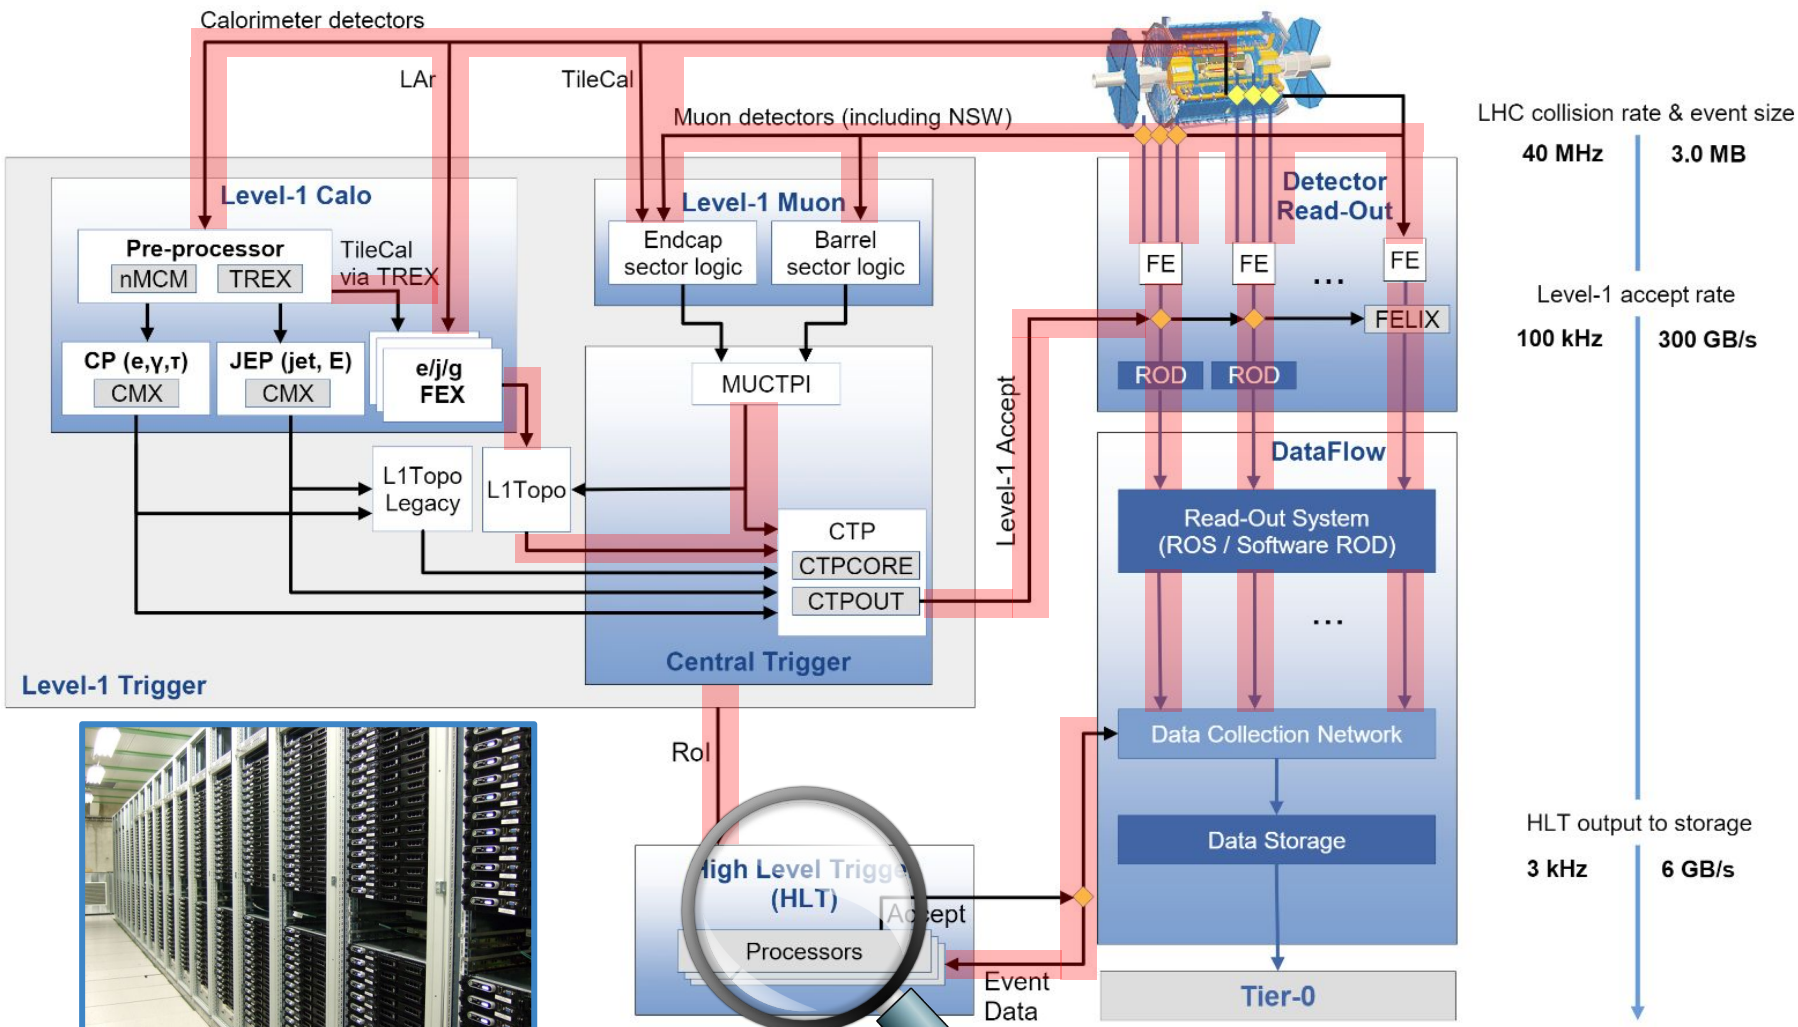
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
L1Topo Legacy
L1Topo
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 Accept
Detector Read-Out
FE
FE
...
FE
FELIX
ROD
ROD
DataFlow
Read-Out System (ROS / Software ROD)
...
Data Collection Network
Data Storage
Tier-0
RoI
High Level Trigger (HLT)
Accept
Processors
Event Data
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 accept rate
100 kHz
300 GB/s
HLT output to storage
3 kHz
6 GB/s

