Triggering in ATLAS for Run 3

31st July 2024

Tim Martin (STFC)

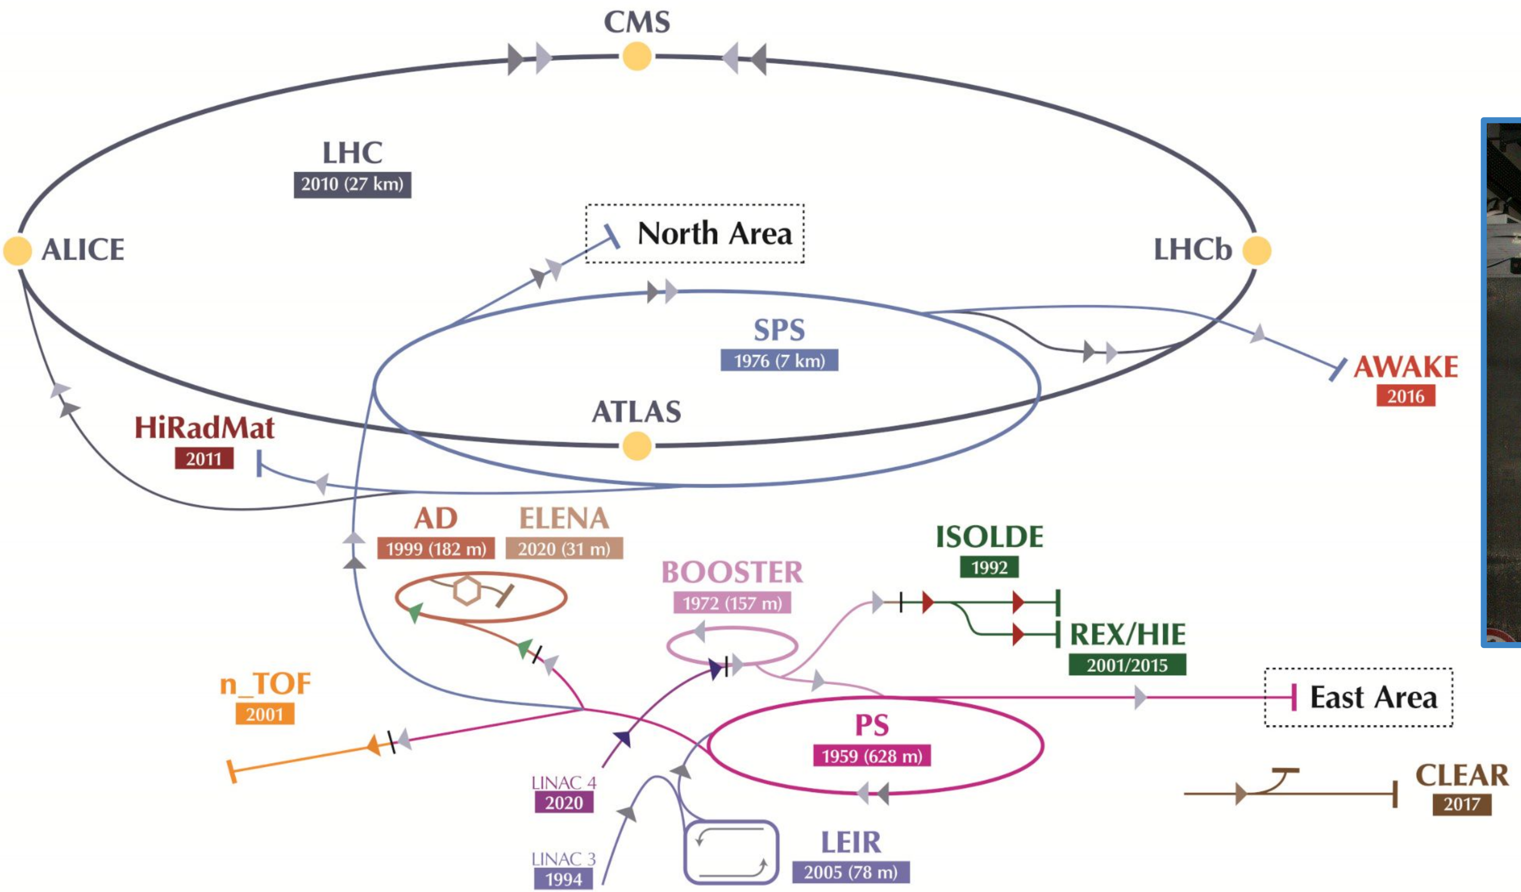
CMS
LHC
2010 (27 km)
ALICE
North Area
LHCb
SPS
1976 (7 km)
AWAKE
2016
ATLAS
HiRadMat
2011
AD
1999 (182 m)
ELENA
2020 (31 m)
ISOLDE
1992
BOOSTER
1972 (157 m)
REX/HIE
2001/2015
n_TOF
2001
East Area
PS
1959 (628 m)
LINAC 4
2020
CLEAR
2017
LEIR
2005 (78 m)
LINAC 3
1994


PREINJECTEUR
LINAC 2

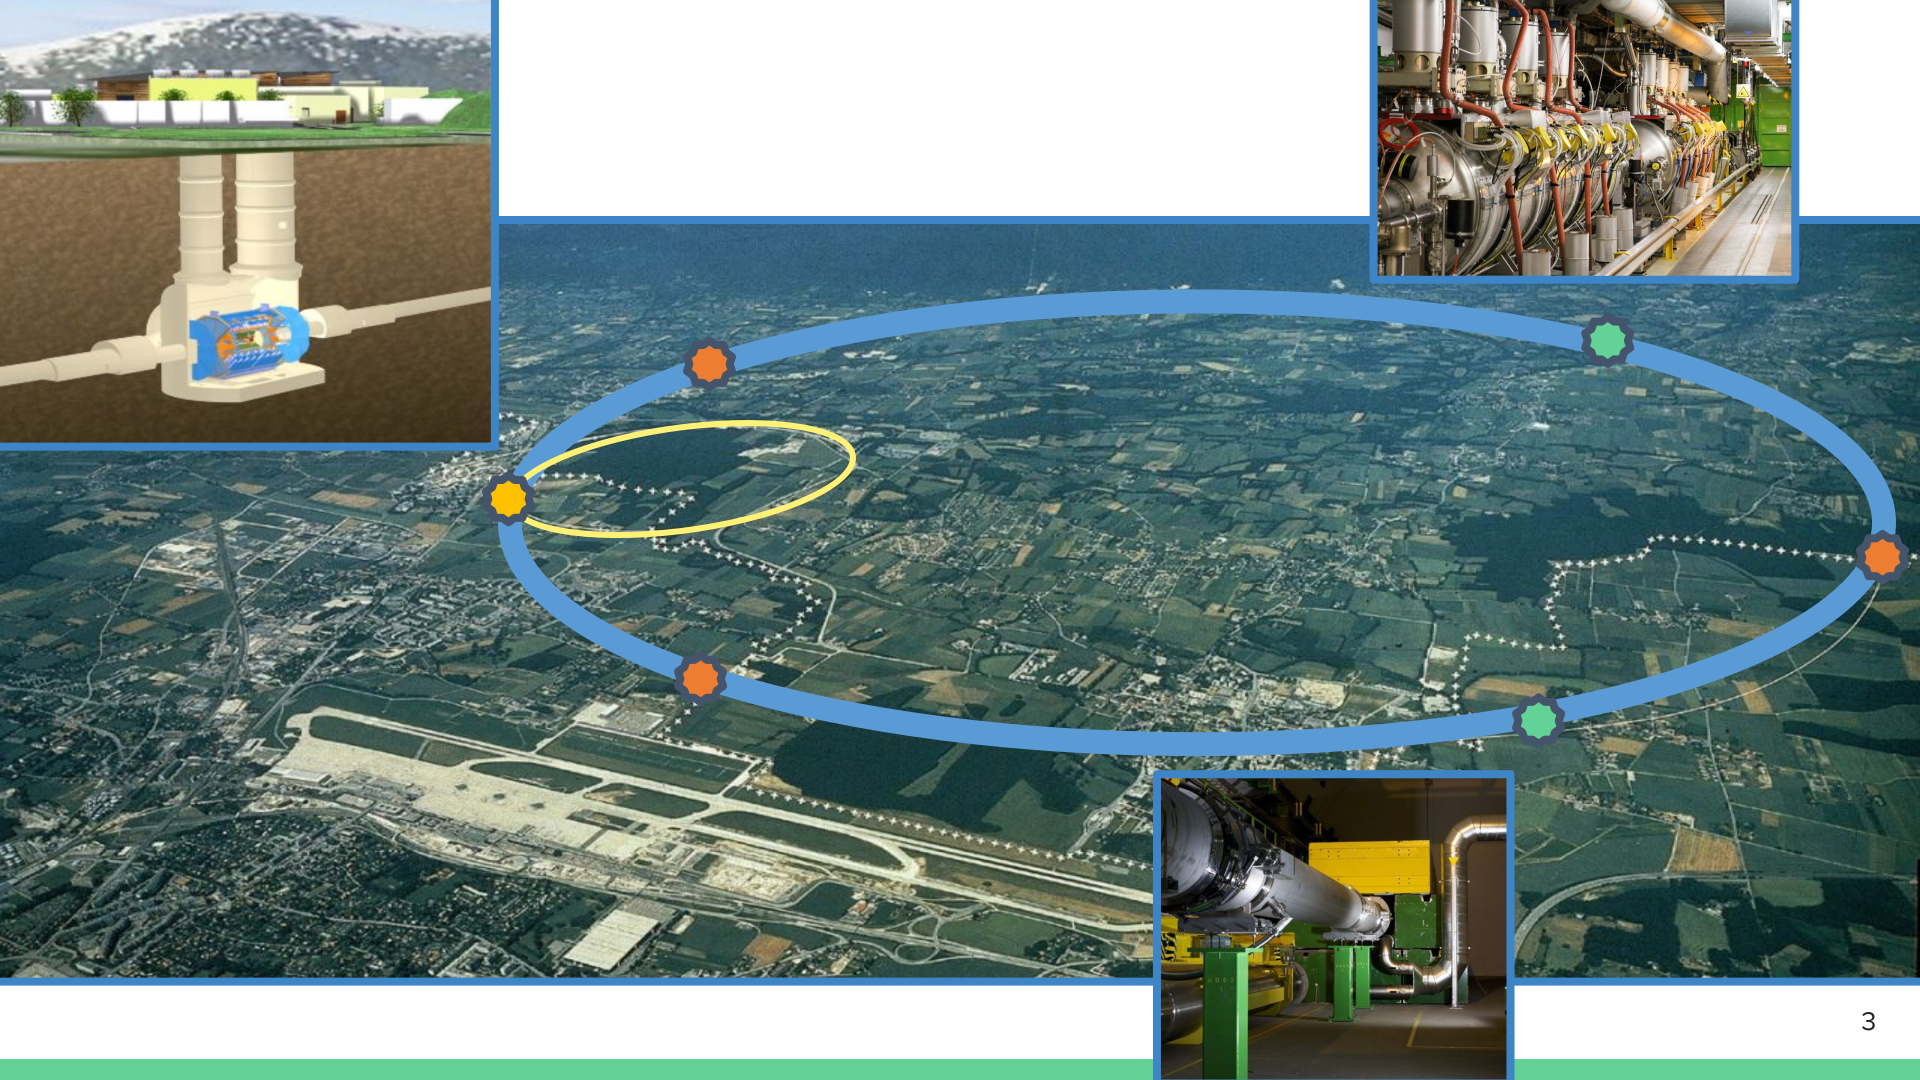

- LHC Radio Frequency is **400.8 MHz**
- LHC Circumference is **26,659 m**
- Ring divides into **35,640** buckets
  - Buckets separated by around **0.75 m**
- Proton bunches from the PS separated by **25 ns**
  - Buckets which may have protons are separated by around **7.5 m**
- The ring is never *entirely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