Software Triggering

An Abstraction

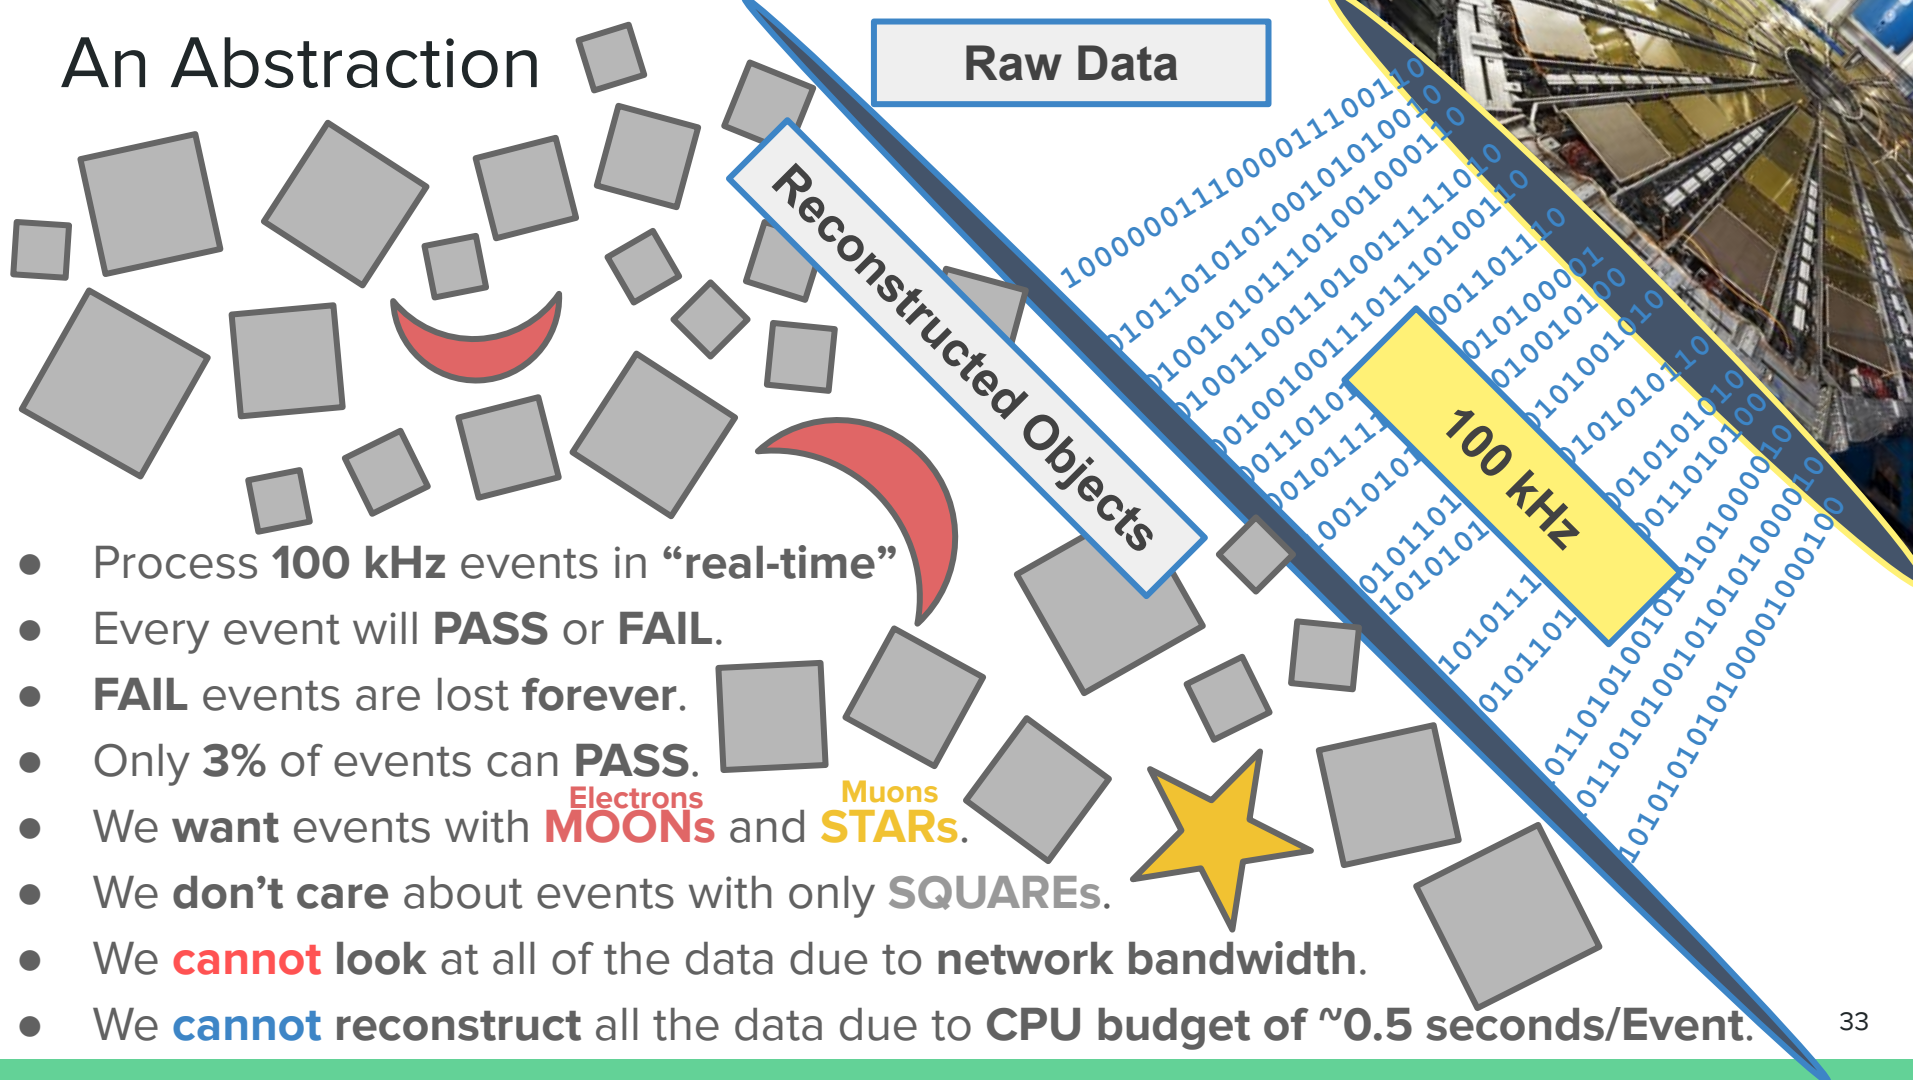

- Process **100 kHz** events in **“real-time”**
- Every event will **PASS** or **FAIL**.
- **FAIL** events are lost **forever**.
- Only **3%** of events can **PASS**.
- We **want** events with Electrons **MOONs** and Muons **STARs**.
- We **don’t care** about events with only **SQUAREs**.
- We **cannot look** at all of the data due to **network bandwidth**.
- We **cannot reconstruct** all the data due to **CPU budget of ~0.5 seconds/Event**.

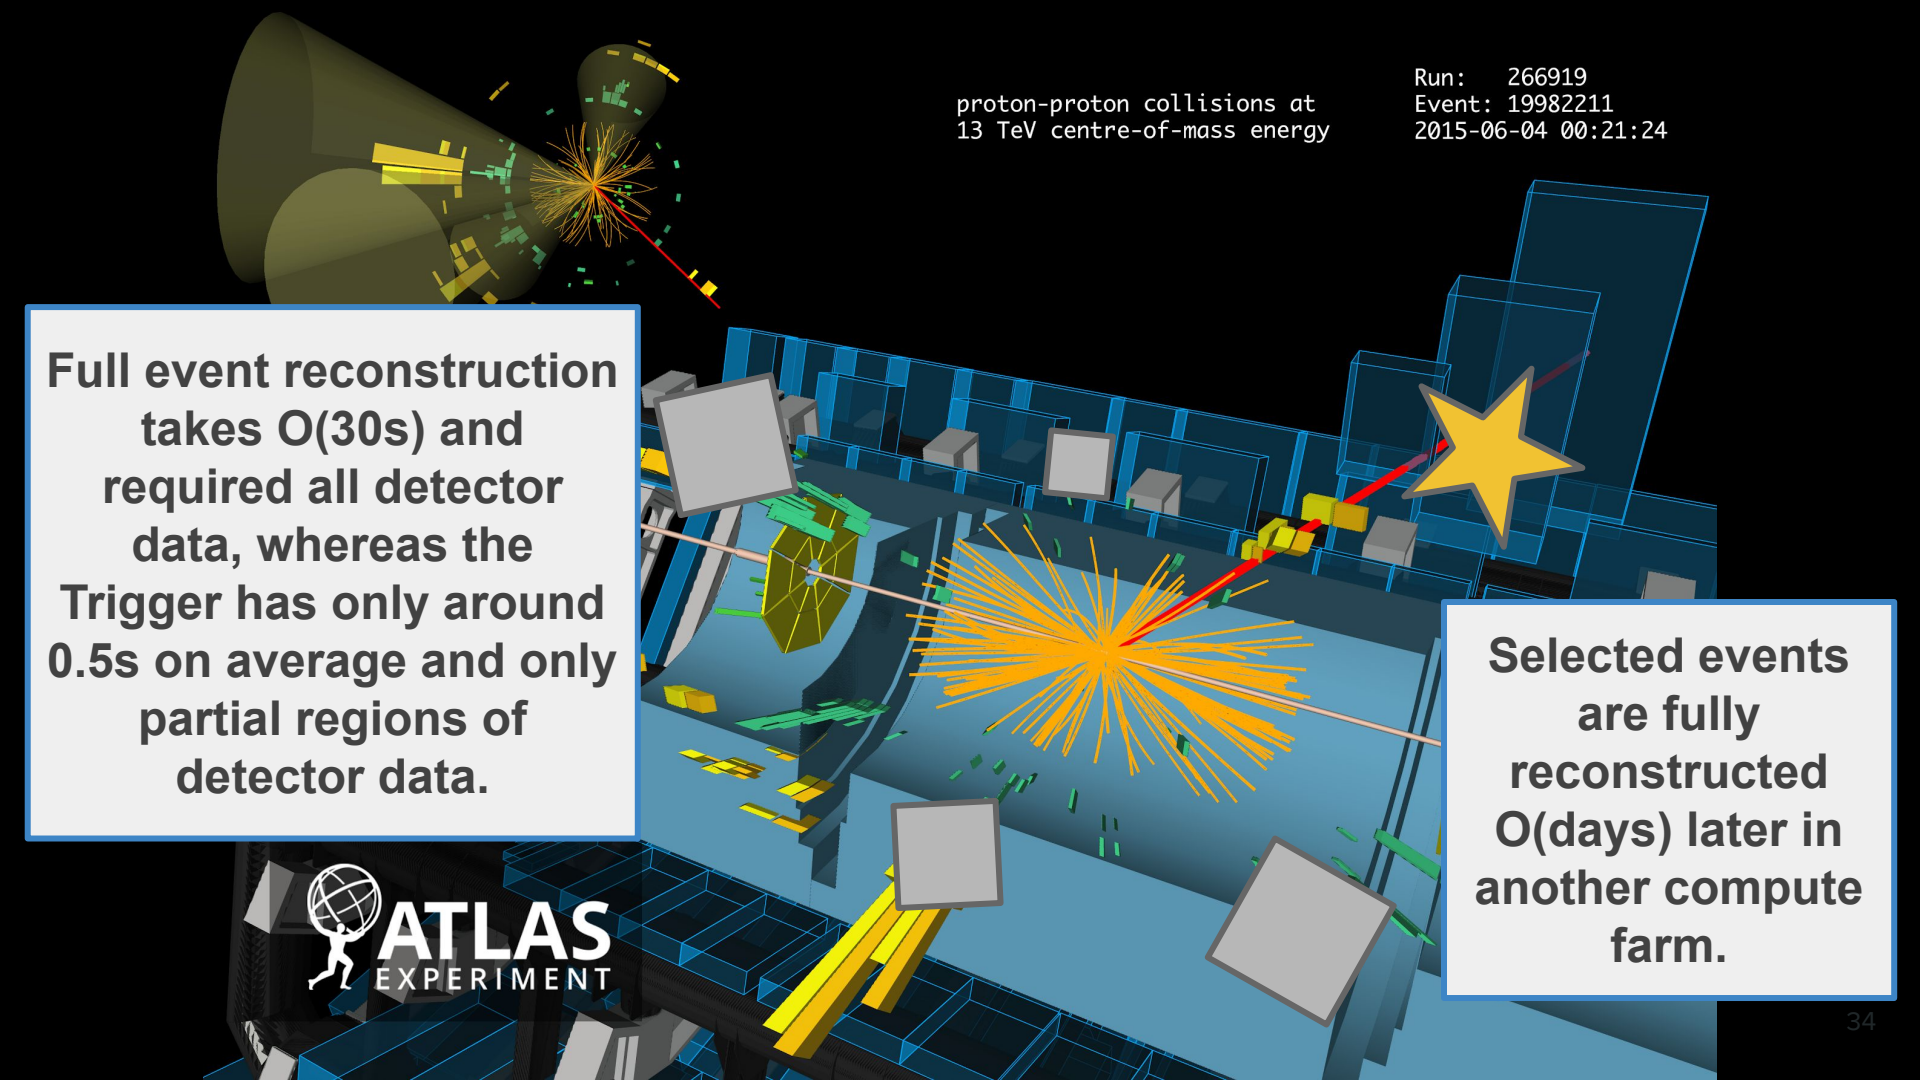
proton-proton collisions at
13 TeV centre-of-mass energy
Run: 266919
Event: 19982211
2015-06-04 00:21:24
Full event reconstruction takes O(30s) and required all detector data, whereas the Trigger has only around 0.5s on average and only partial regions of detector data.
Selected events are fully reconstructed O(days) later in another compute farm.
ATLAS
EXPERIMENT

Key Principles of the Trigger

- **Regional Reconstruction**
  - We **cannot** look at all of every event due to bandwidth.
  - Restrict to running **reconstruction algorithms** within **Regions of Interest**, identified in the 1st level hardware trigger.
- **Early Rejection**
  - Split reconstruction up into multiple **Steps,** called **Chains**.
  - **Filtering** occurs after each **Step** via **Hypothesis Algorithms**
  - **Early** steps are **fast**, but **coarse**.
  - **Later** steps **take more time**, but are **detailed**.
  - **Stop** reconstructing an **object** as soon as it fails a selection at the end of a **Step**.
  - **Stop** reconstructing the **event** when all objects are **rejected**.

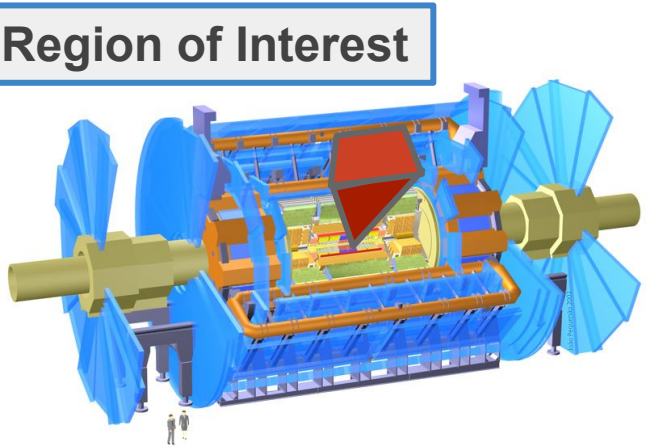

High Level Trigger Example

- Suppose three **types of selection**:  and  . Each **Chain** will follow **one of these three paths**, with the chain's configuration controlling object quality & object size requirements.
- We build a **dependency graph** of the algorithms required to perform the reconstruction. This is divided into different **steps**.
- Three classes of algorithm are used to **control the execution**.
  - **Filter Algorithm** **Always runs** at the start of each step. Responsible for implementing **Early Rejection**. Returns a boolean **Filter Decision**.
  - **'Input Maker' Algorithm** Provides concrete starting point for reconstruction algorithms. Responsible for restricting reconstruction to **Regions of Interest.**
  - **Hypothesis Algorithm** Executes **hypothesis testing** for all active **chains**. Provides input to next Step's Filter(s).
- Graph nodes evaluated **asynchronously** & in **parallel:** multi-threaded FW.

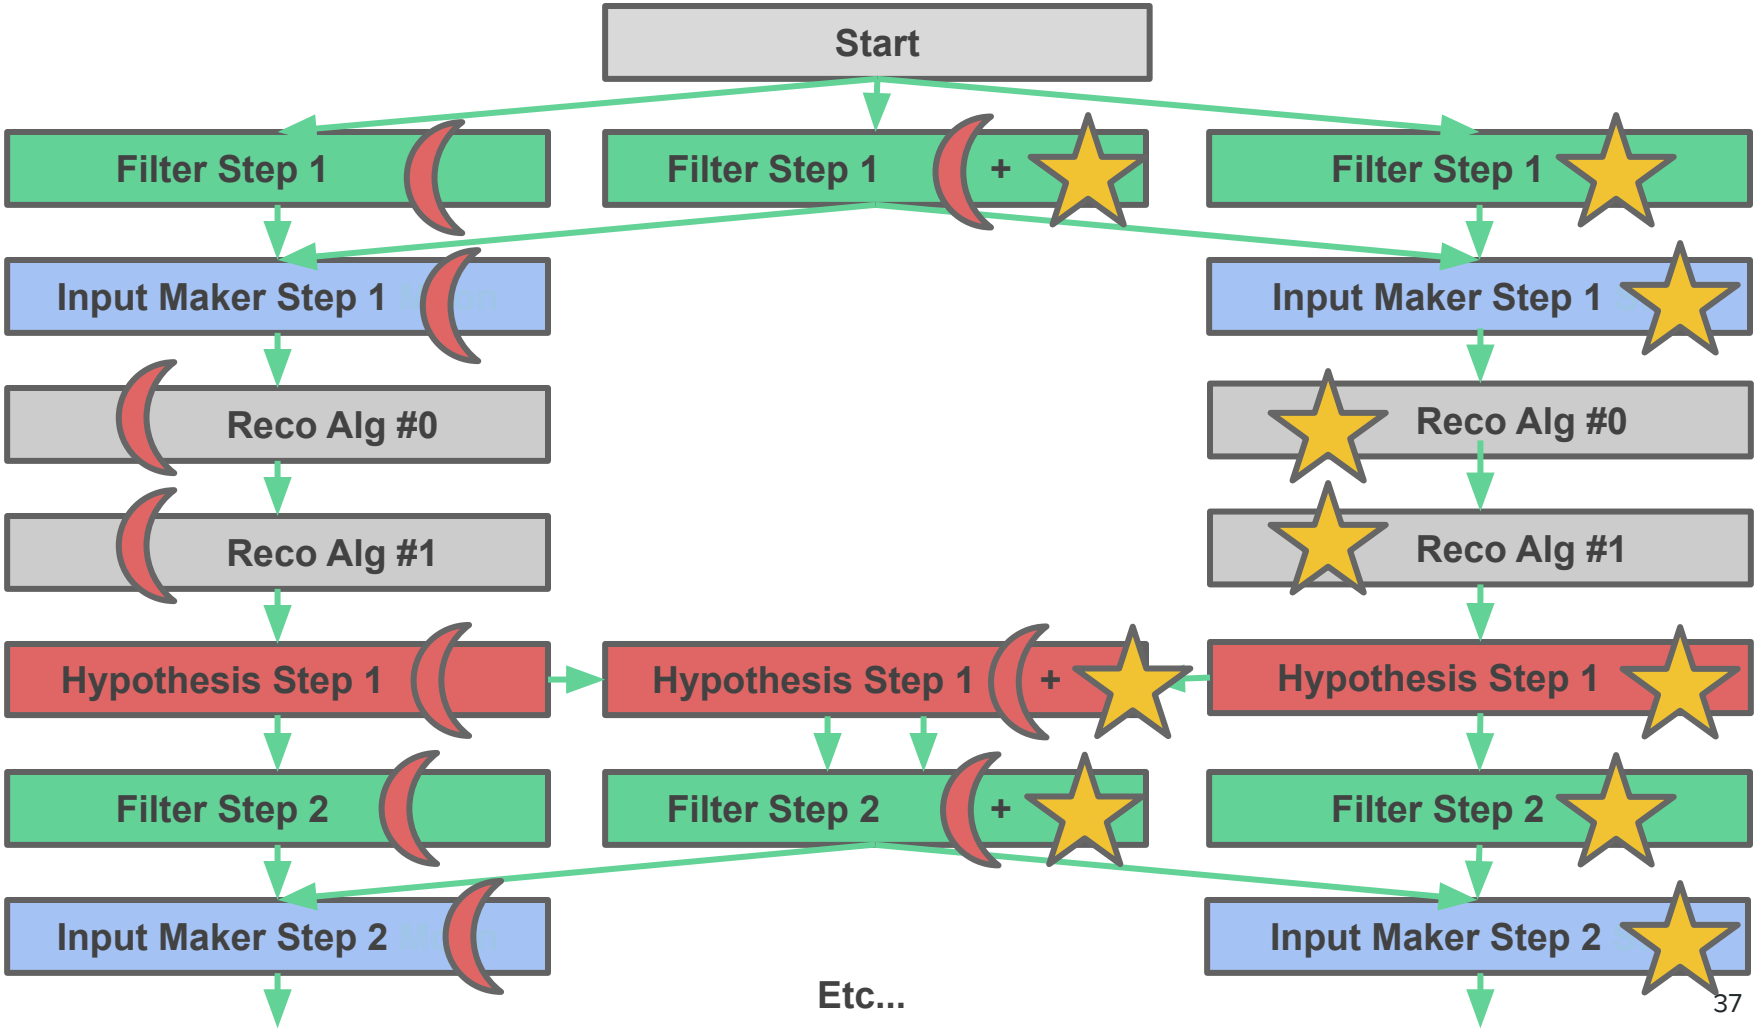
Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
Input Maker Step 1
Reco Alg #0
Reco Alg #0
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

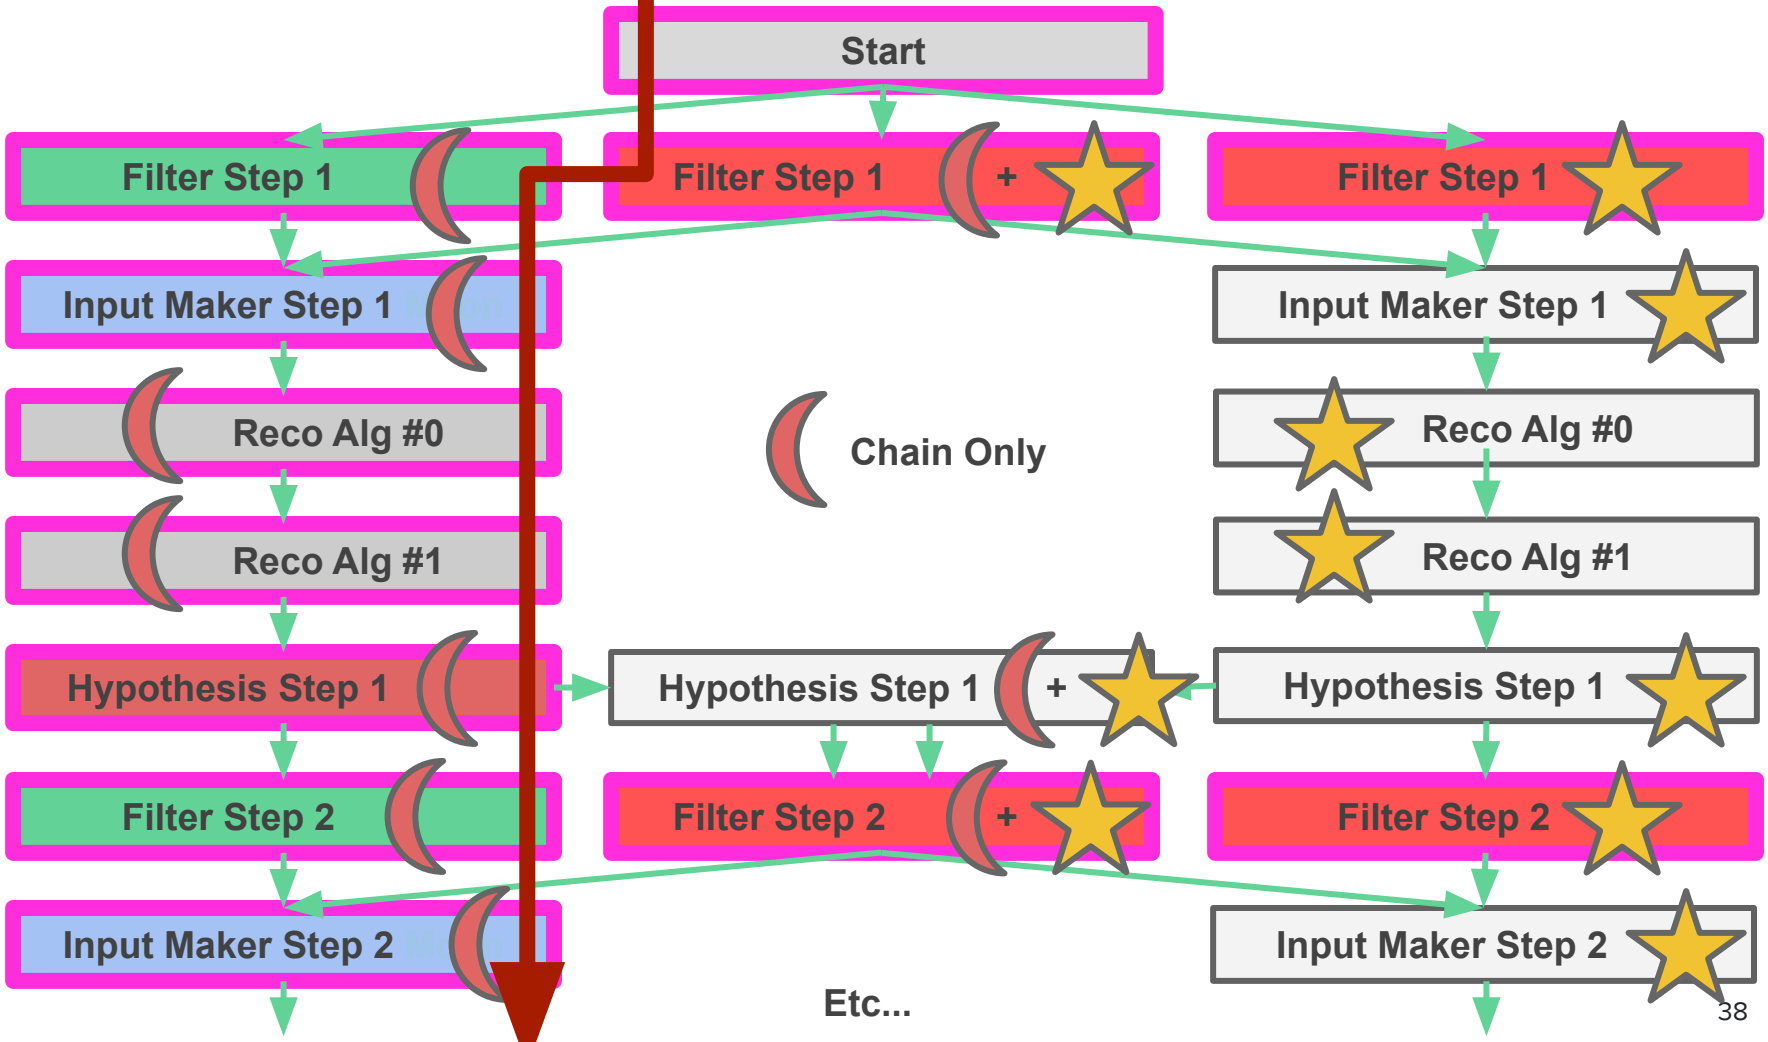
Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
Input Maker Step 1
Reco Alg #0
Chain Only
Reco Alg #0
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