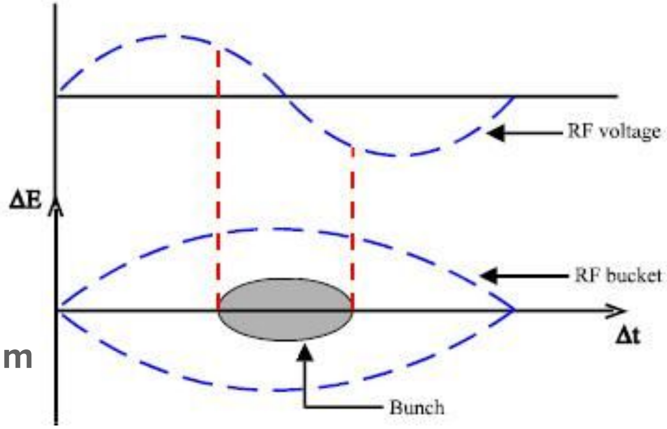

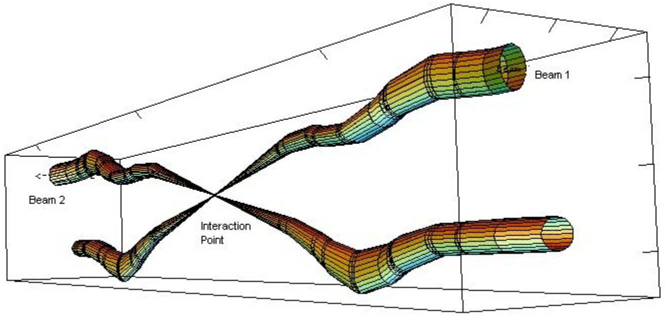

Relative beam sizes around IP1 (Atlas) in collision

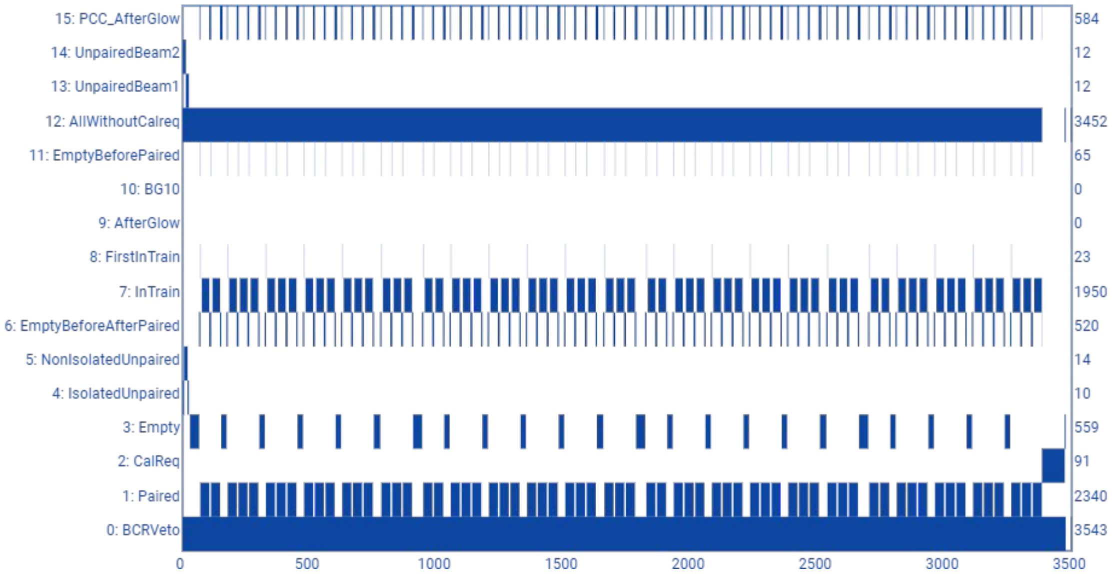

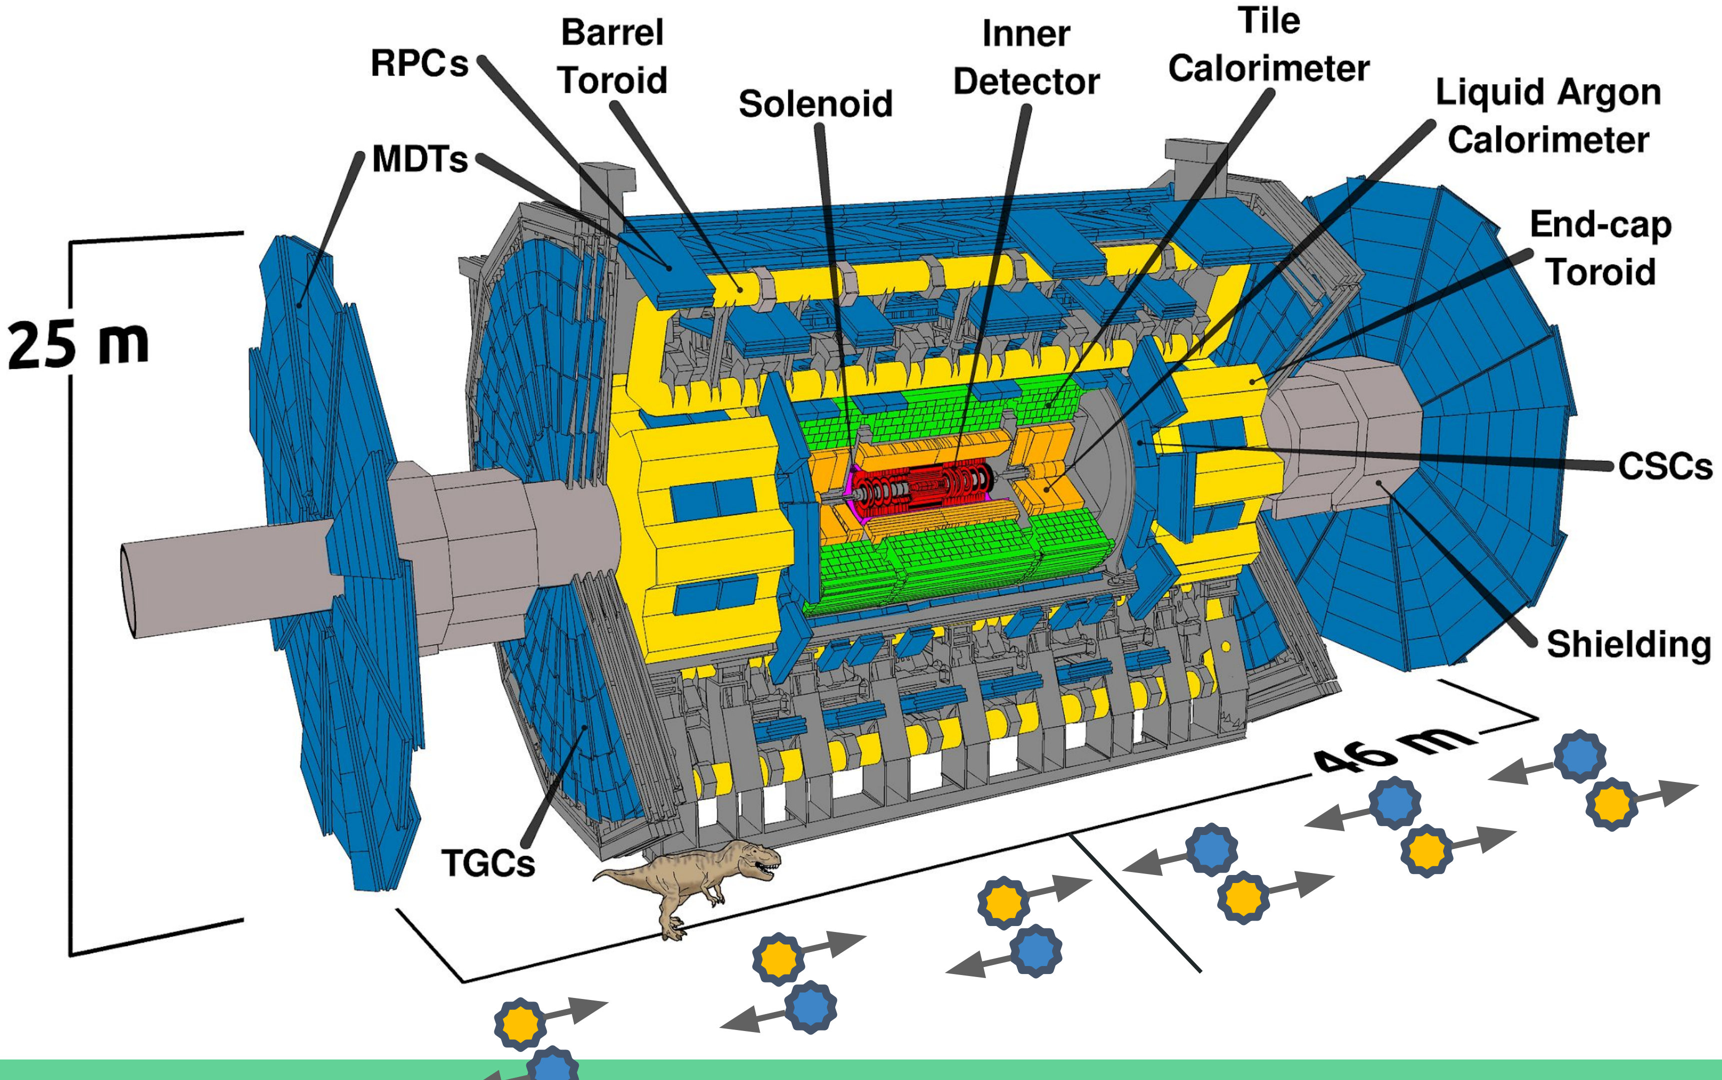
Barrel Toroid
RPCs
Solenoid
Inner Detector
Tile Calorimeter
Liquid Argon Calorimeter
MDTs
End-cap Toroid
25 m
CSCs
Shielding
46 m
TGCs

One barn, $\mathbf{10^{-24}\ m^2}$, is approximately the cross sectional area of a Uranium nucleus.

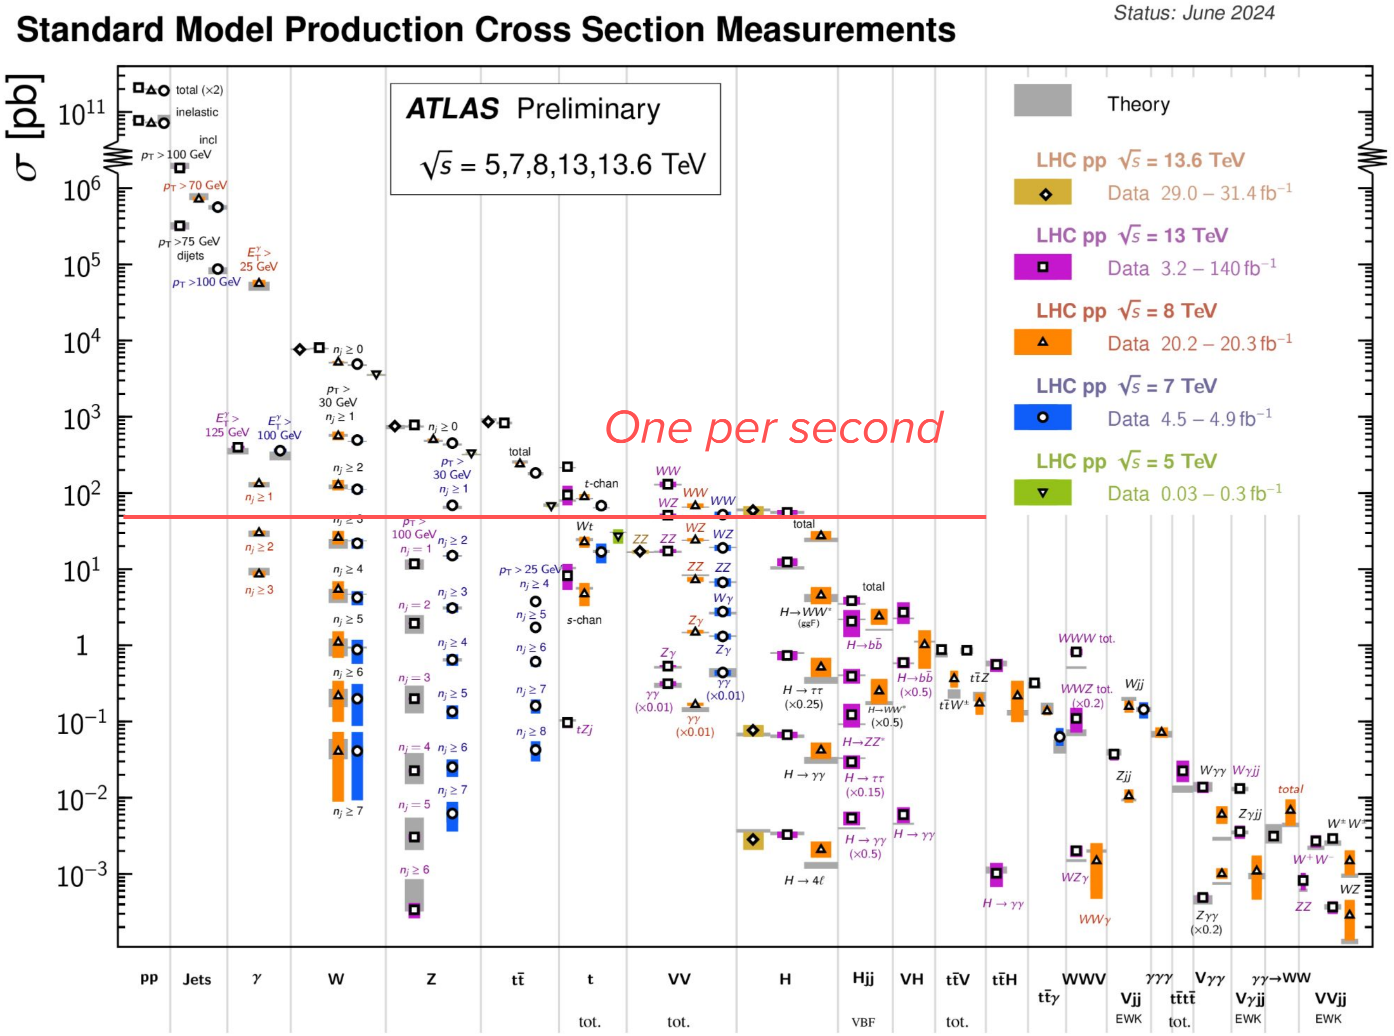

The LHC currently delivers an interaction flux (*luminosity*) of $\mathbf{2x10^{34}\ cm^{-2}\ s^{-1}}$

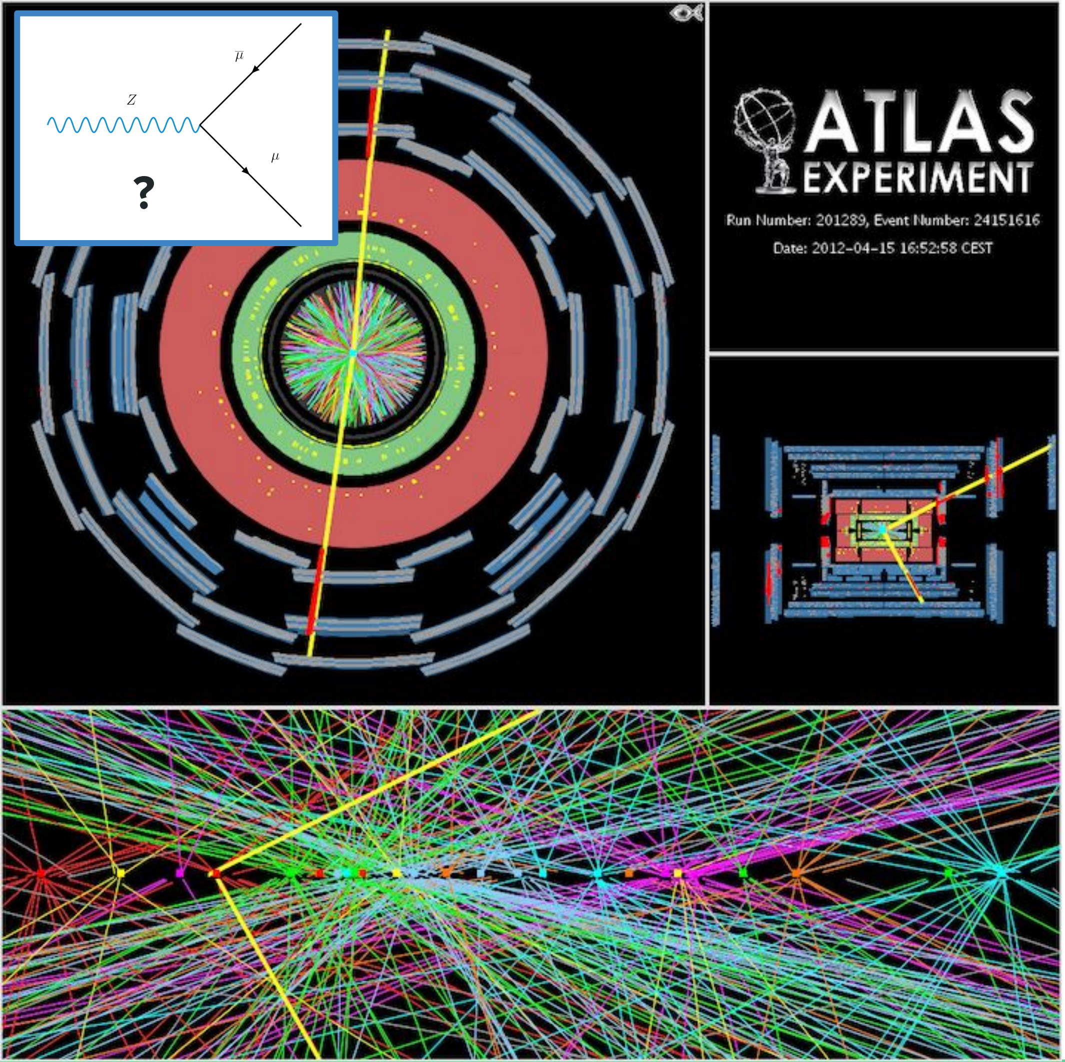

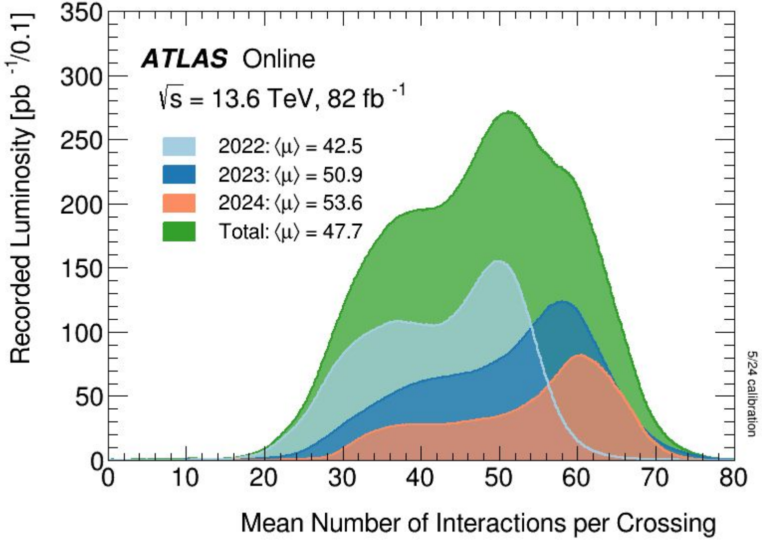

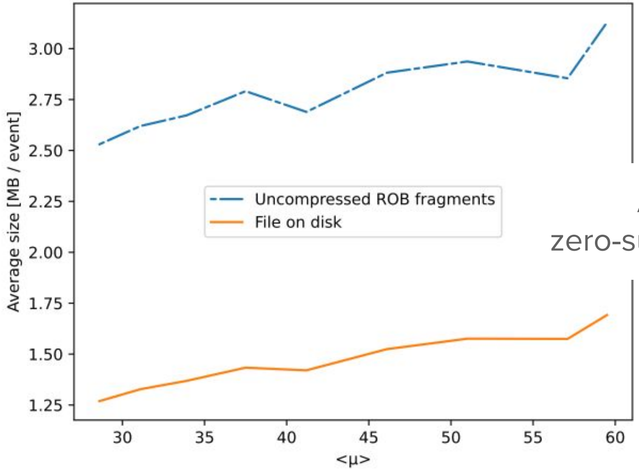

After
zero-suppression

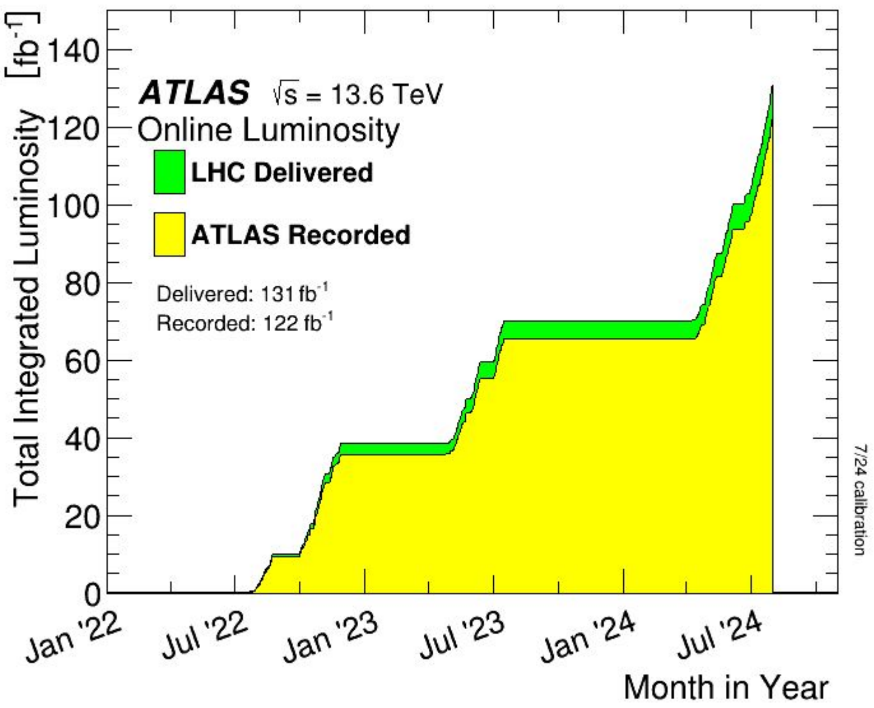
ATLAS √s = 13.6 TeV
Online Luminosity
LHC Delivered
ATLAS Recorded
Delivered: 131 fb⁻¹
Recorded: 122 fb⁻¹
Total Integrated Luminosity [fb⁻¹]
Month in Year
7/24 calibration