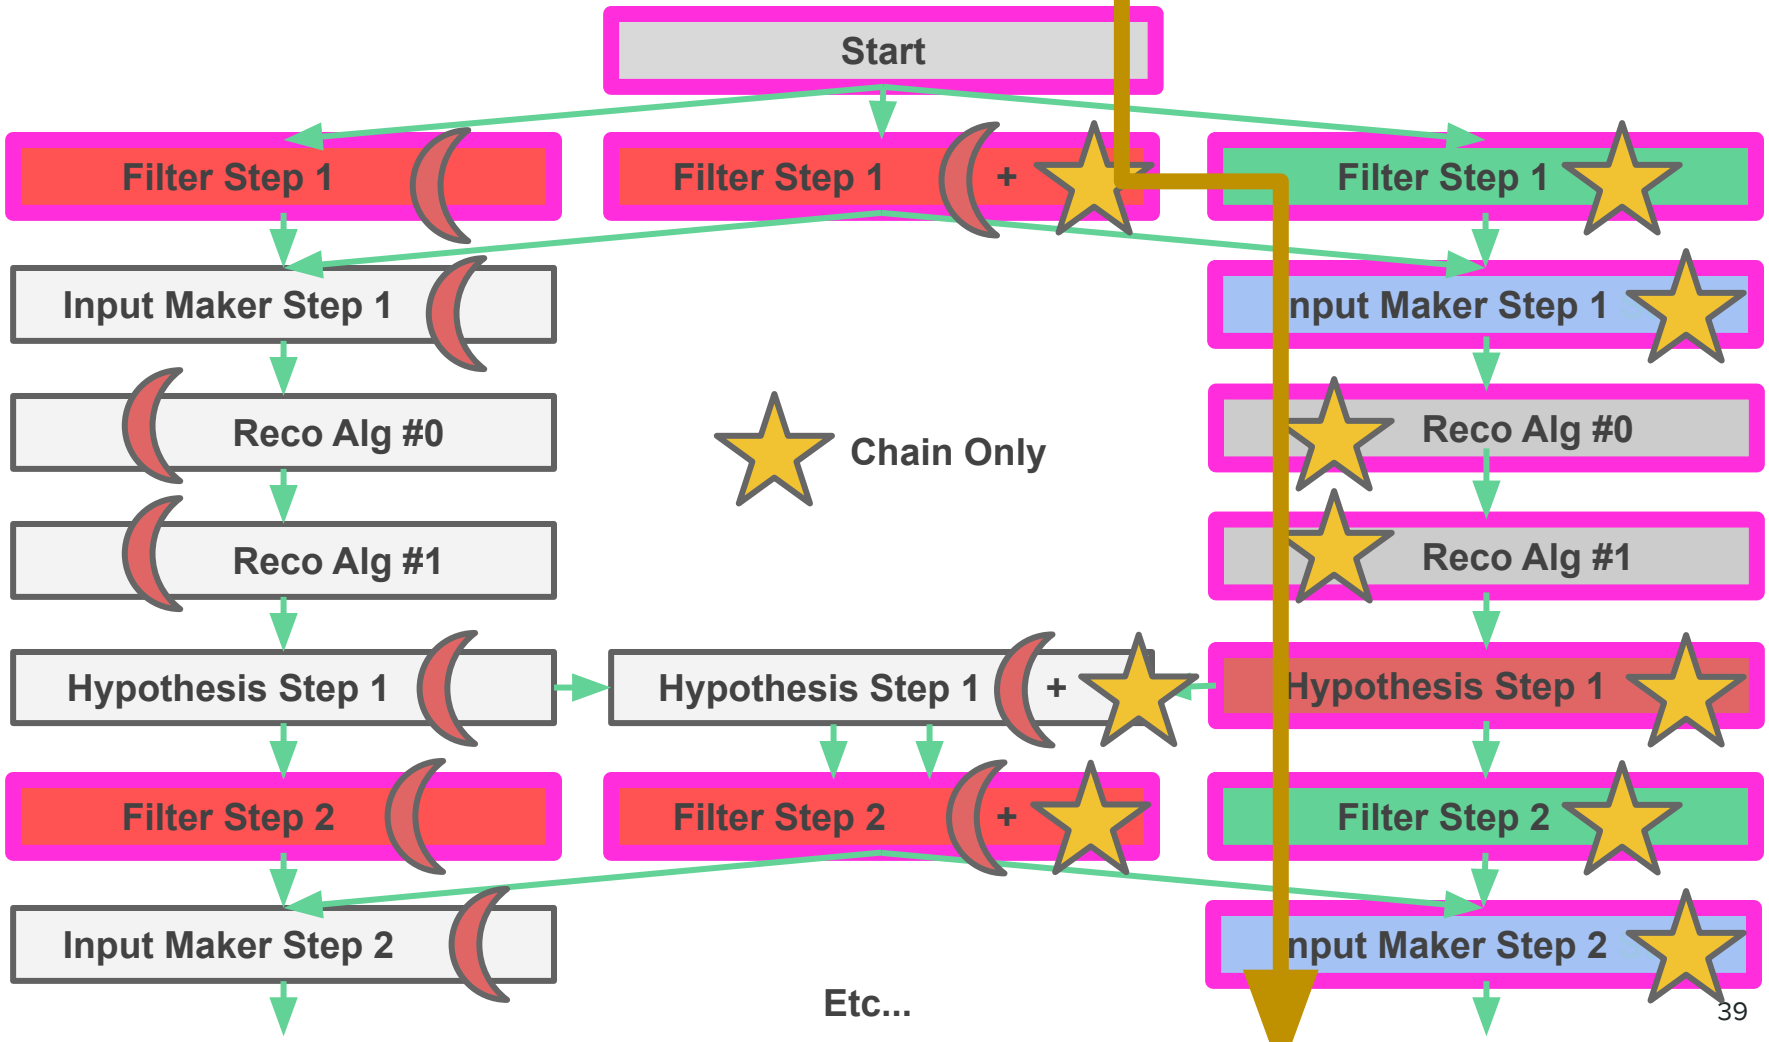
Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
nput Maker Step 1
Reco Alg #0
Reco Alg #0
Chain Only
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
nput Maker Step 2
Etc...

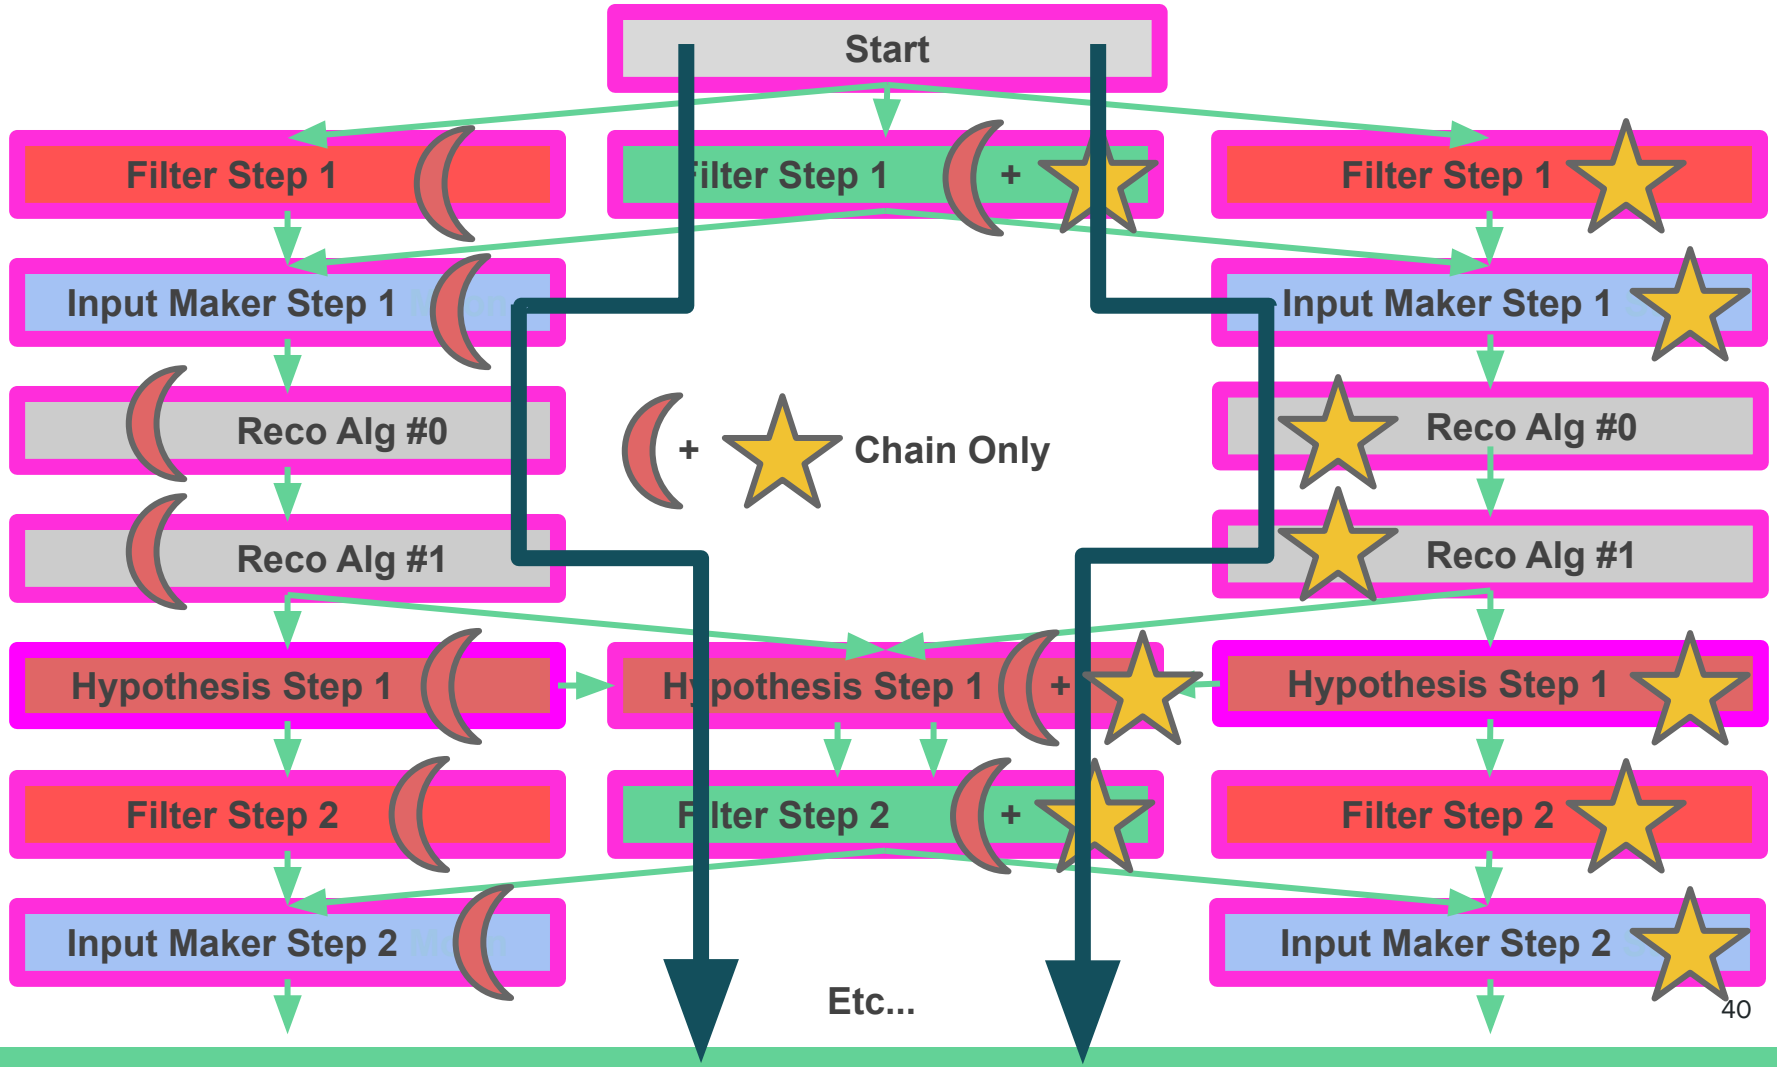
Start
Filter Step 1
Filter Step 1
Filter Step 1
Input Maker Step 1
Input Maker Step 1
Reco Alg #0
Reco Alg #0
+
Chain Only
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
Hypothesis Step 1
Filter Step 2
Filter Step 2
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

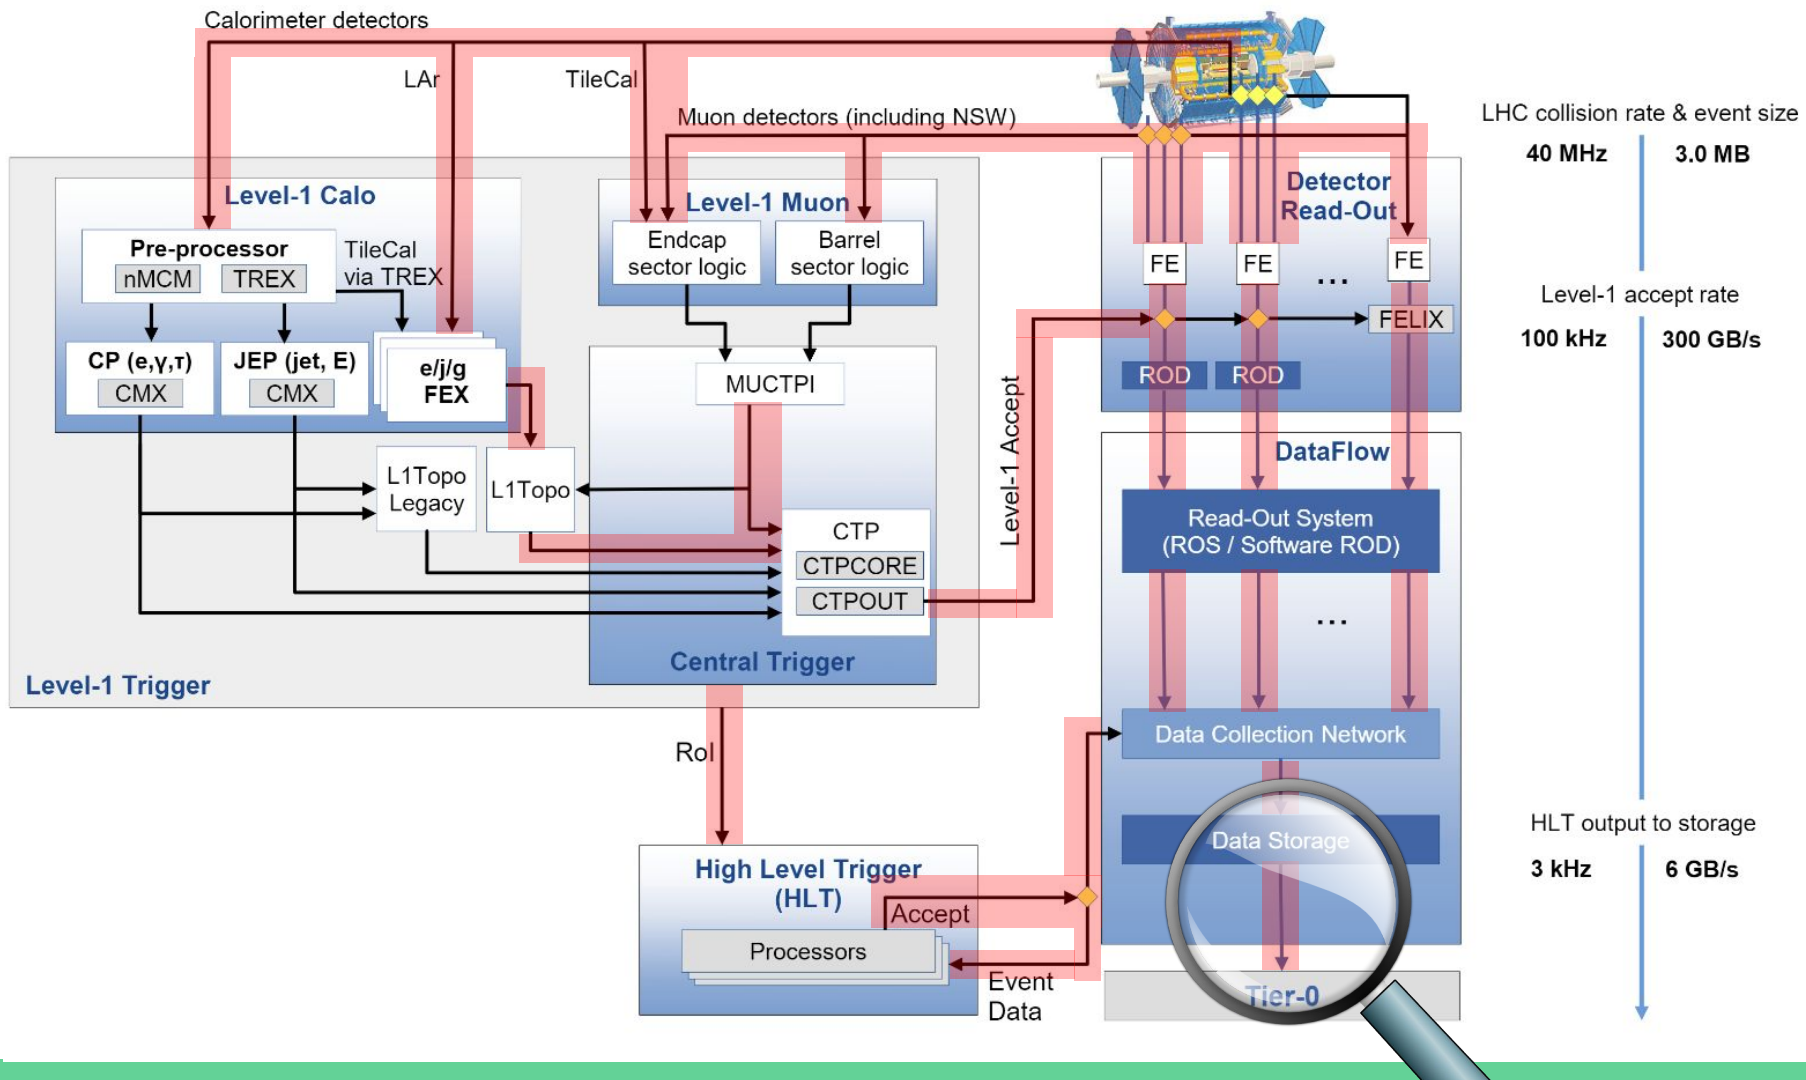
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 Accept
Detector Read-Out
FE
FE
...
FE
FELIX
ROD
ROD
DataFlow
Read-Out System (ROS / Software ROD)
...
Data Collection Network
Data Storage
Tier-0
RoI
High Level Trigger (HLT)
Processors
Accept
Event Data
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 accept rate
100 kHz
300 GB/s
HLT output to storage
3 kHz
6 GB/s

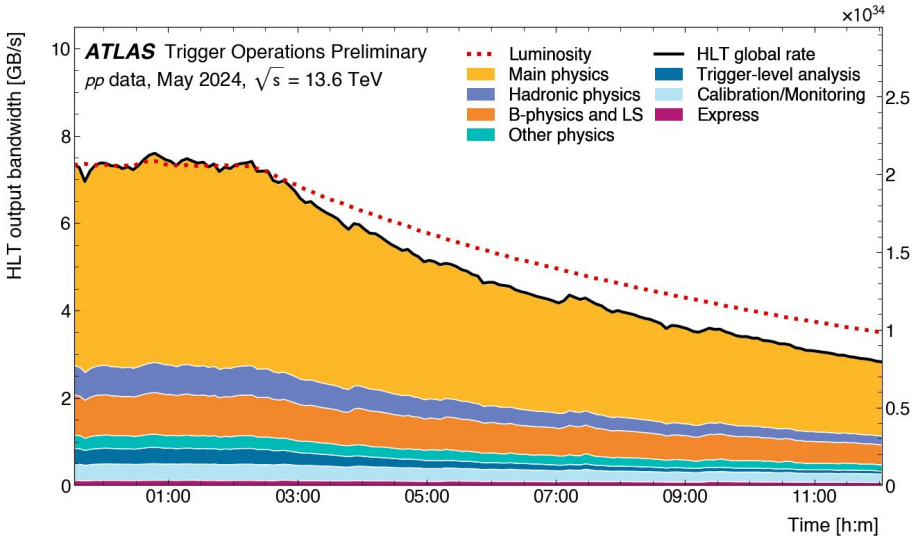

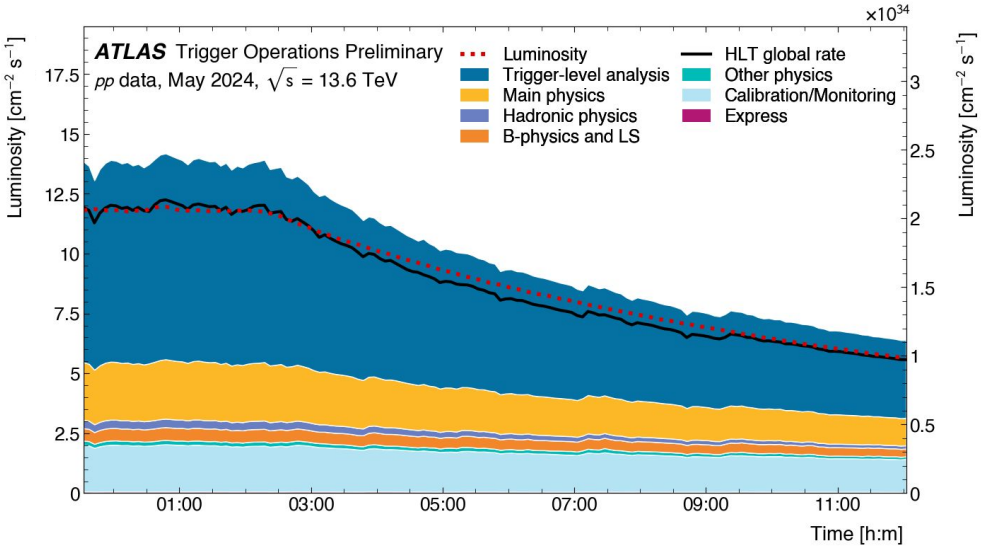

- ATLAS' HLT Trigger Menu records around **3 kHz full-events to disk** (x32 rejection).
- The **largest bandwidth** is on the **"Main"** stream of physics events, but the largest rate is for the **"Trigger Level Analysis"** stream.
    - *Only* save the High Level Trigger reconstructed physics objects, and none of the raw detector data.