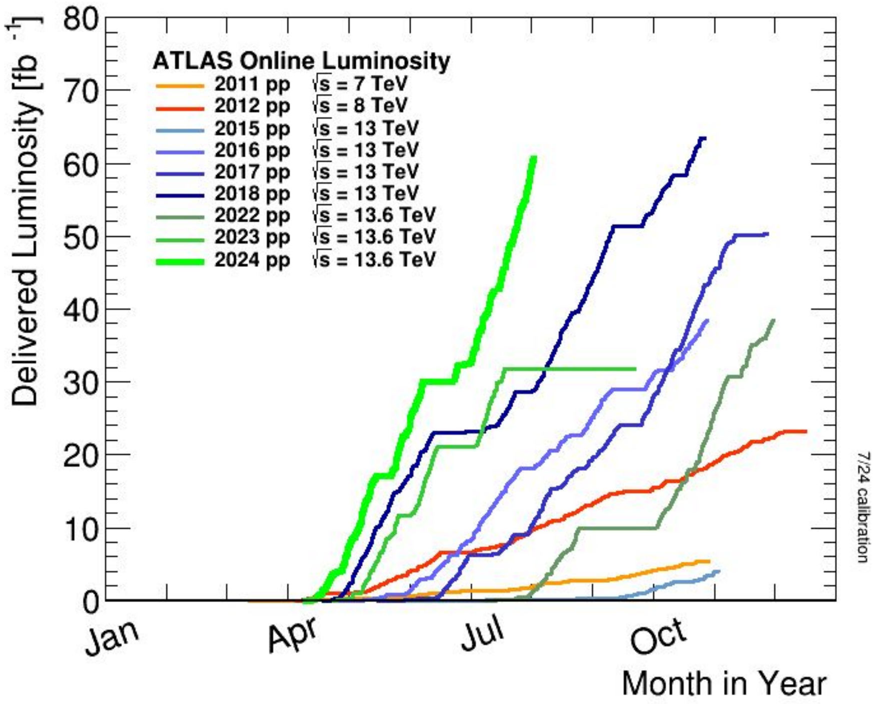
ATLAS Online Luminosity
2011 pp √s = 7 TeV
2012 pp √s = 8 TeV
2015 pp √s = 13 TeV
2016 pp √s = 13 TeV
2017 pp √s = 13 TeV
2018 pp √s = 13 TeV
2022 pp √s = 13.6 TeV
2023 pp √s = 13.6 TeV
2024 pp √s = 13.6 TeV
Delivered Luminosity [fb⁻¹]
Month in Year
7/24 calibration

Trigger in ATLAS is around 300 active personnel, and is currently focused on delivery of Run 3 (2022-2025).

Hardware-based trigger, and all upgrade activities are located in TDAQ

ATLAS Organisation

Organisation

ATLAS Management
- Spokesperson: *Andreas Hoecker (CERN)*
- Deputy Spokesperson: *Manuella Vincter (Carleton)*, *Stephane Willocq (Massachusetts)*
- Technical Coordinator: *Martin Aleksa (CERN)*
- Resources Coordinator: *David Francis (CERN)*
- Upgrade Coordinator: *Benedetto Gorini (CERN)*

Collaboration Board
- Collaboration Board Chair: *Maria Jose Costa Mezquita (Valencia)*
- Collaboration Board Deputy Chair: *Lucia Di Ciaccio (Annecy LAPP)*

Executive Board
- Spokesperson: *Andreas Hoecker (CERN)*
- Technical Coordinator: *Martin Aleksa (CERN)*

Technical Coordination
- Technical Coordinator: *Martin Aleksa (CERN)*
- Technical Deputy Coordinator: *Michel Raymond (CERN)*, *Stefan Schlenker (CERN)*

Upgrade Coordination
- Upgrade Coordinator: *Benedetto Gorini (CERN)*
- Upgrade Deputy Coordinator: *Jinlong Zhang (Argonne)*

Boards and Committees

Authorship Committee
- Authorship Committee Chair: *Helena Santos (Portugal 1-LIP)*

Diversity and Inclusion (D&I)
- Contact on Diversity and Inclusion: *Steven Goldfarb (Melbourne)*, *Flavia De Almeida Dias (Nikhef)*, *Simon Connell (Johannesburg)*, *Maria Teresa Dova (La Plata)*

Early Career Scientist Board (ECSB)
- Early Career Scientist Board (ECSB) member: *Maximilian J Swiatlowski (TRIUMF)*, *Kevin Thomas Greif (UCI)*, *Christian Appelt (Tel Aviv)*, *Laura Elaine Bruce (Harvard)*, *Name withheld*, *Petar Bokan (CERN)*, *Hannah Arnold (Stony Brook)*

Panel for Operation Task Sharing (POTS)
- POTS Chair: *Manuella Vincter (Carleton)*

Physics Office
- Physics Office Coordinator: *Gabriela Alejandra Navarro (Bogota UAN)*

Publications Committee
- Publication Committee Chair: *Aurelio Juste Rozas (Barcelona)*
- Publication Committee Deputy Chair: *Marie-Helene Genest (Grenoble LPSC)*

Speakers Committee
- Speakers Committee Chair: *Alain Bellerive (Carleton)*
- Speakers Committee Deputy Chair: *Hernan Pablo Wahlberg (La Plata)*

Speakers Committee Advisory Board (SCAB)
- SCAB Chair: *Erich Ward Varnes (Arizona)*

Other

Activities

Data Preparation
- Data Preparation Coordinator: *Heather Gray (Berkeley LBNL)*
- Data Preparation Deputy Coordinator: *Chris Malena Delitzsch (Dortmund)*

Physics
- Physics Coordinator: *Monica Dunford (Heidelberg KIP)*
- Physics Deputy Coordinator: *Fabio Cerutti (Berkeley LBNL)*

Run Coordination
- Run Coordinator: *Catrin Bernius (SLAC)*
- Run Deputy Coordinator: *Andrej Gorisek (Ljubljana)*

Software and Computing
- Computing co-Coordinator 1: *Zach Marshall (Berkeley LBNL)*
- Computing co-Coordinator 2: *David South (DESY)*

Trigger
- Trigger Coordinator: *Savanna Shaw (Manchester)*
- Trigger Deputy Coordinator: *Teng Jian Khoo (Berlin HU)*

Outreach
- Outreach Coordinator: *Hans Peter Beck (Bern)*, *Dilia Maria Portillo Quintero (TRIUMF)*

Detector Projects

Forward Detectors
- Forward Detectors System Project Leader: *Maciej Trzebinski (Krakow IFJ PAN)*
- Forward Detectors System Deputy Project Leader: *Mario Campanelli (London UC)*

High Granularity Timing Detector (HGTD)
- HGTD System Project Leader: *Joao Barreiro Guimaraes Da Costa (Beijing IHEP)*
- HGTD System Deputy Project Leader: *Stefan Guindon (CERN)*, *Sebastian Grinstein (Barcelona)*

Inner Detector (ID)
- ID System Project Leader: *Kerstin Lantzsch (Bonn)*

Inner Tracker for Phase-II (ITk)
- ITk System Project Leader: *Petra Riedler (CERN)*
- ITk System Deputy Project Leader: *Thomas Koffas (Carleton)*

Liquid-Argon Calorimeter (LAr)
- LAr System Project Leader: *Emmanuel Monnier (Marseille CPPM)*
- LAr System Deputy Project Leader: *Luis Hervas (CERN)*

Muon Spectrometer
- Muon System Deputy Project Leader: *Kostas Ntekas (UCI)*
- Muon System Project Leader: *Paolo Iengo (Napoli)*

Tile Calorimeter
- Tile System Project Leader: *Henric Wilkens (CERN)*
- Tile System Deputy Project Leader: *Bruce Mellado Garcia (Witwatersrand)*

Trigger and Data Acquisition (TDAQ)
- TDAQ System Project Leader: *Stefano Veneziano (Roma I)*
- TDMT Member: *Thorsten Wengler (CERN)*, *Stefano Veneziano (Roma I)*, *Wainer Vandelli (CERN)*

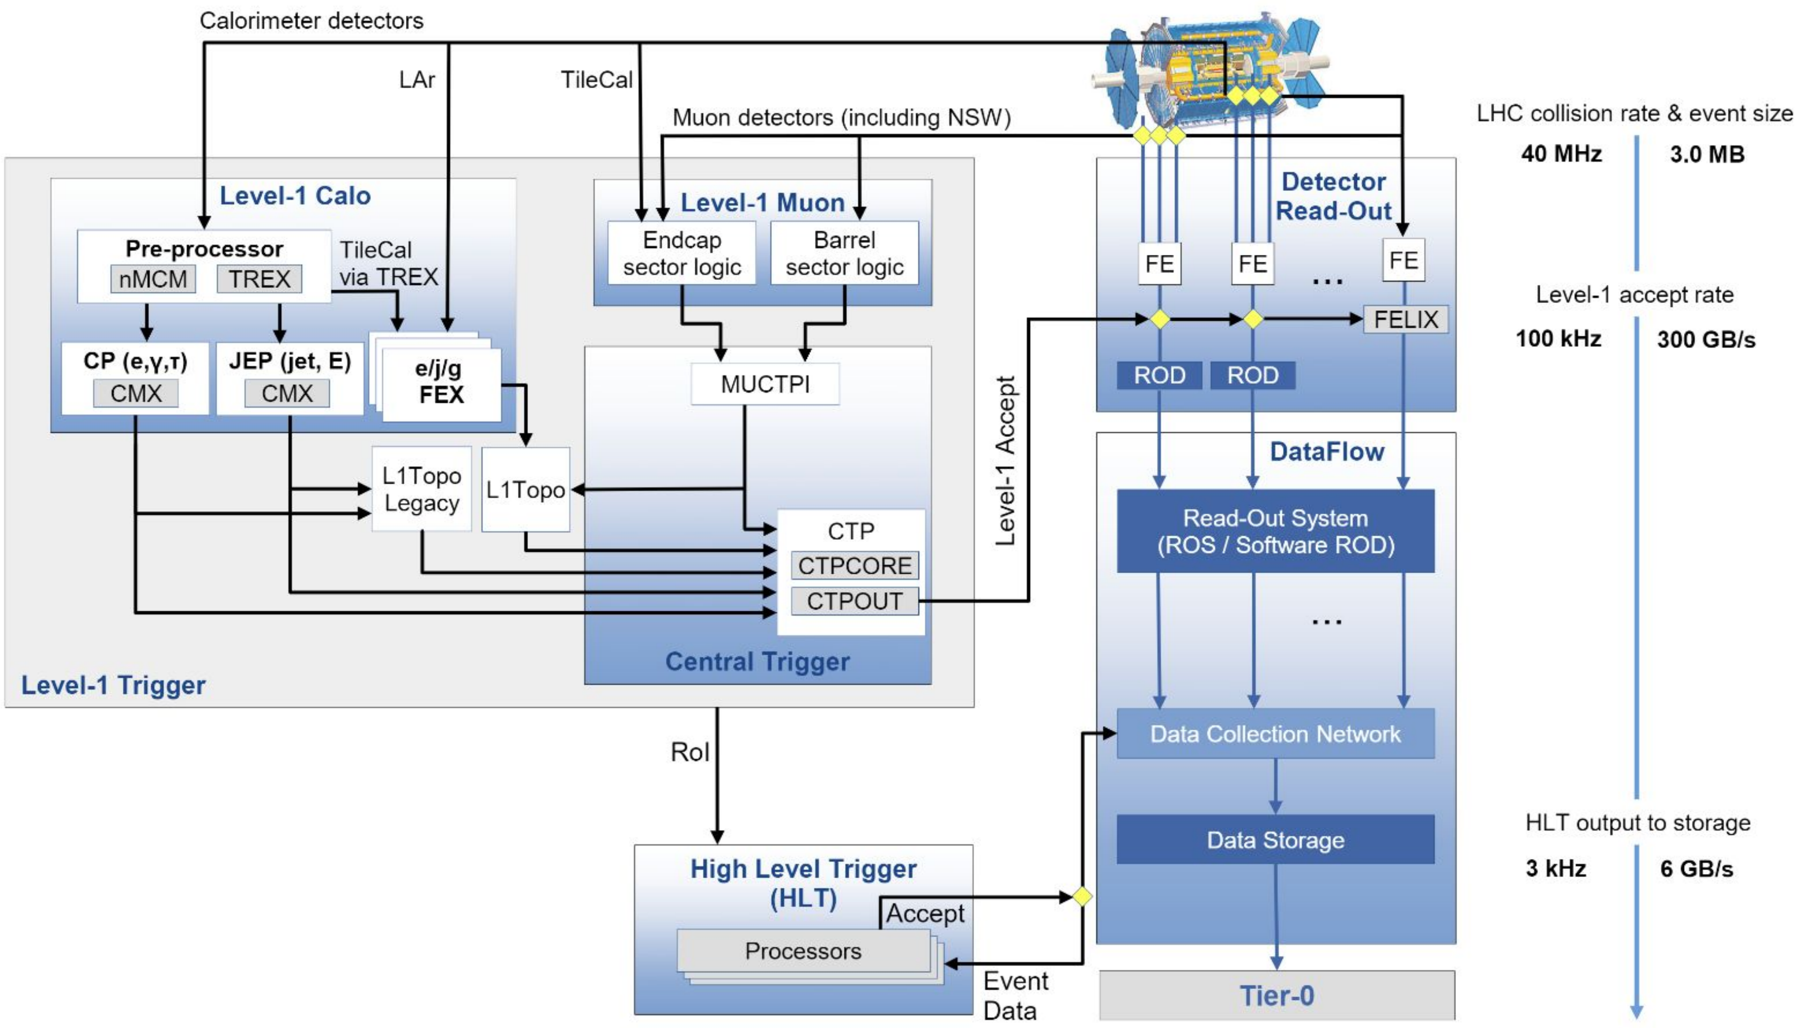
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
LHC collision rate & event size
40 MHz
3.0 MB
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
L1Topo Legacy
L1Topo
CTP
CTPCORE
CTPOUT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