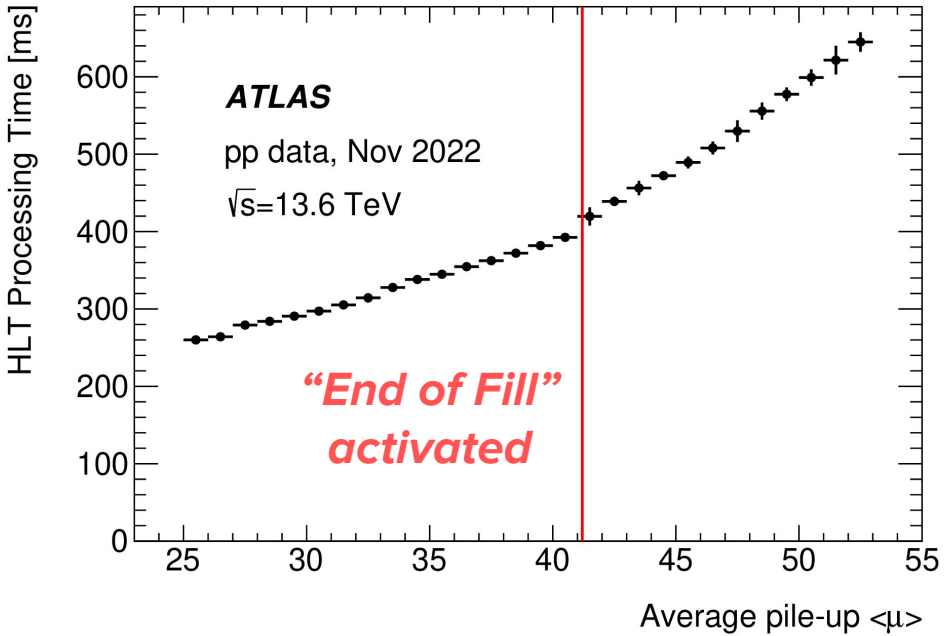

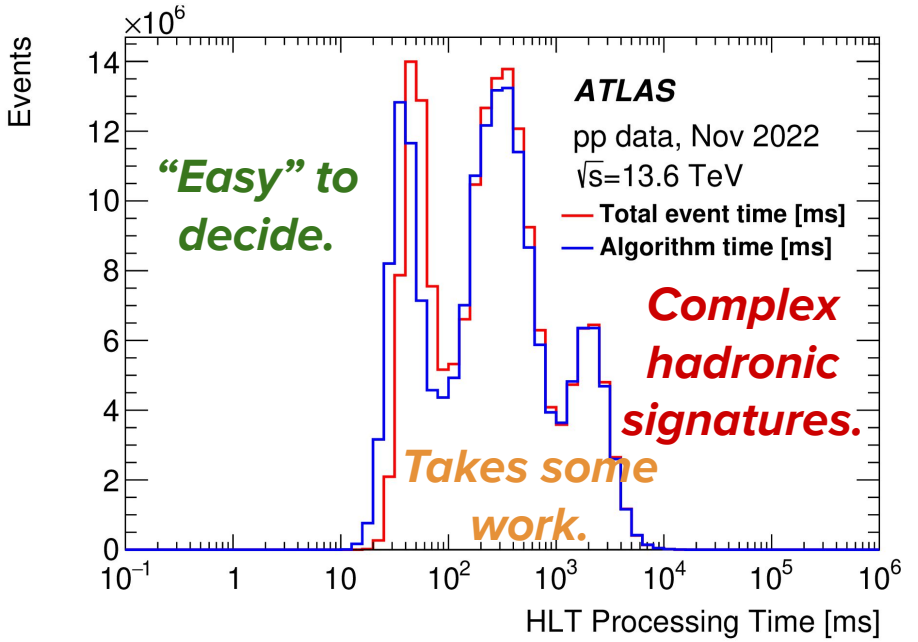

- Dual-CPU machines, **128 GB RAM**, **2x AMD EPYC 7302**. Total of **32 physical** and **64 physical+virtual** (hyper-threading) threads per machine.
- Total of around **120,000 threads** in the farm.

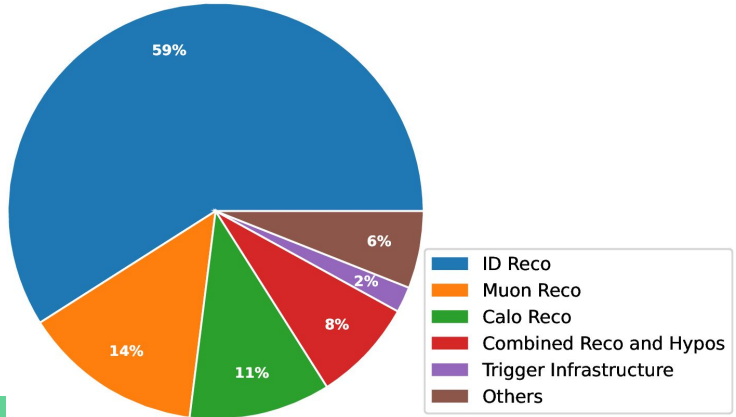

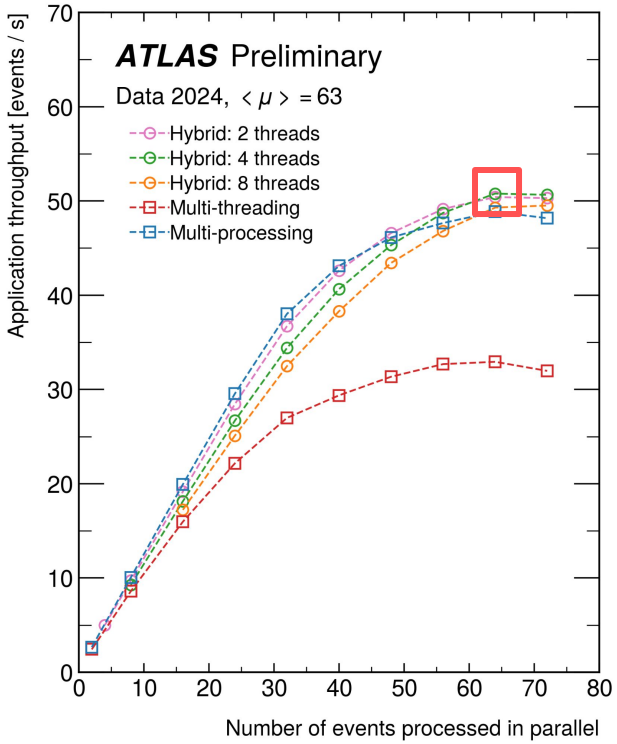

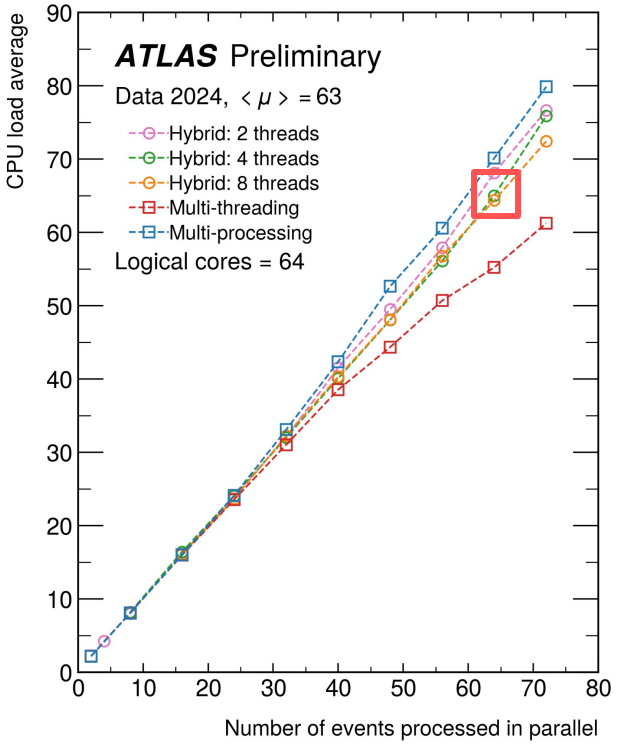

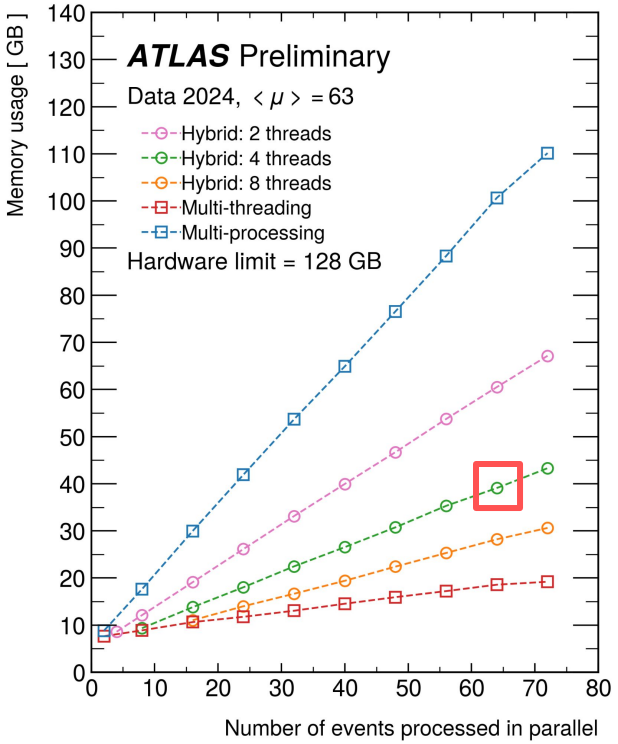

- Multi-threaded execution is still relatively new to ATLAS (2022+).
- With our MT-system we can fine tune the number of ***forks*** (multi-process) and the number of **threads** *(= concurrent events)* (multi-threading) used by each process.
- We continue to work to reduce thread contention & maximise MT operation.

Triggers In Other Detectors

Other Systems: LHCb

- From 2022, LHCb has essentially **done away with its hardware trigger**.
- Instead, all detector data are read out directly into High Level Trigger (software), which features GPU based acceleration.
- Possible due to (relatively) small event size (approx. 10% of CMS/ATLAS), but involves **very complex reconstruction at high rates in HLT**.