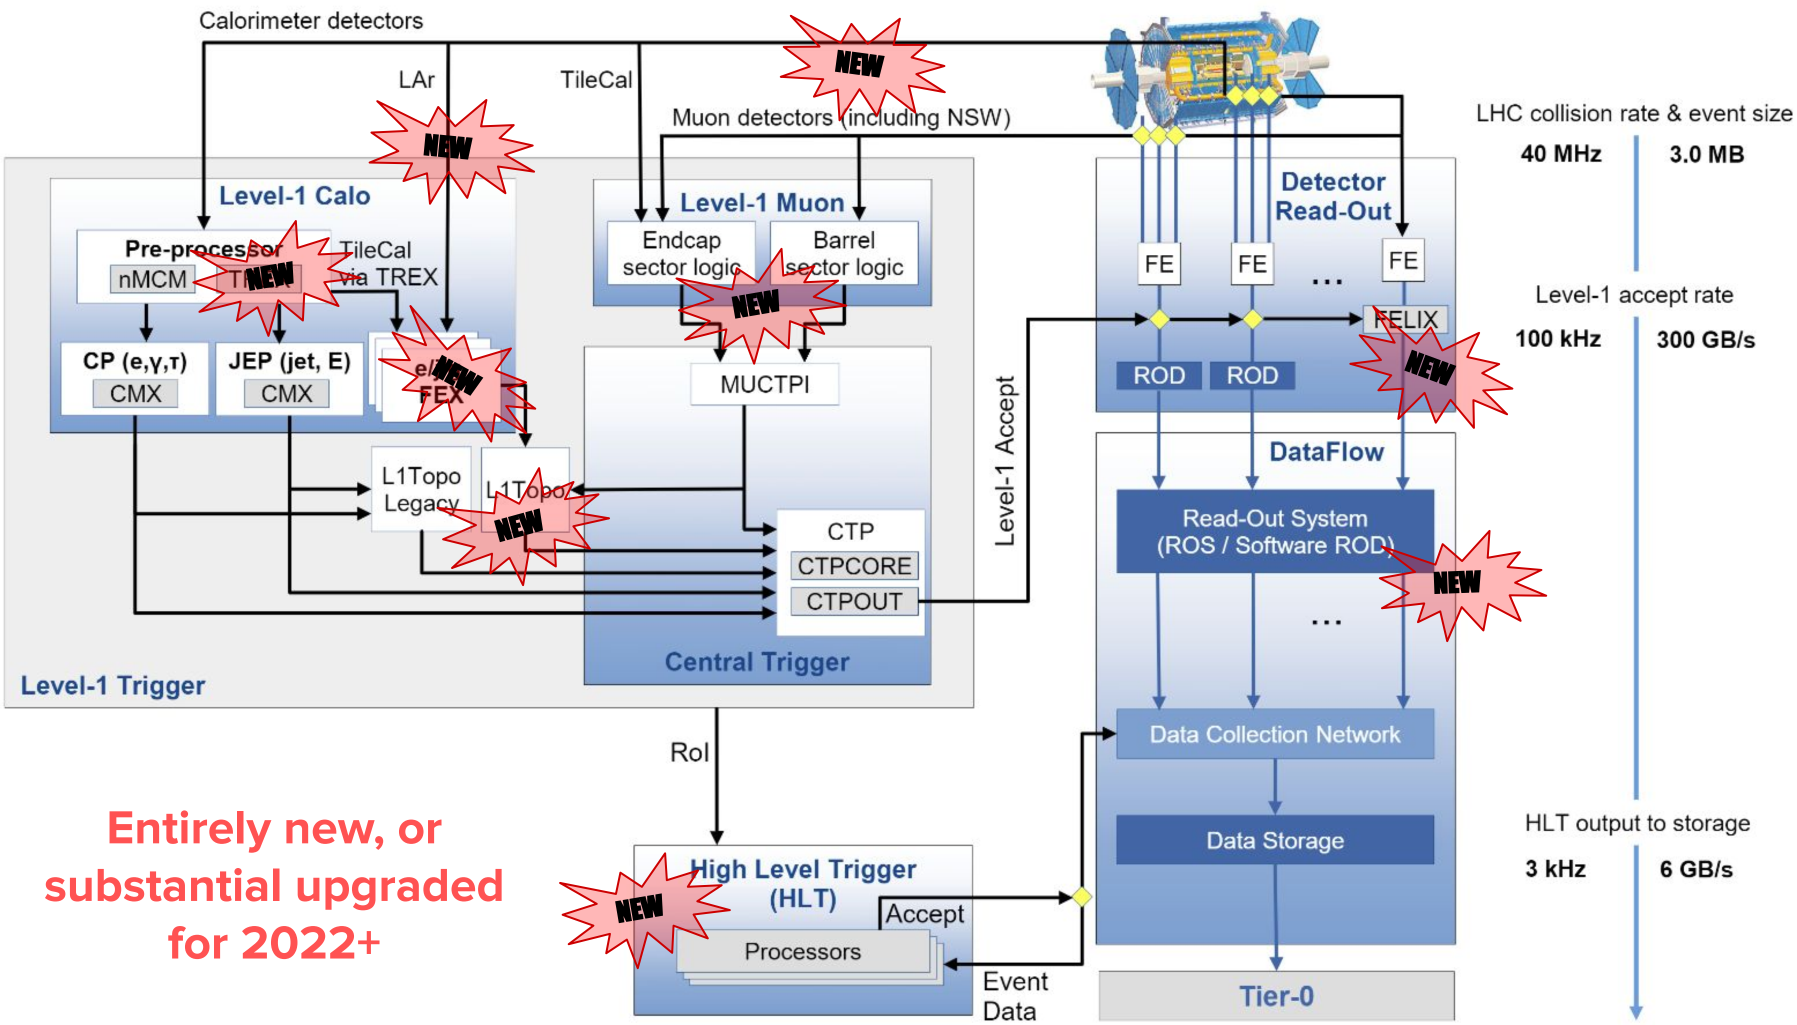

Entirely new, or substantial upgraded for 2022+

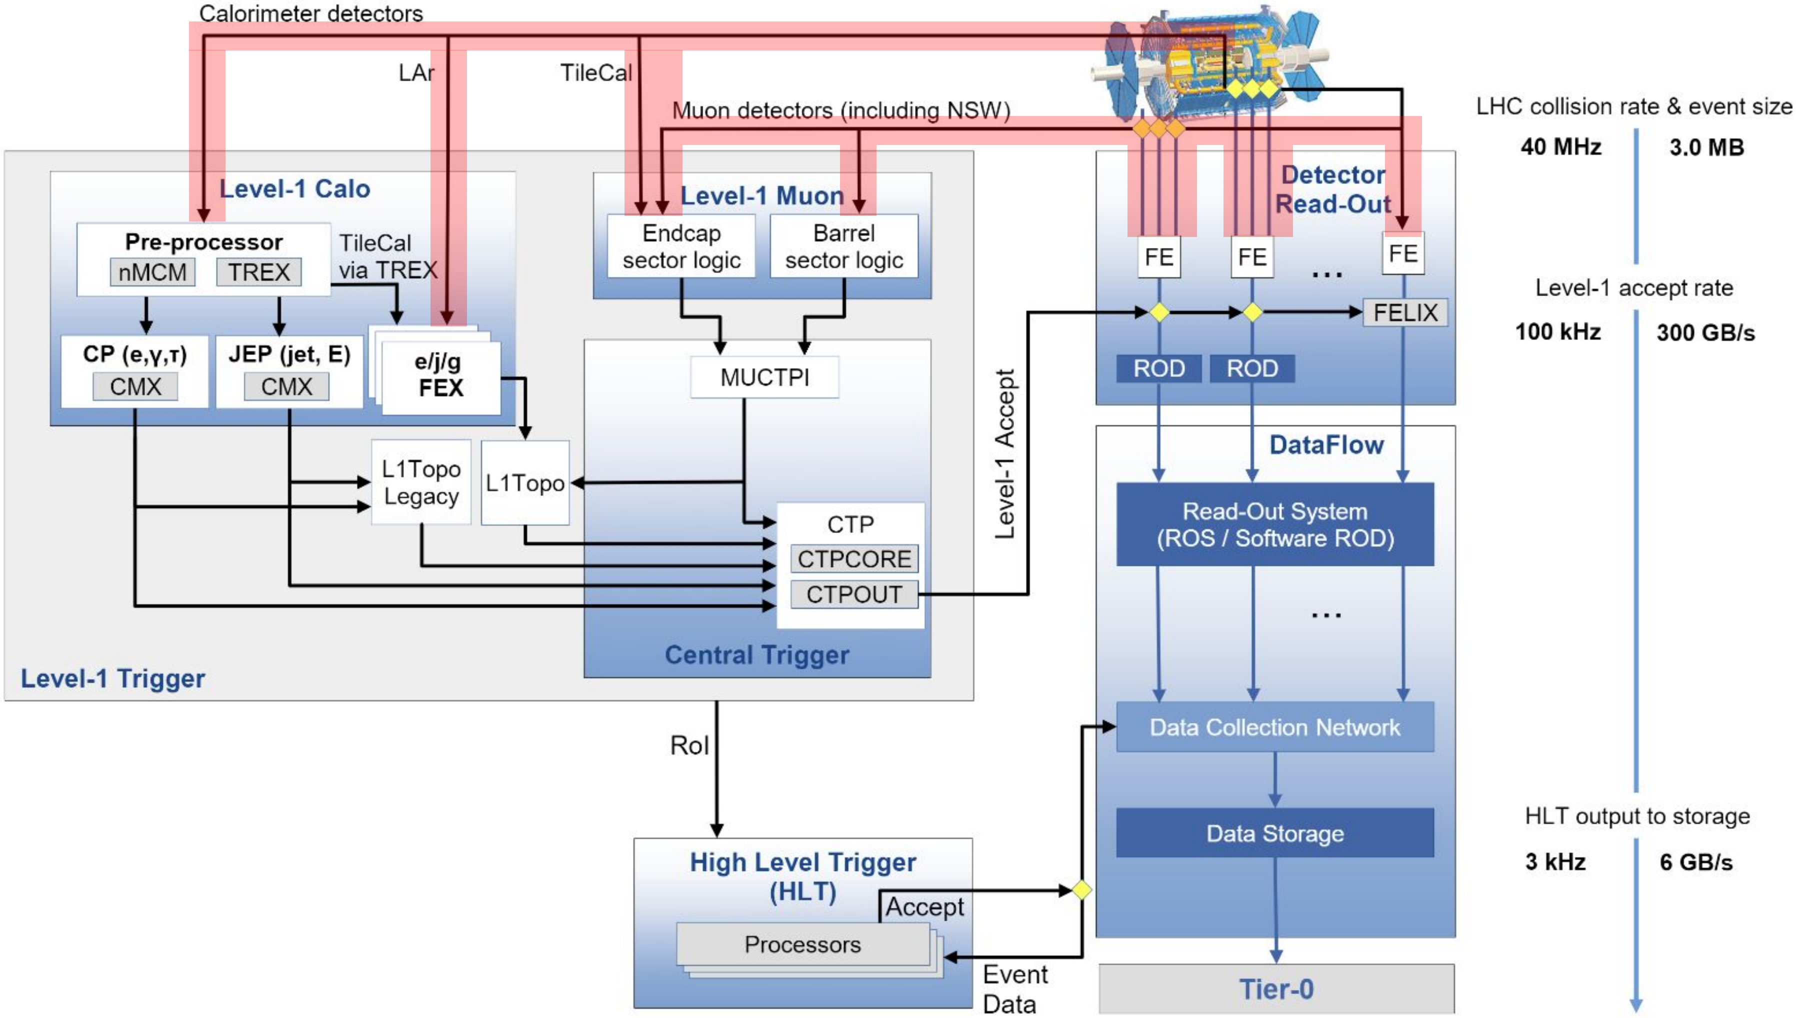
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 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)
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

- 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**.

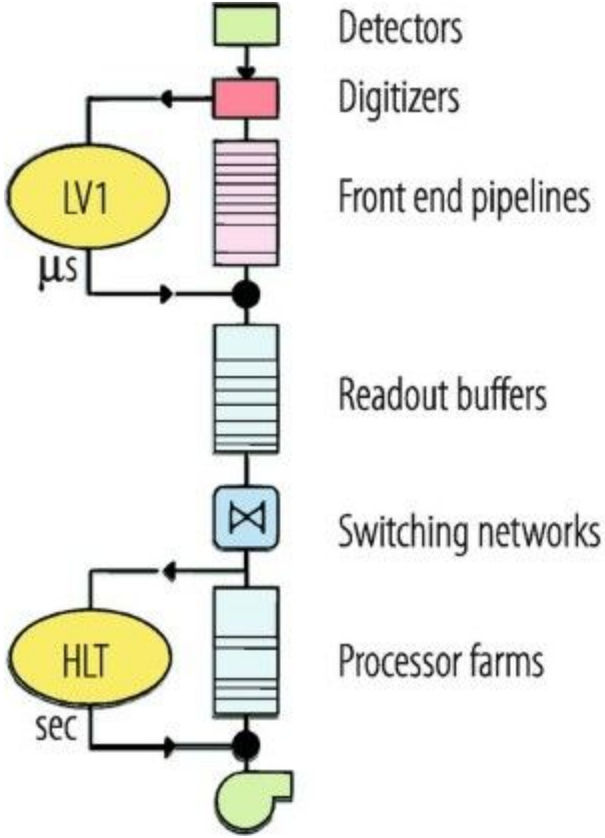

Electromagnetic Feature EXtractor (eFEX)