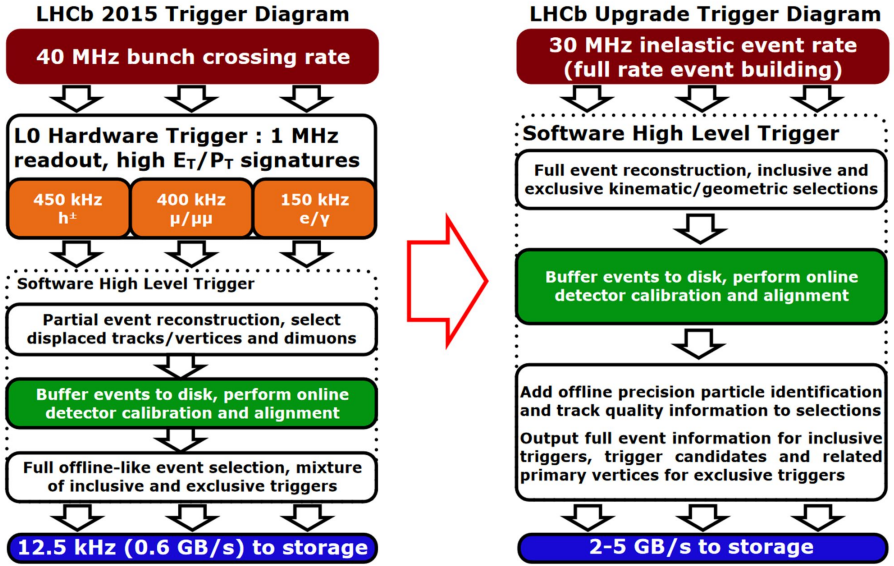

Other Systems: ALICE

- From 2022, moved to integrated online-offline ‘O2’ system with **continuous readout**.
- Introduction of new Fast Interaction Trigger (FIT) minimum
  bias trigger with high precision timing and luminometry.
- Making use of GPU-accelerated reconstruction.

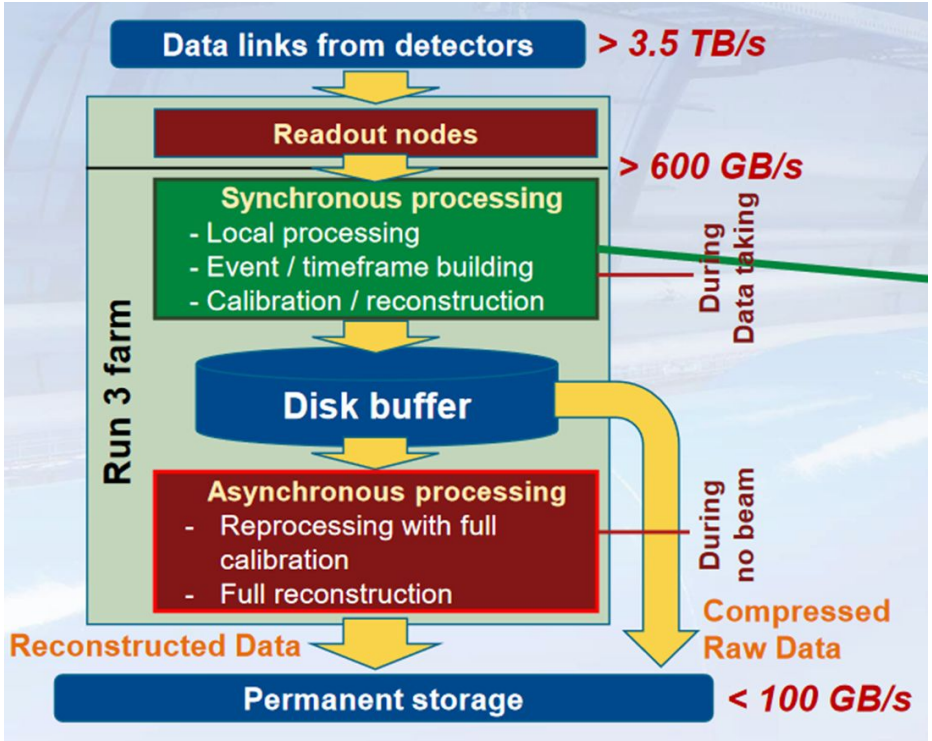

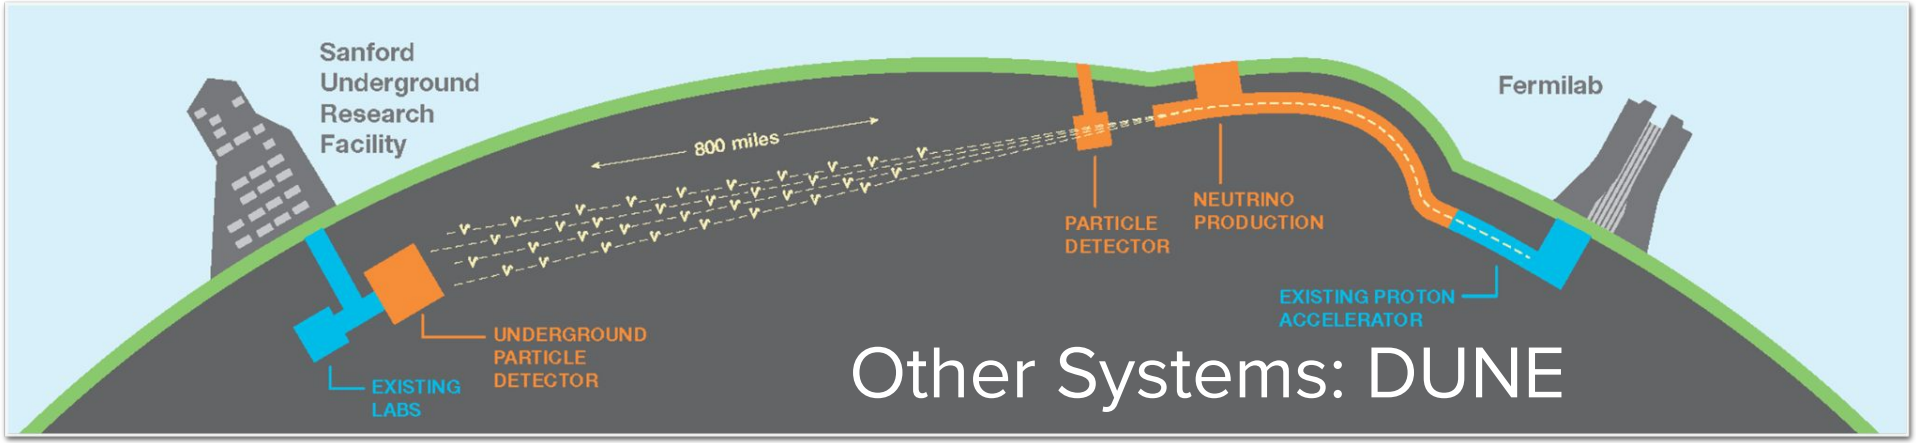

Other Systems: DUNE

- Single trigger level with two modes of operation:
  - **Interaction triggers** (interesting localized activity somewhere in the detector)
  - **Beam triggers**, cosmic rays and photon detection.
- Supernova Neutrino Burst triggers
  - All data are stored for **100 sec window including O(10 s) before the trigger signal**.
  - Interfaces to the **Supernova Early Warning System**

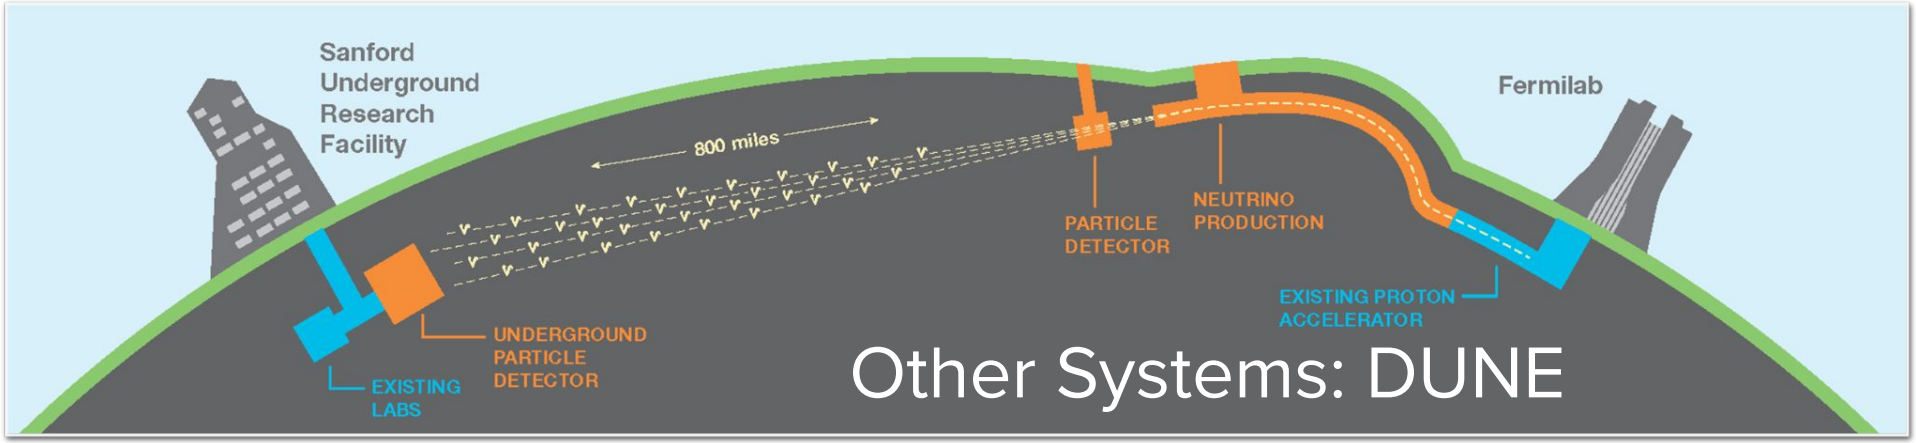

Other Systems: DUNE

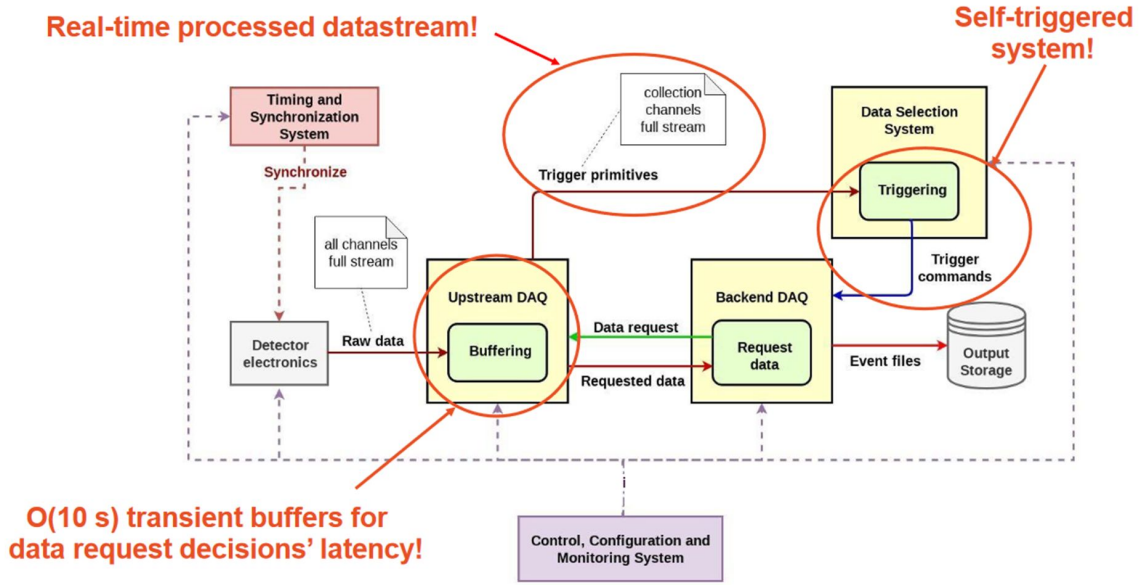

Future at the High Luminosity LHC

To The Future! (Phase-II)

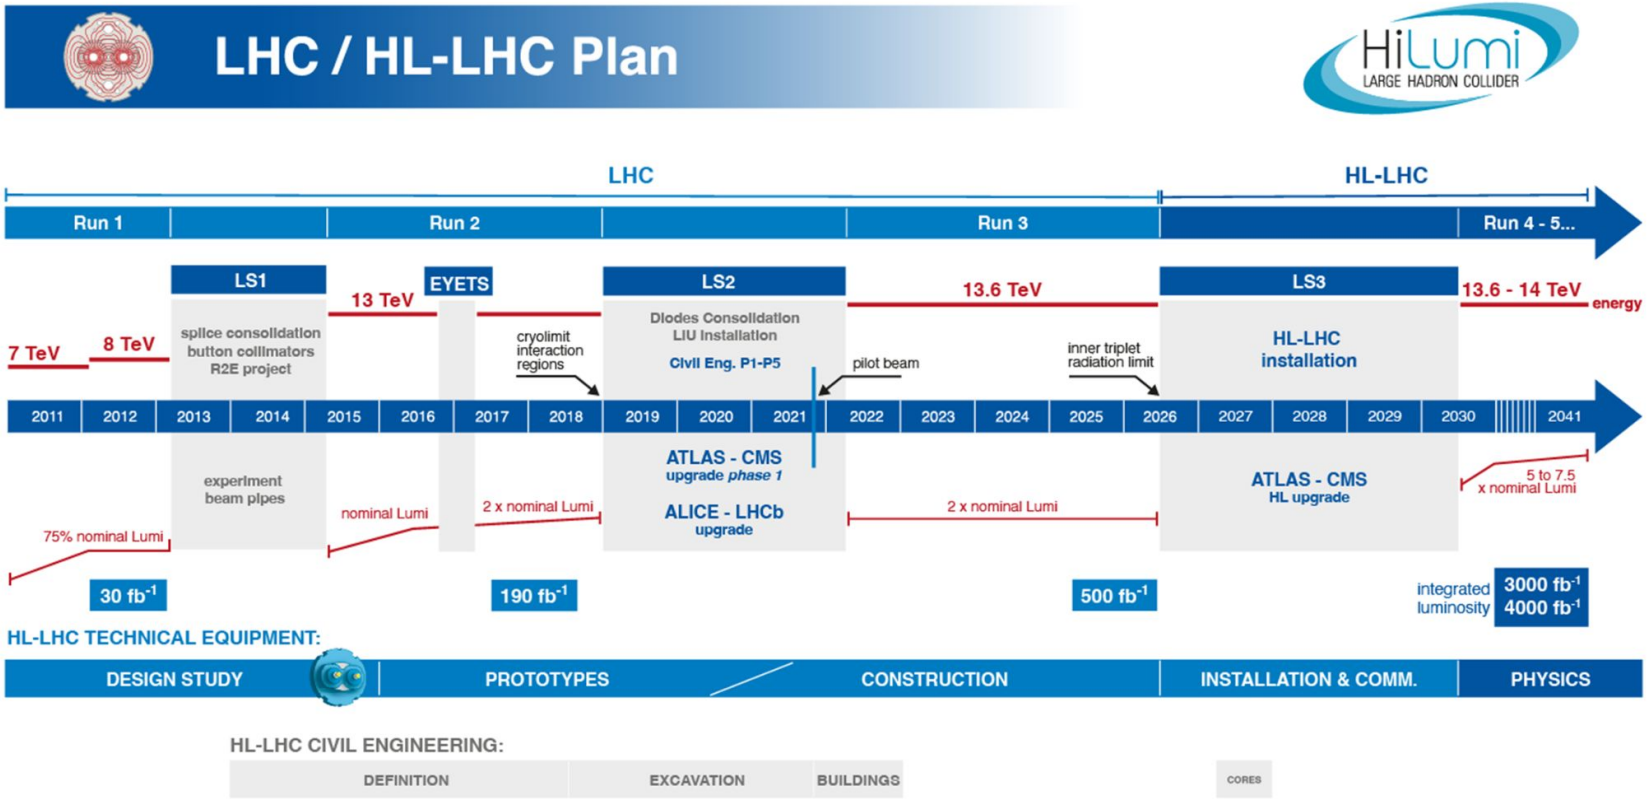

Phase-II in ATLAS

- **x10** and **x3.3** increases in **L1** and **HLT** output bandwidth.
- New **time multiplexed Global FPGA** trigger at L1, with access to calorimeter **cells**.
- HLT processing nodes have access to **all event data fragments**.
- HLT processing nodes will be augmented with **GPUs** for **track reconstruction** (and more).
- Greater use of **machine learning** in hardware & software levels.

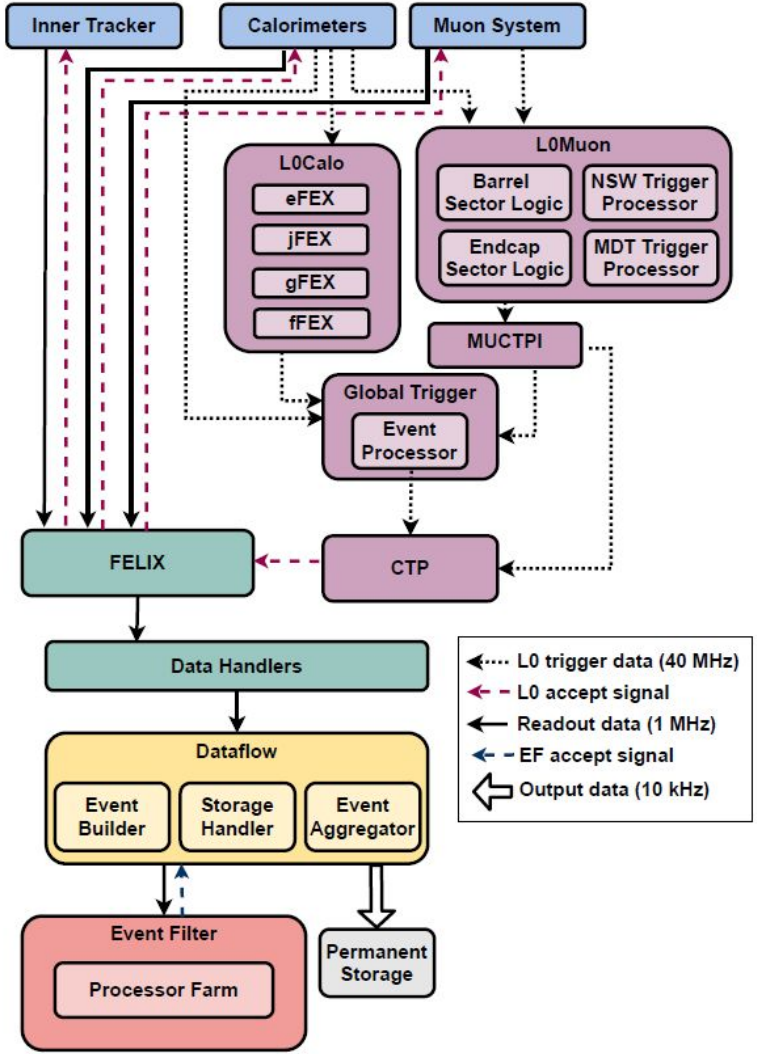

Phase-II in CMS

- L1 latency increases **4 → 12.5 μs**
- Event size increase to **7.4 MB**
- L1 trigger rate increases **100 → 750 kHz**
- Rate to storage **~7.5 kHz**
- Introduce **tracking trigger in hardware** at L1.
- Targeting **50-80% GPU offload**.
- Increased used of **Trigger Level Analysis**.
- Calibrated, **offline-like quality objects** in the trigger.

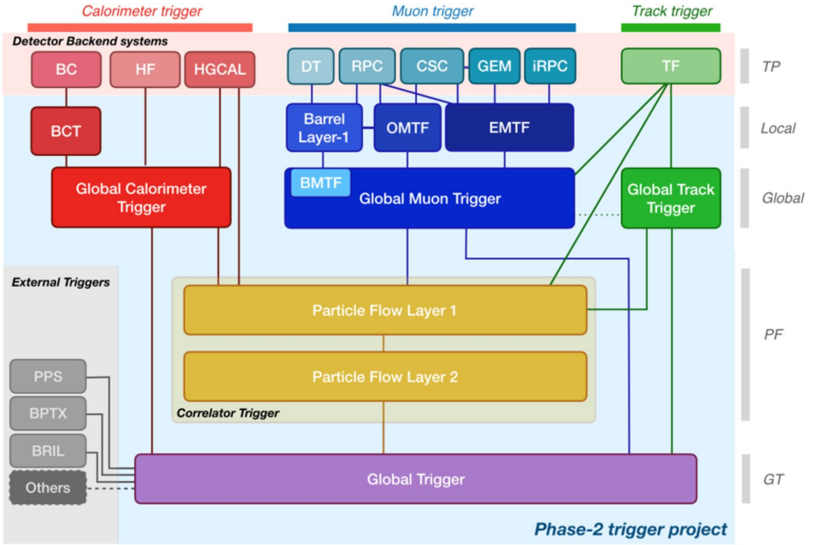

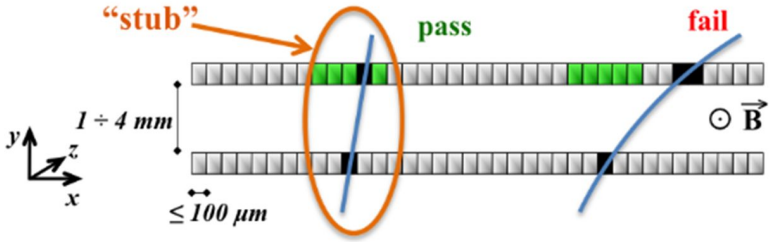

Summary

- *“A system that uses simple criteria to rapidly decide which events in a particle detector to keep when only a small fraction of the total can be recorded.”*
- Trigger strategy is a **trade-off** between **physics requirements** and **affordable systems** and technologies.
- Main trigger requirements:
  - **High efficiency** but with control of rates.
  - Knowledge of trigger selection on signal and background events.
  - Flexibility and redundancy.
- Most triggers are underpinned by some fundamental concepts, but there is a lot of innovation happening now right across the field as technology continues to evolve.