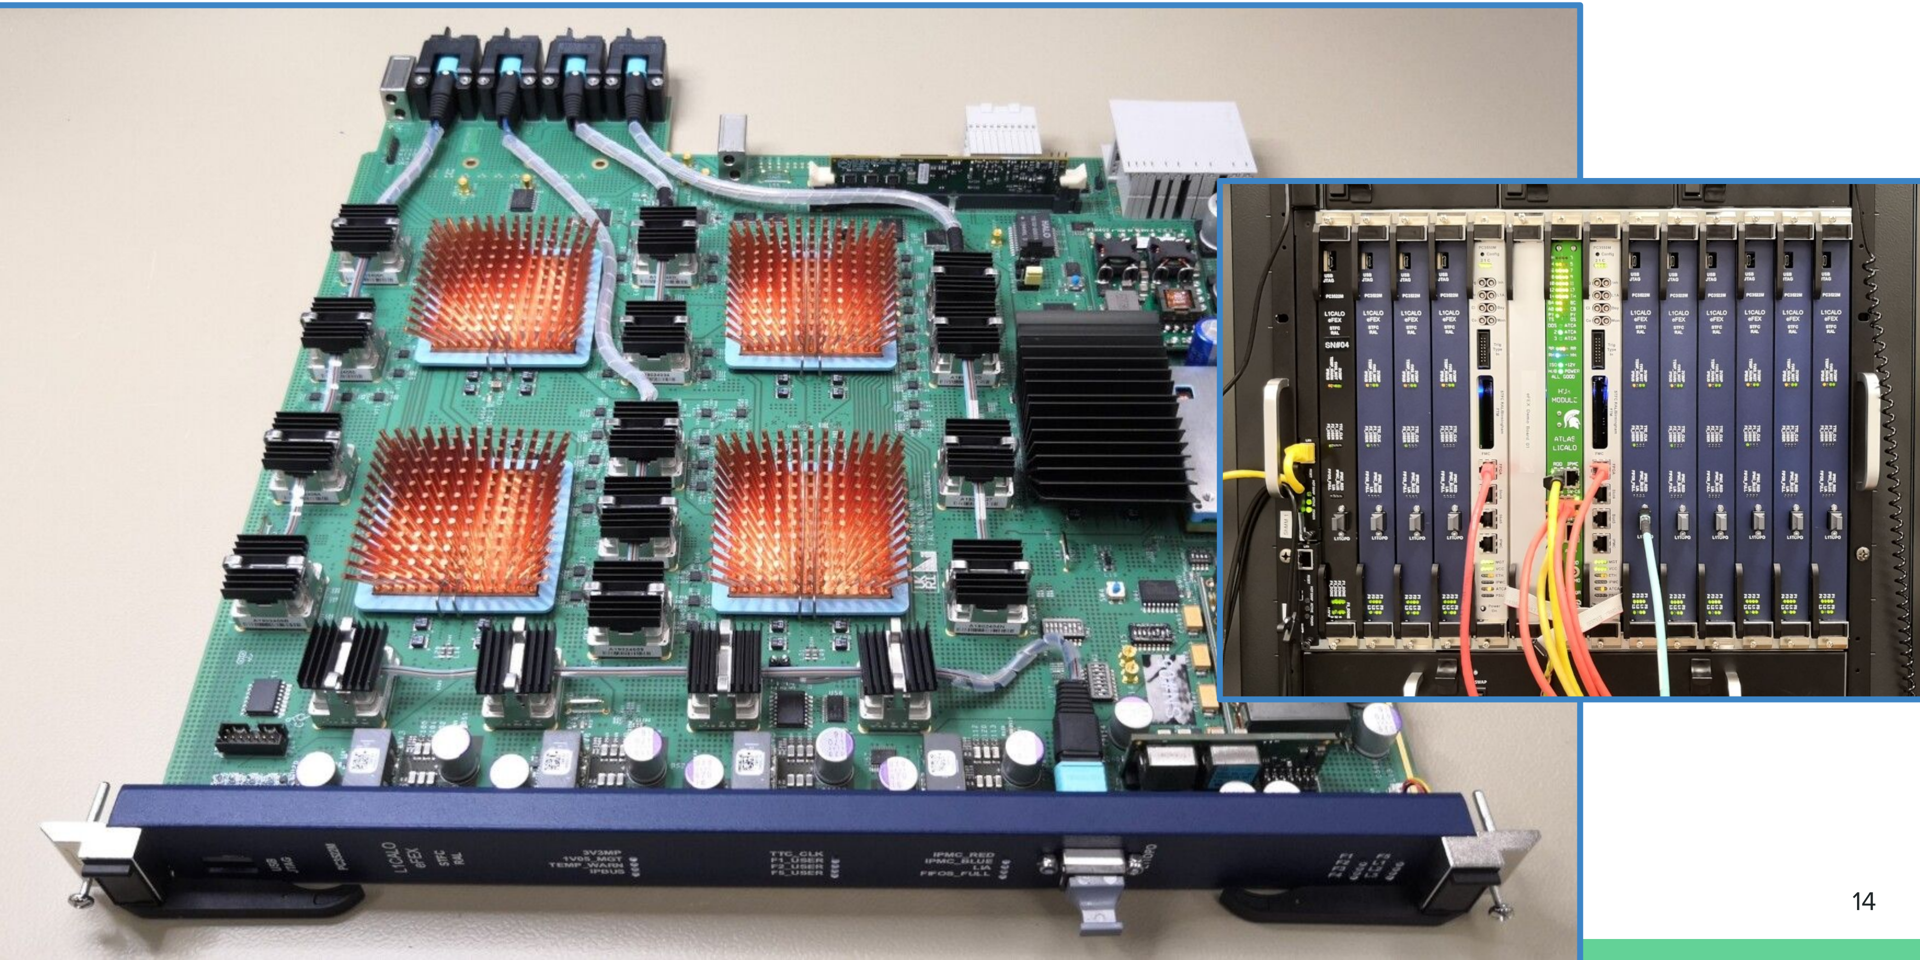

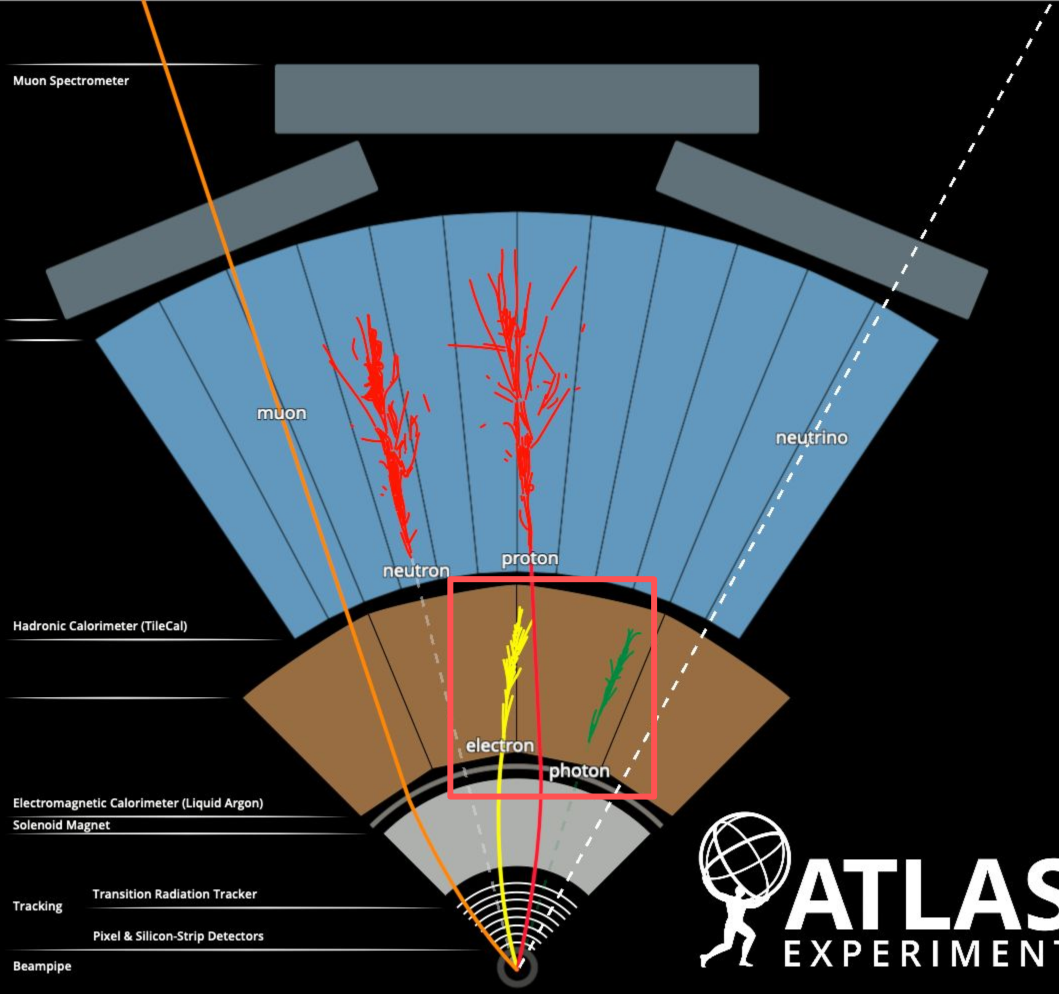

ATLAS
EXPERIMENT

- Before 2022, the L1 trigger only had a single e.m. energy reading per 0.1 x 0.1 region.
- In our new system, eFEX gets **10 e.m. 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 first electromagnetic layer.

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

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

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

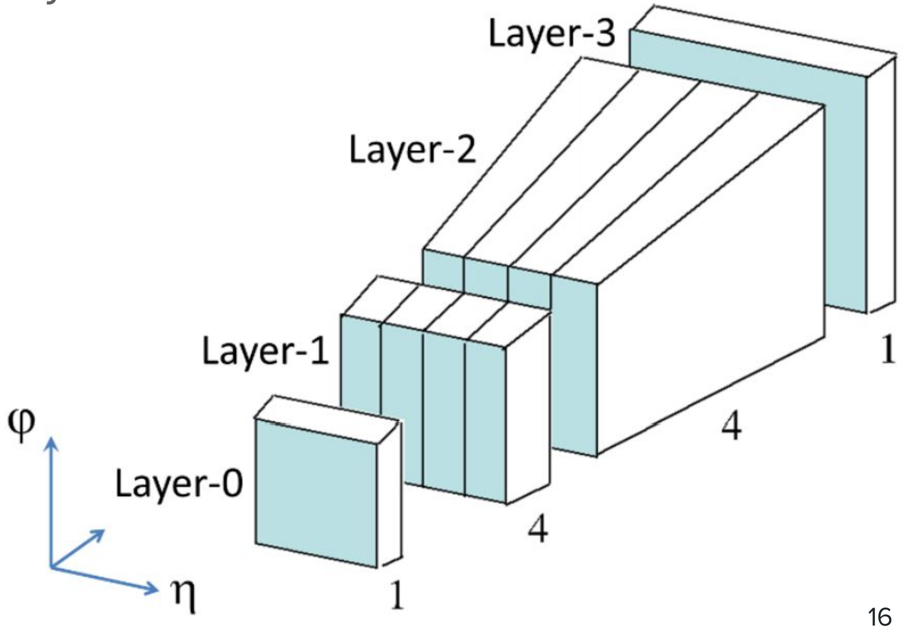

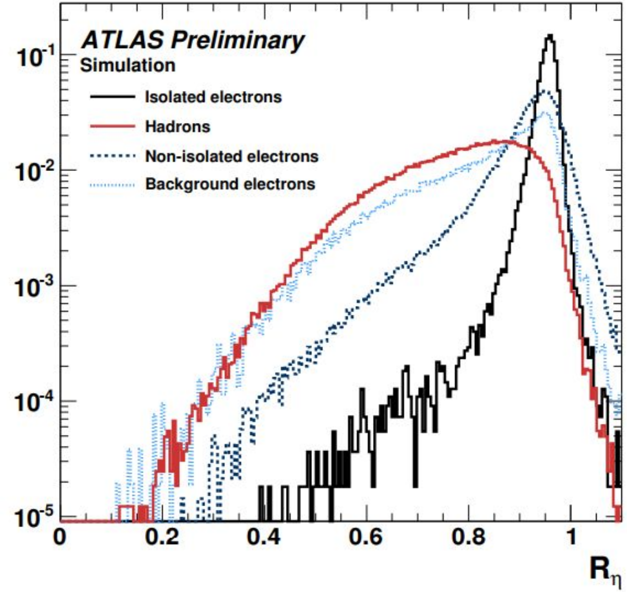

“Good” isolated electrons are from a W or Z

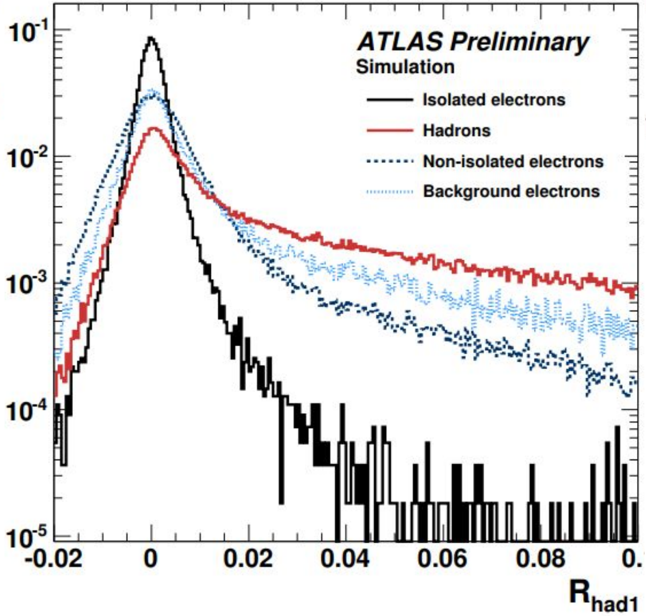

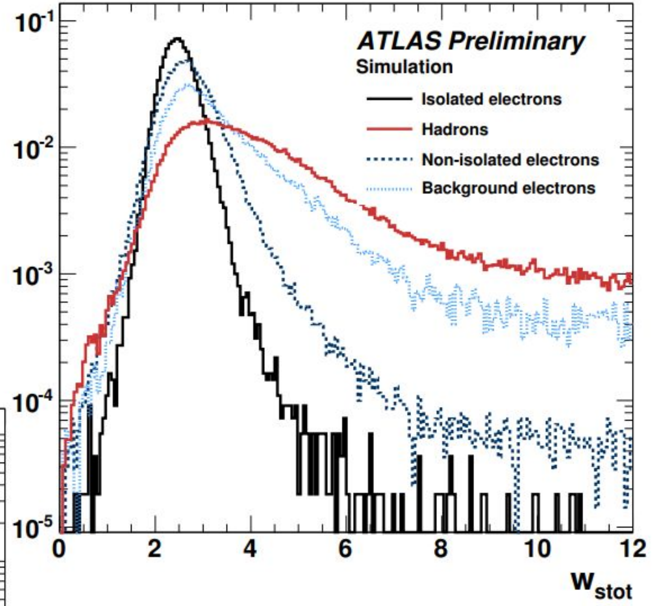

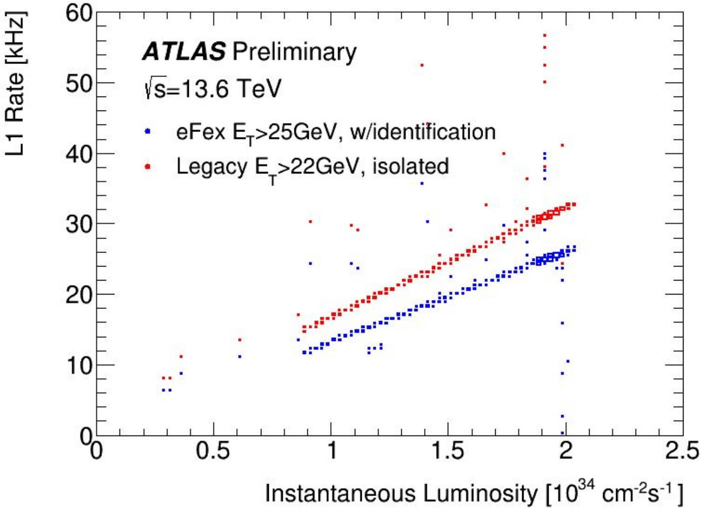

- Lower rate for the same effective energy cut.
- And a higher efficiency too!
- And less of a dependency of on pileup!
- **eFEX: Both purer and more efficient.**
- *(Nomenclature: EM22VHI ~= eEM26M)*

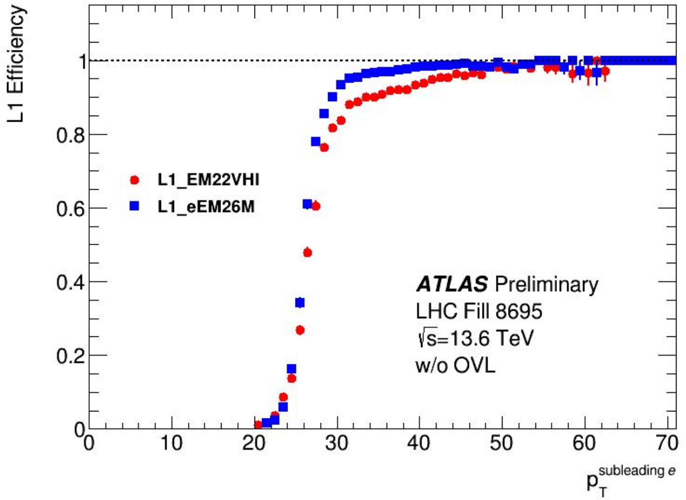

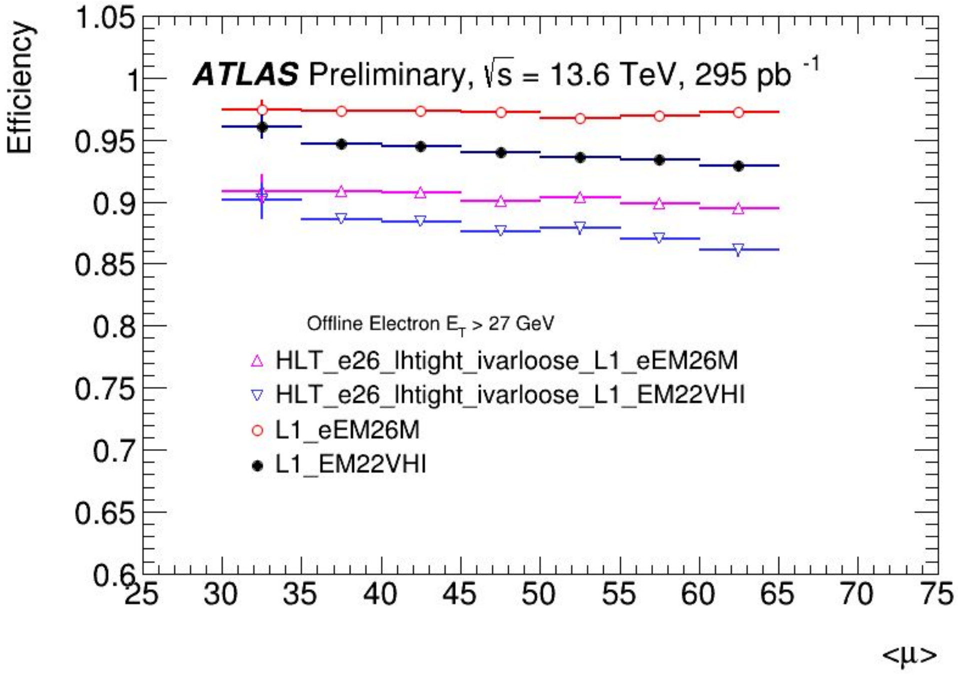

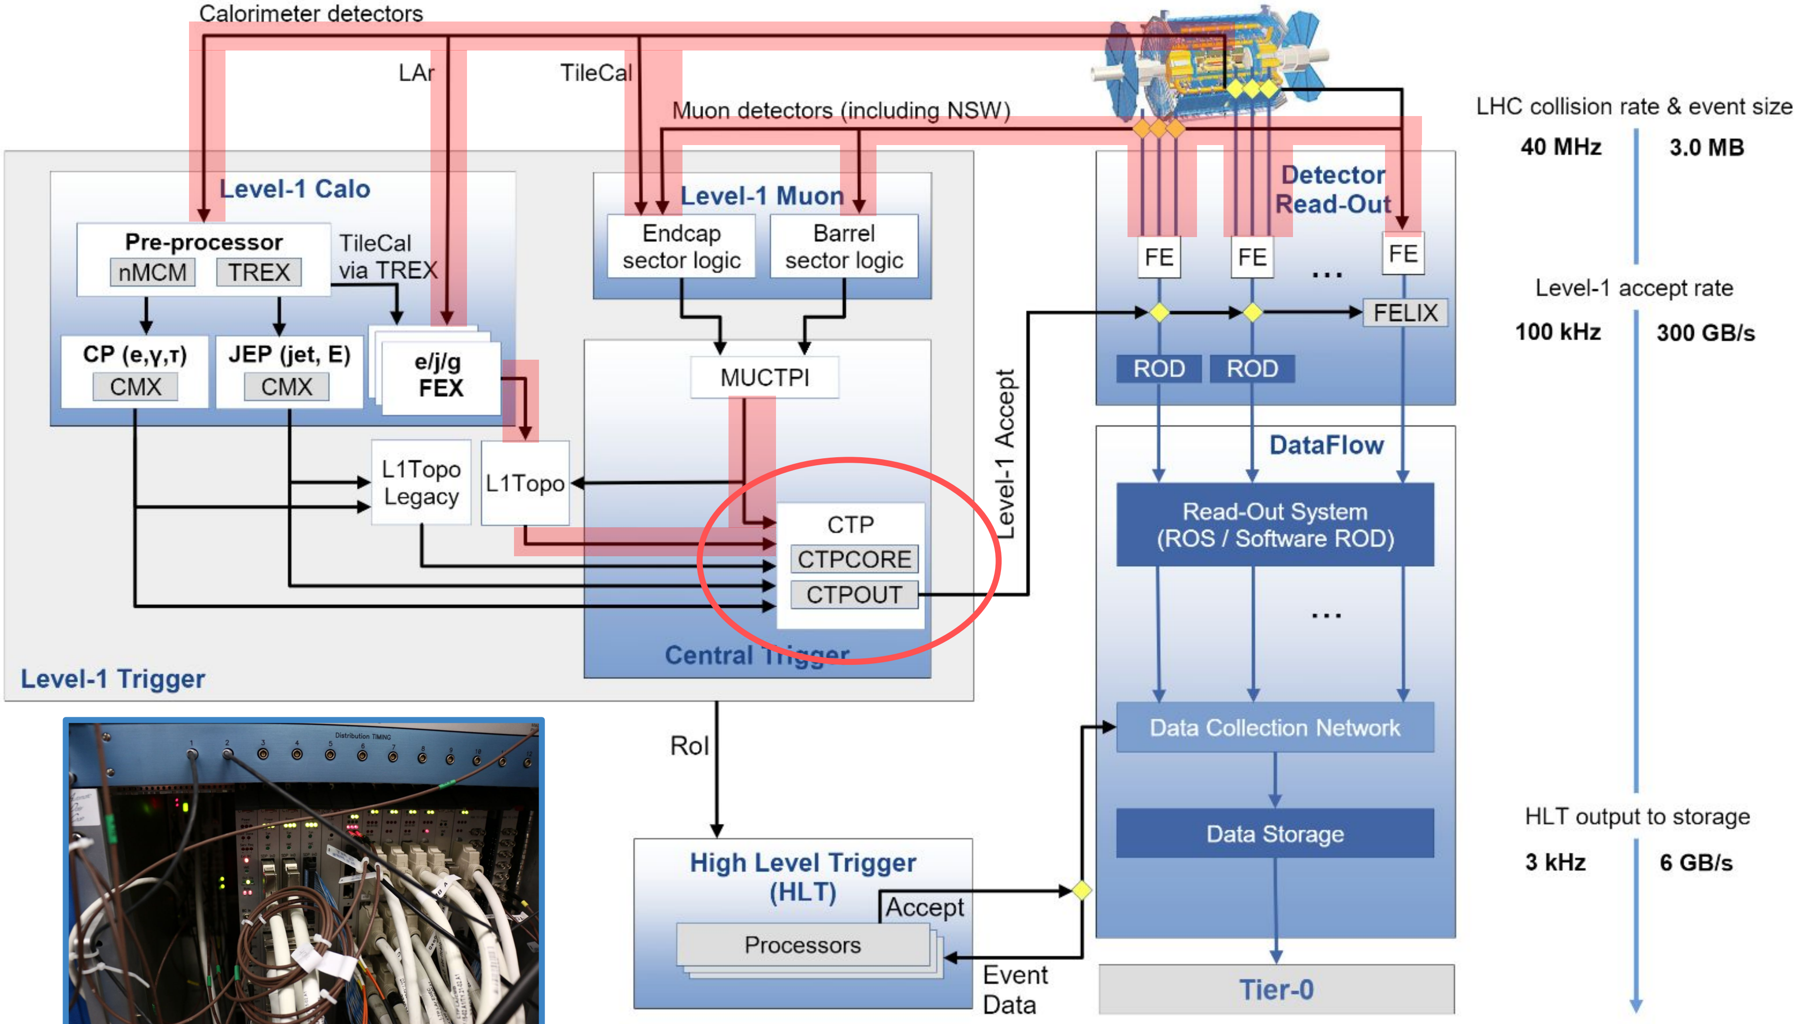

Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 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)
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

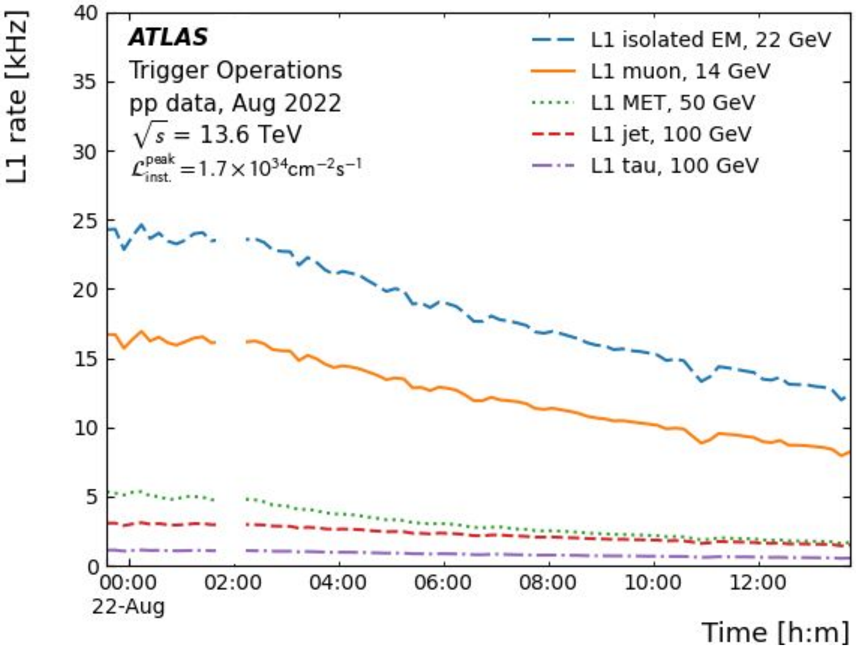

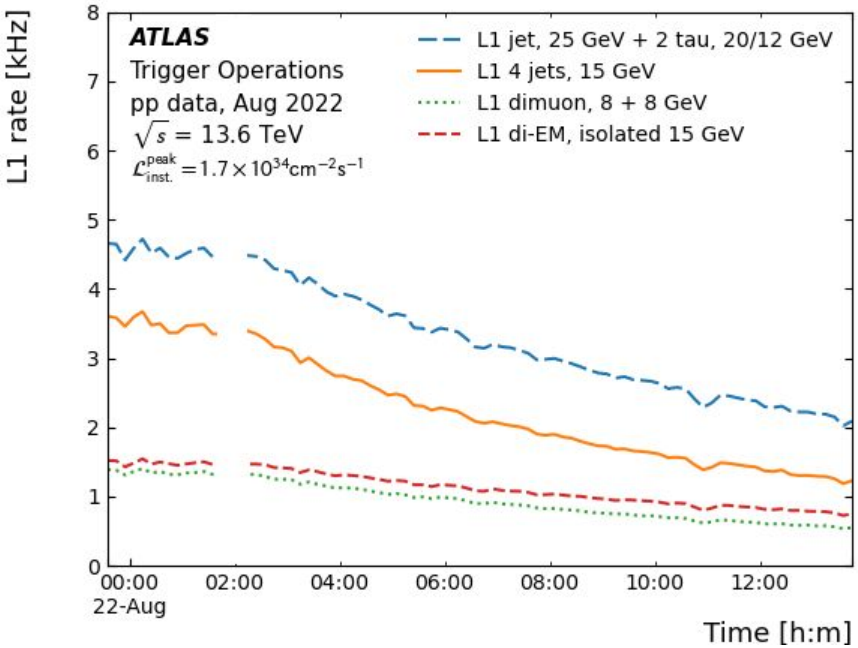

- A *Trigger Menu* converts the collaboration's physics goals into **momentum & identification requirements** on objects identified by the L1 trigger.
- ATLAS' limitation at L1 has always been **100 kHz** (x260 rejection), it will rise to **1 MHz** in 2029/30. We choose to operate at 95 kHz *after veto*, to minimise losses.

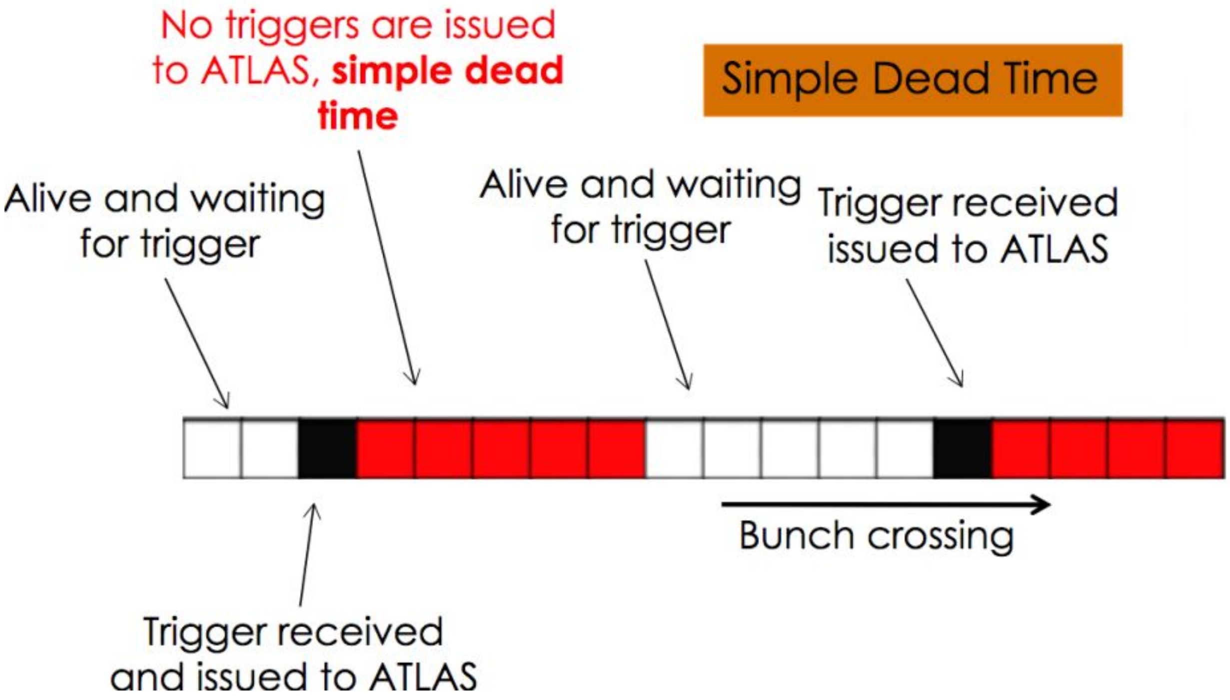

- A *Trigger Menu* converts the collaboration's physics goals into **momentum & identification requirements** on objects identified by the L1 trigger.
- ATLAS' limitation at L1 has always been **100 kHz** (x260 rejection), it will rise to **1 MHz** in 2029/30. We choose to **target 95 kHz *after veto***, to minimise losses.

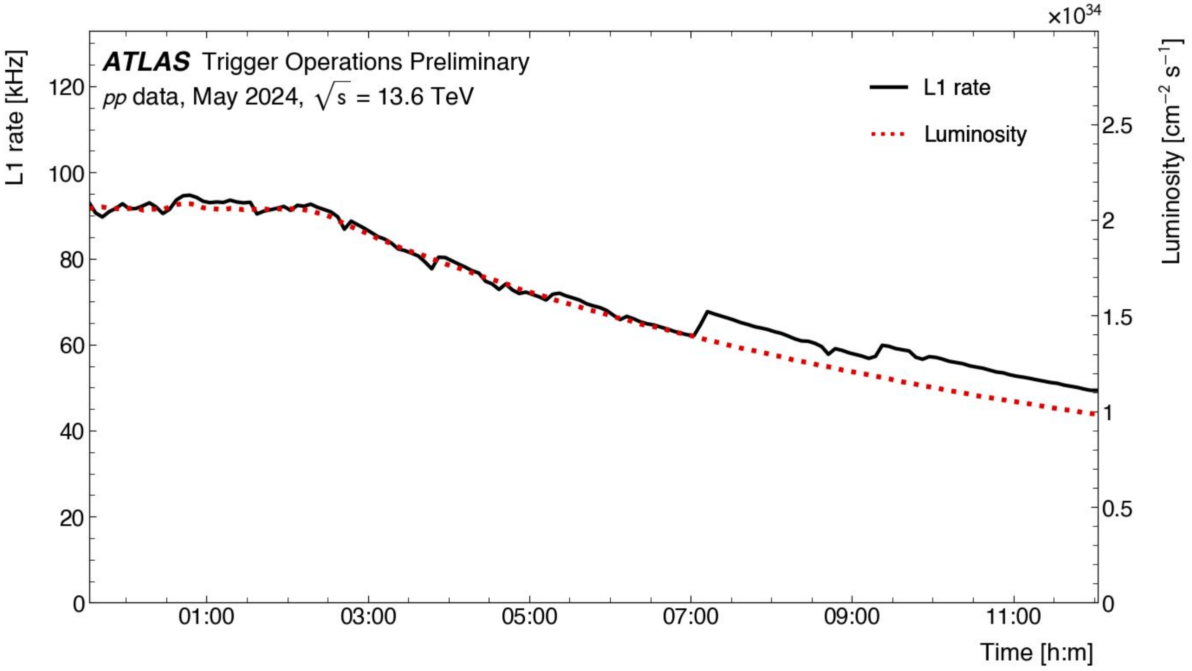

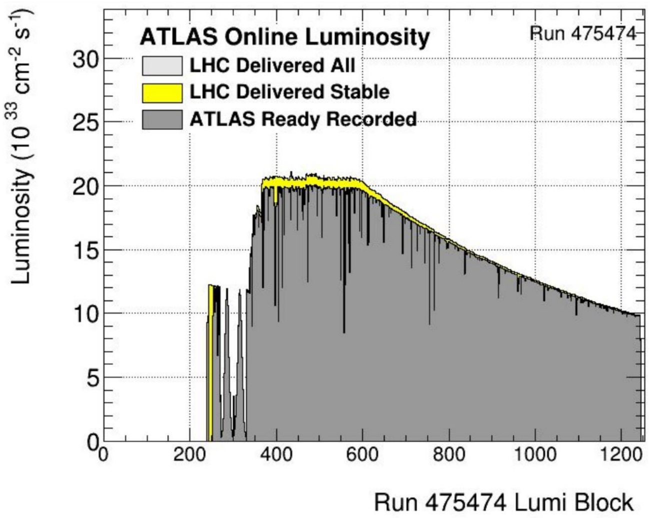

- A *Trigger Menu* converts the collaboration’s physics goals into **momentum & identification requirements** on objects identified by the L1 trigger.
- ATLAS’ limitation at L1 has always been **100 kHz** (x260 rejection), it will rise to **1 MHz** in 2029/30. We choose to **target 95 kHz *after veto***, to minimise losses.

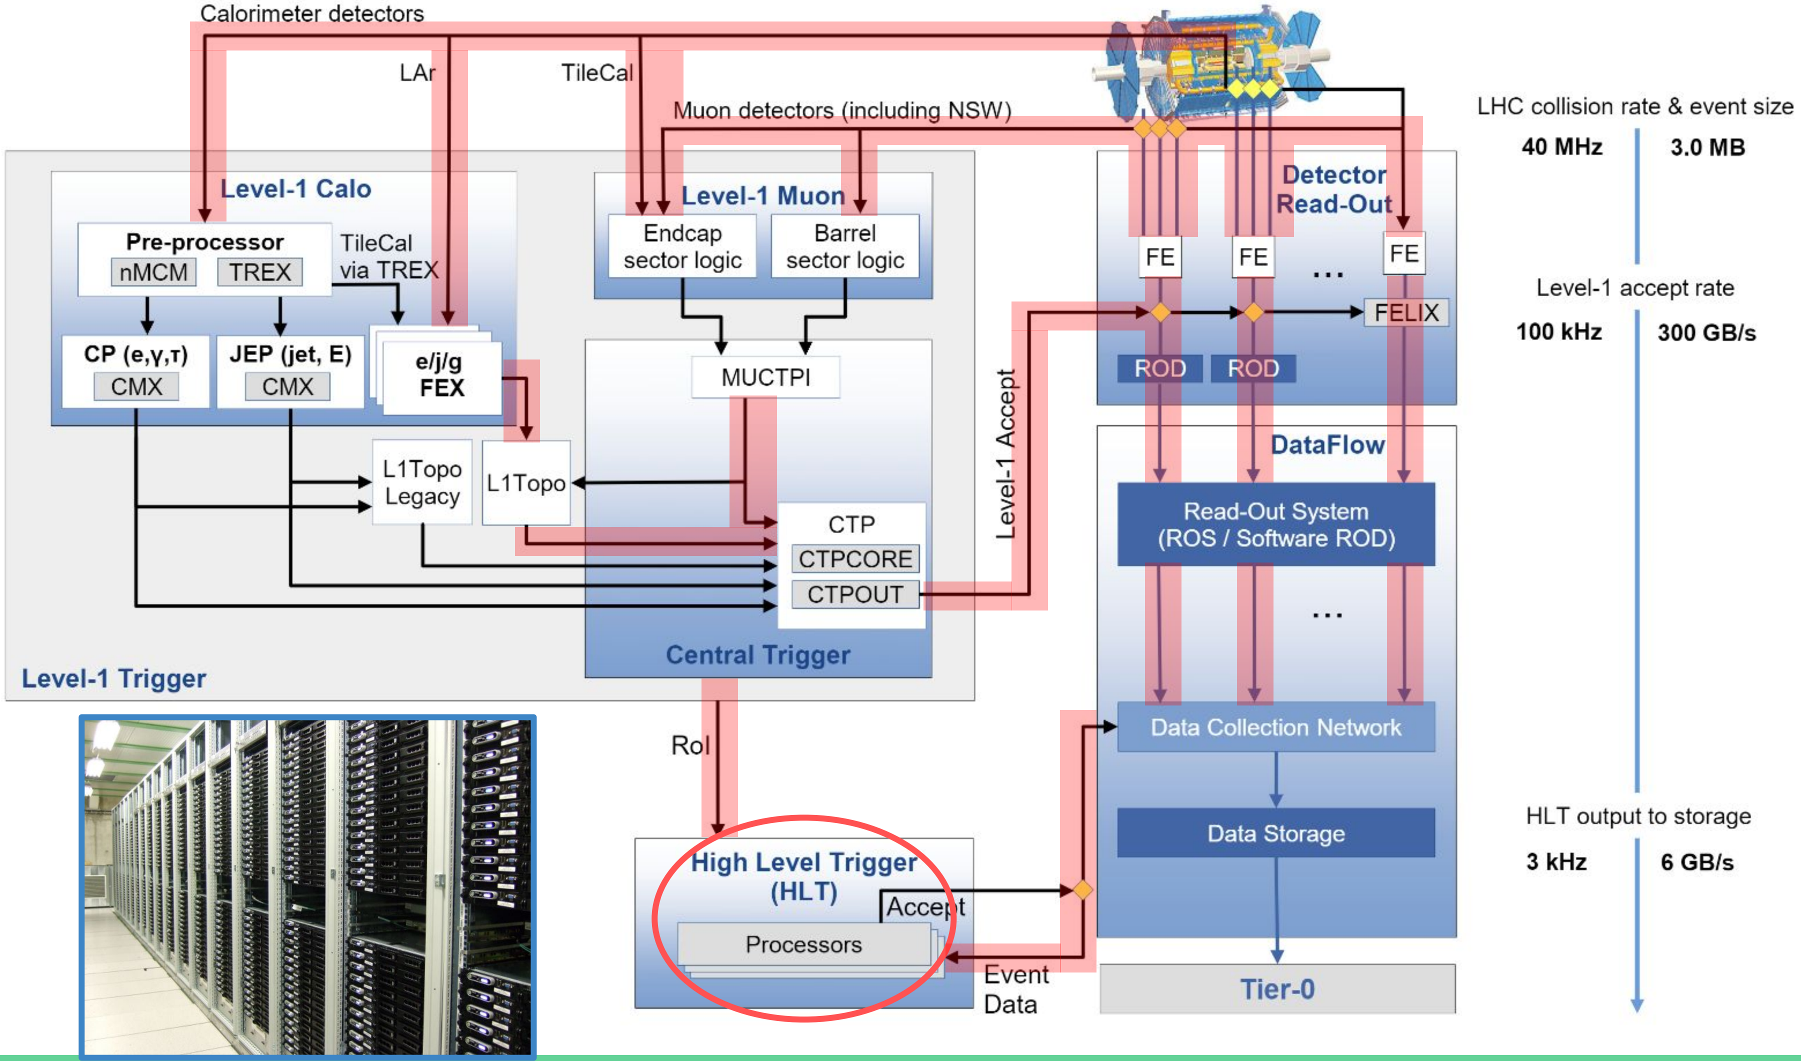
Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
Level-1 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)
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

An Abstraction

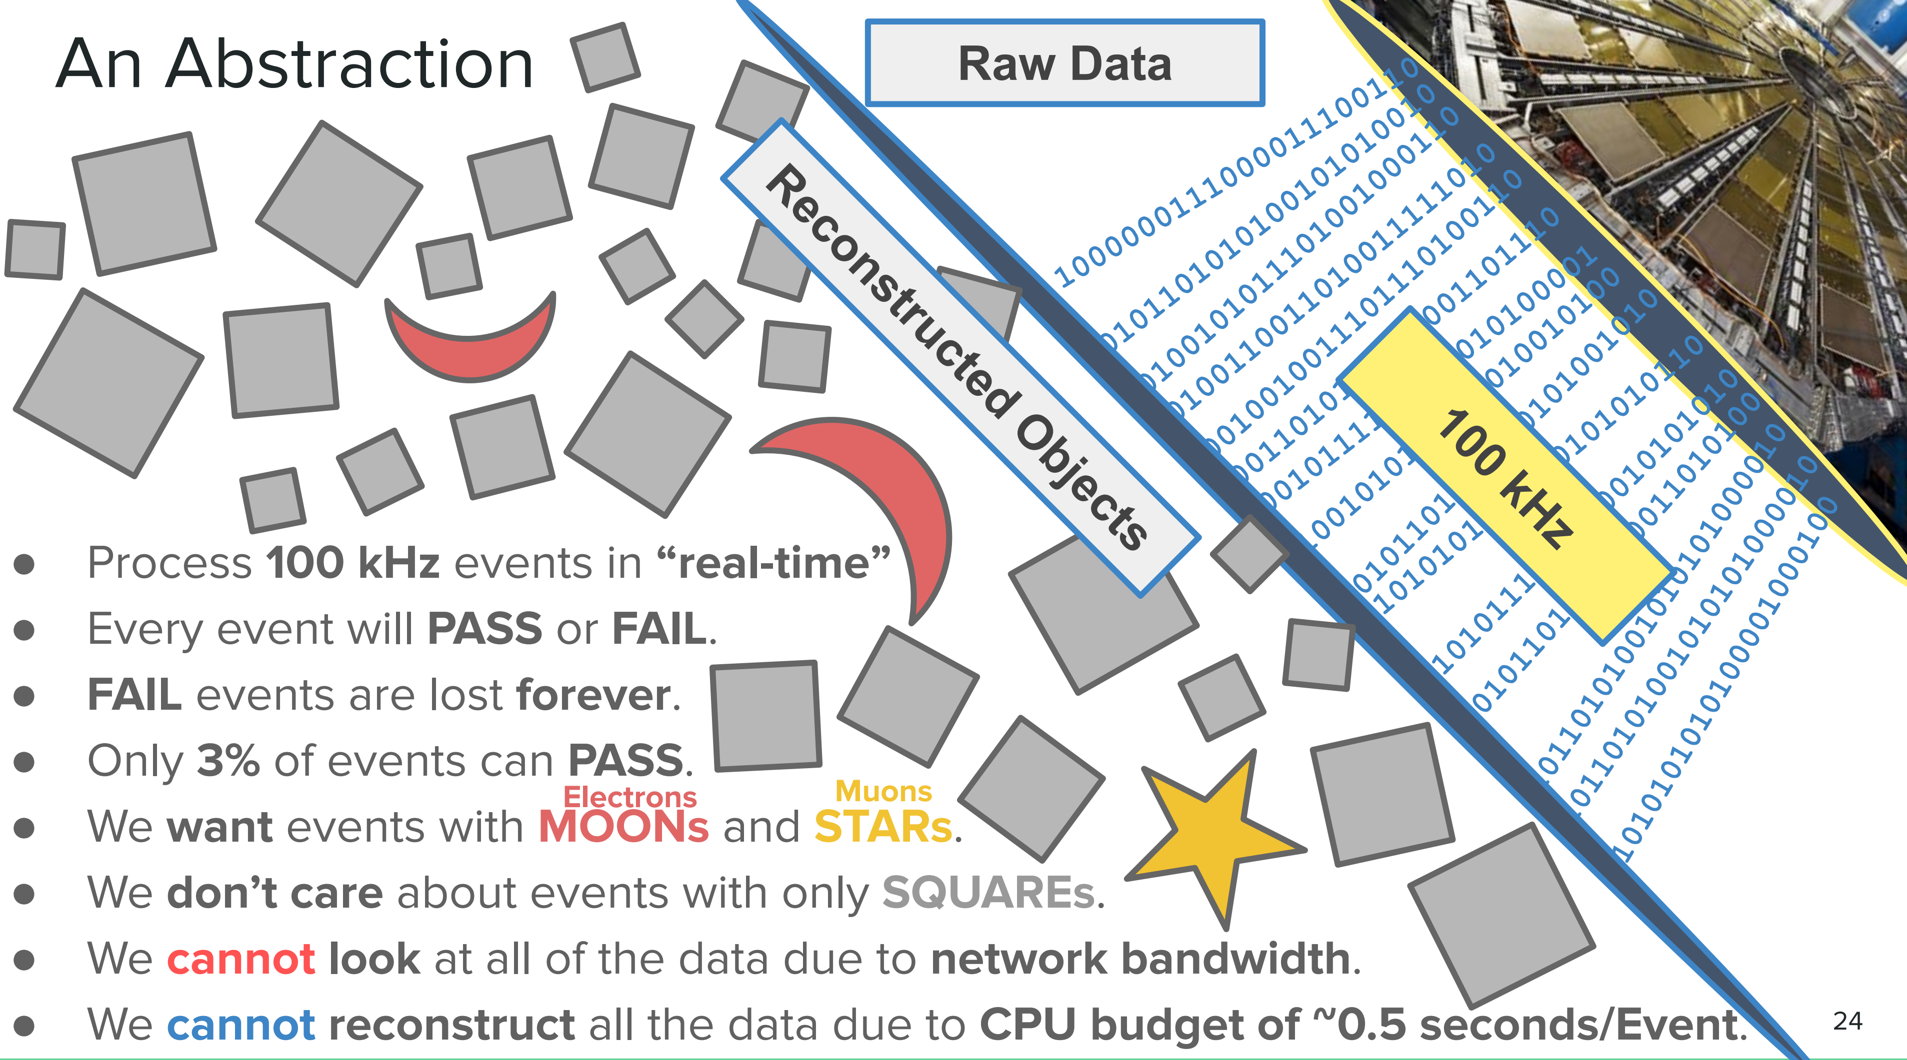

- Process **100 kHz** events in **“real-time”**
- Every event will **PASS** or **FAIL**.
- **FAIL** events are lost **forever**.
- Only **3%** of events can **PASS**.
- We **want** events with **MOONs** (Electrons) and **STARs** (Muons).
- We **don’t care** about events with only **SQUAREs**.
- We **cannot 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**.

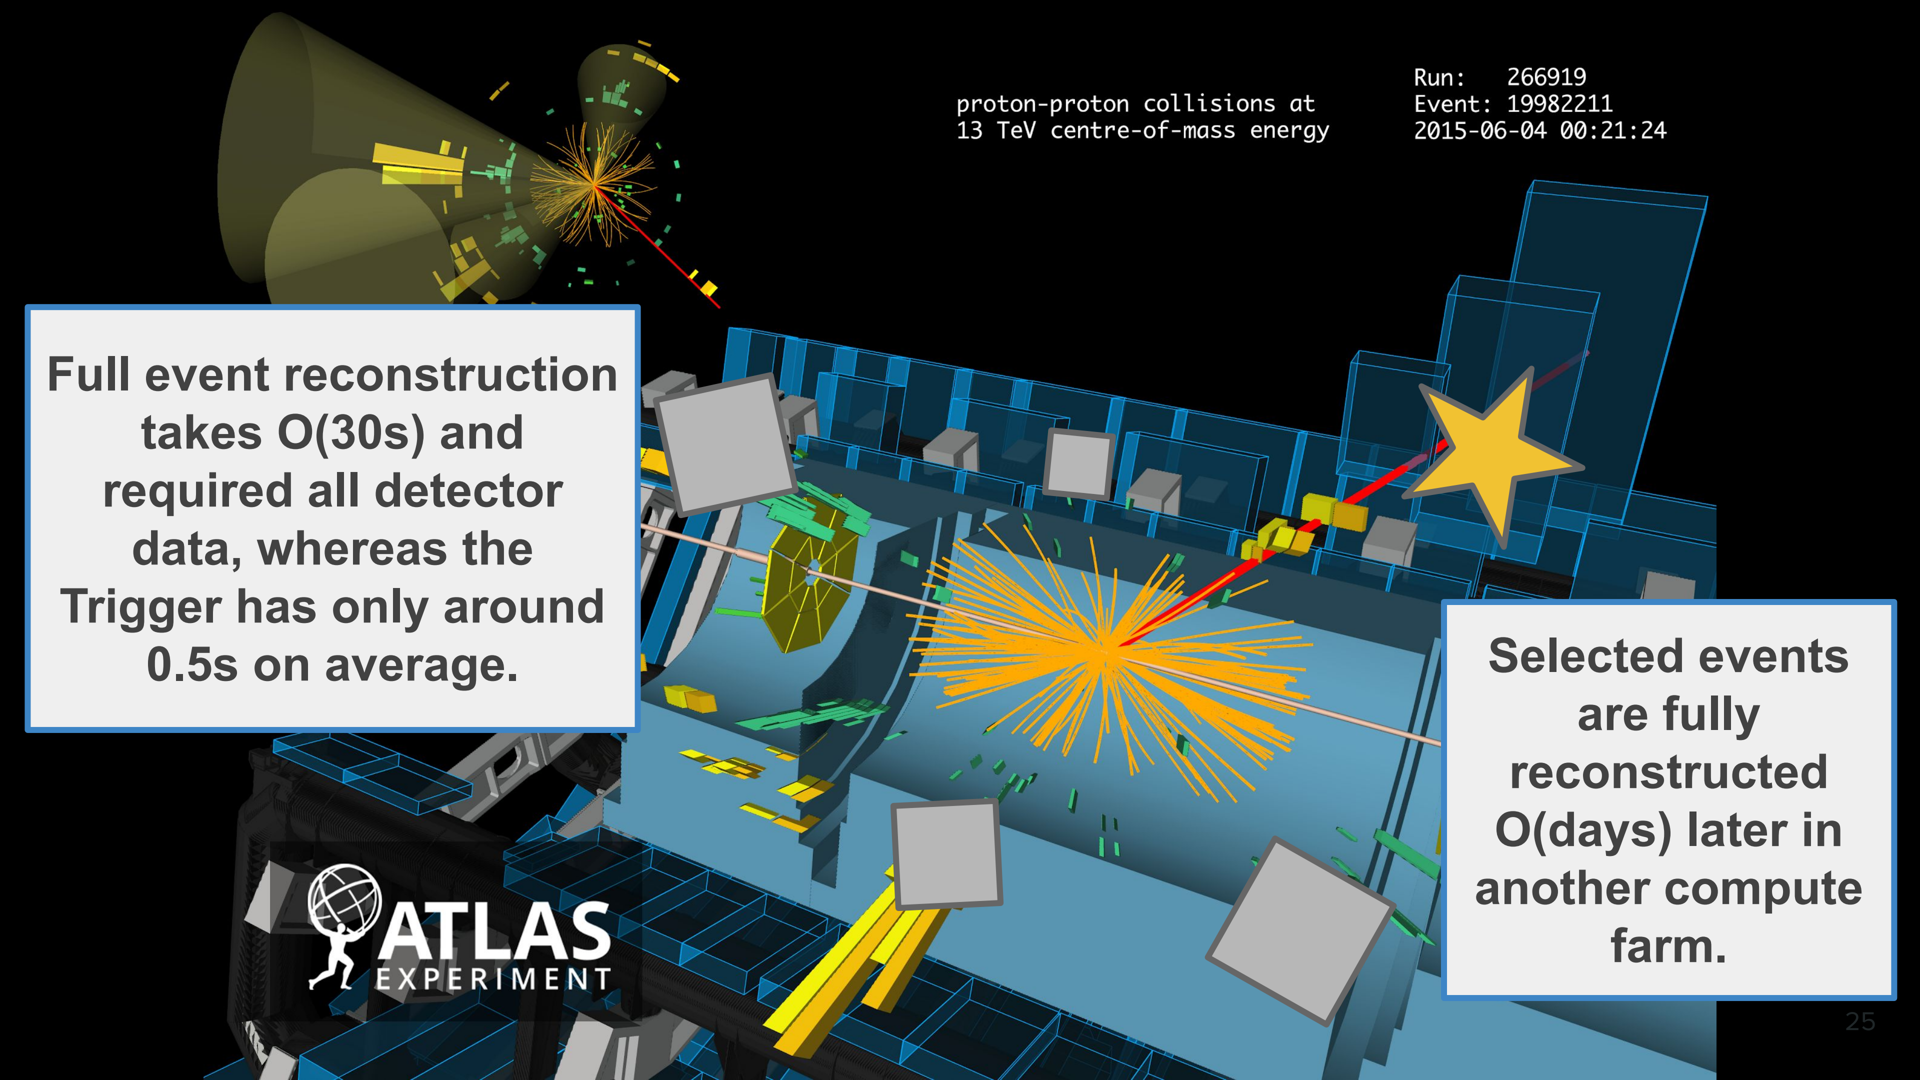

Full event reconstruction takes O(30s) and required all detector data, whereas the Trigger has only around 0.5s on average.

Selected events are fully reconstructed O(days) later in another compute farm.

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**.

Region of Interest

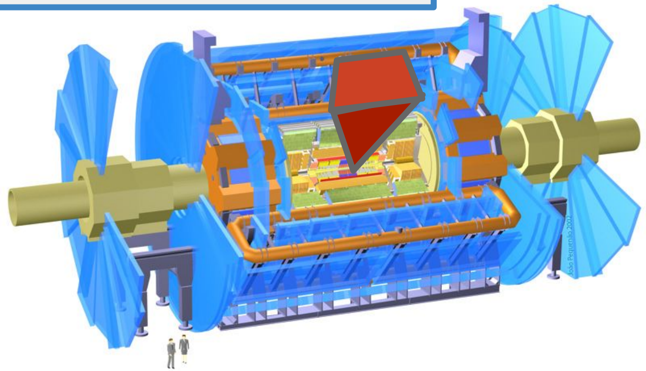

Hypothesis Algorithms

Multi Threaded Athena (Software Framework)

- **ATLAS**, including the **High Level Trigger**, went multi-threaded in 2022.
- **AthenaMT** is built on the **Gaudi Hive (Intel TBB)** multi-threaded architecture.
- Offers **Intra-Event parallelisation**.
  - An **Algorithm Scheduler** is configured with the **Input** and **Output Data Handles** of all algorithms. Builds a **Data Dependency** graph.
  - Multiple algorithms within an event can run in parallel, *provided that their input dependencies are satisfied*.
- Offers **Inter-Event parallelisation**.
  - **Multiple events** may be being processed simultaneously: **"in flight"**.
  - Optimal memory efficiency if all algorithms are **re-entrant**, i.e. **stateless** and able to run on **multiple concurrent events** (alternate: cloneable).

The ATLAS HLT was rewritten for 2022 to be steered by the Gaudi MT scheduler, with no HLT-specific code wrappers required.

Multi Threaded Athena

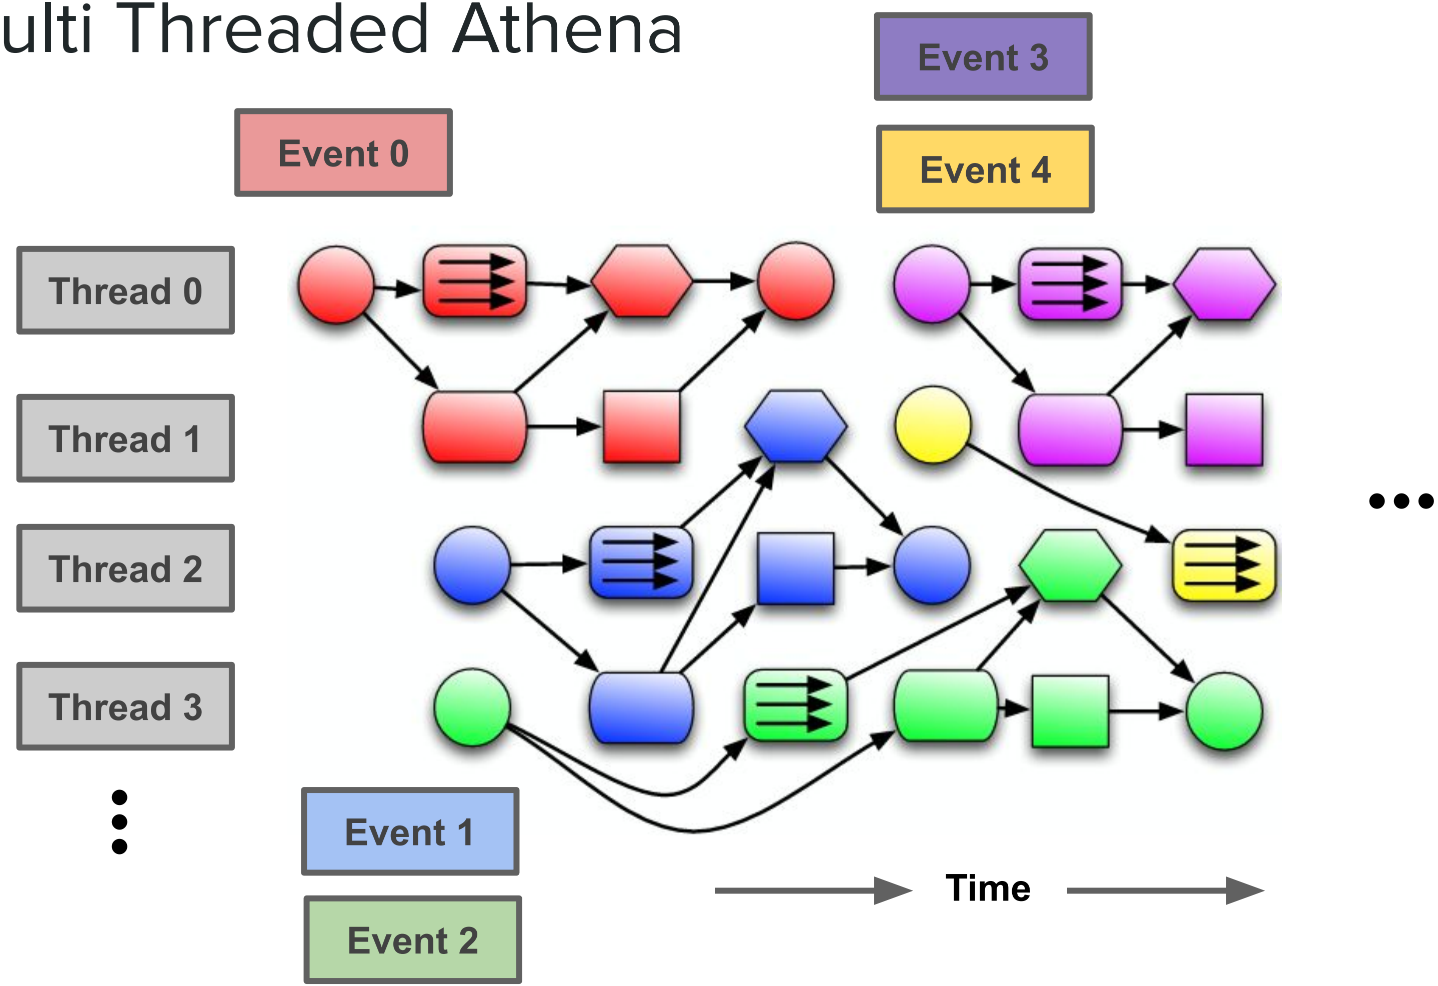

AthenaMT : Controlling Execution

- 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** to the **Gaudi MT Scheduler**.
  - **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).

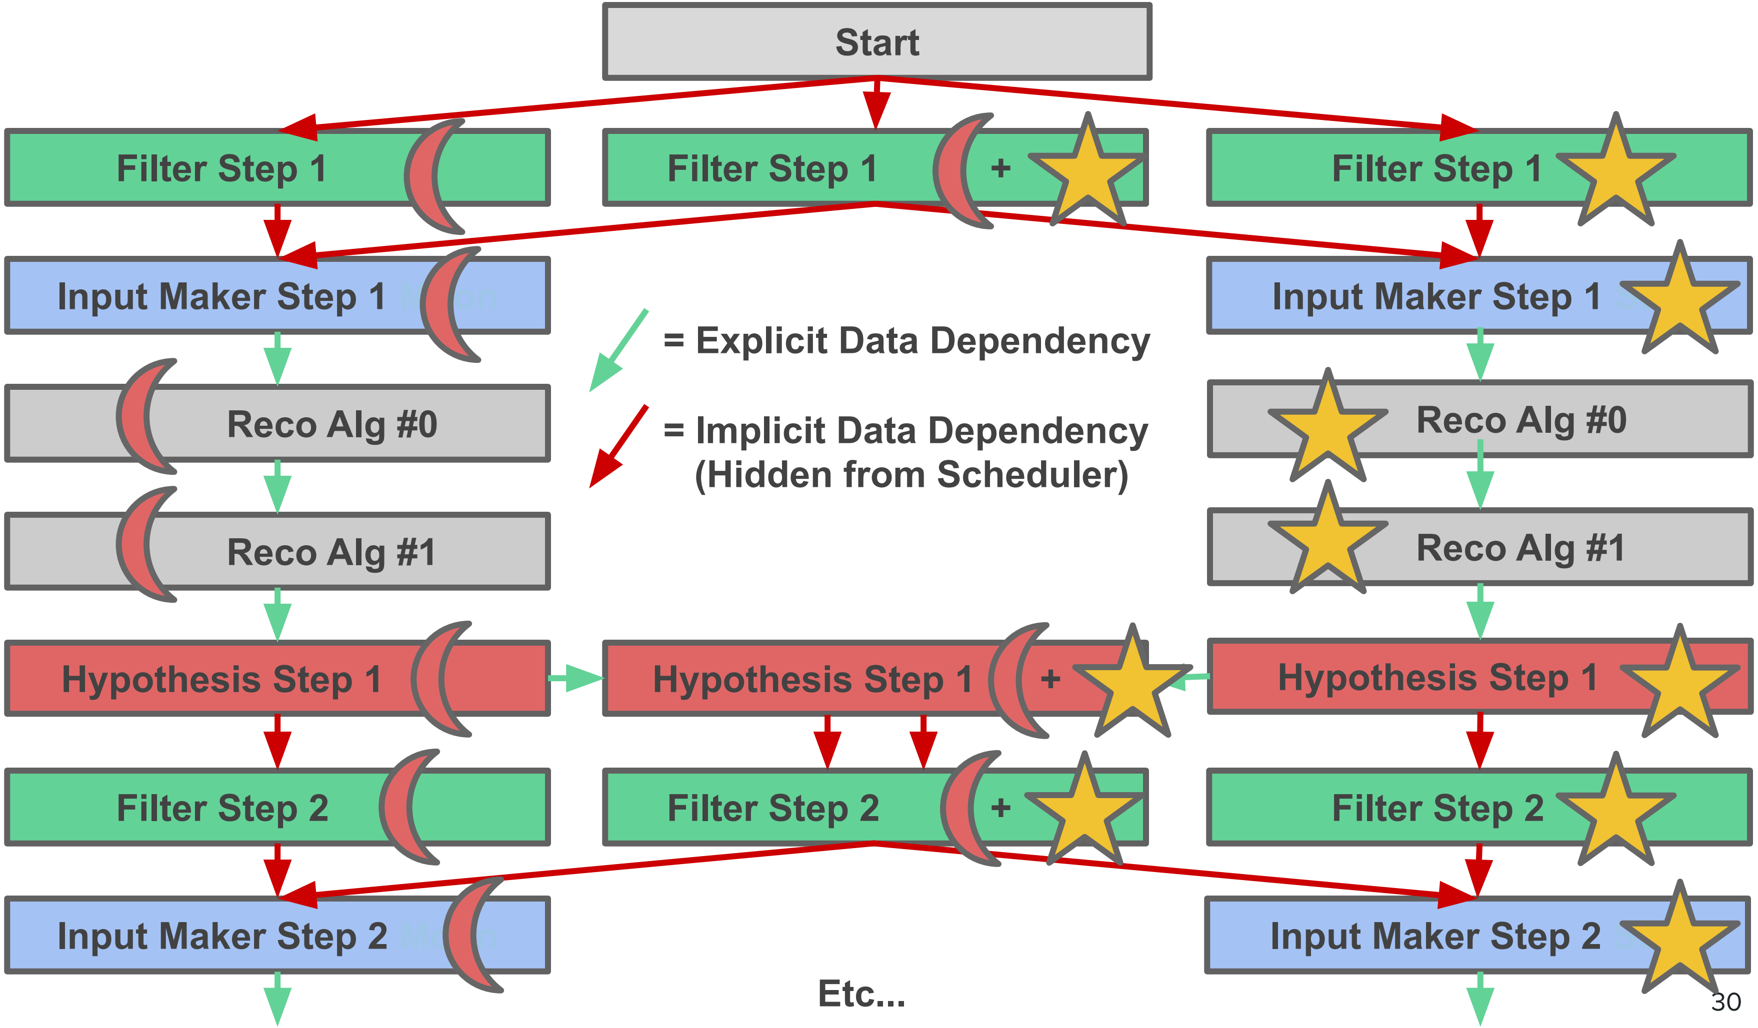

Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
Input Maker Step 1
= Explicit Data Dependency
Reco Alg #0
Reco Alg #0
= Implicit Data Dependency
(Hidden from Scheduler)
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

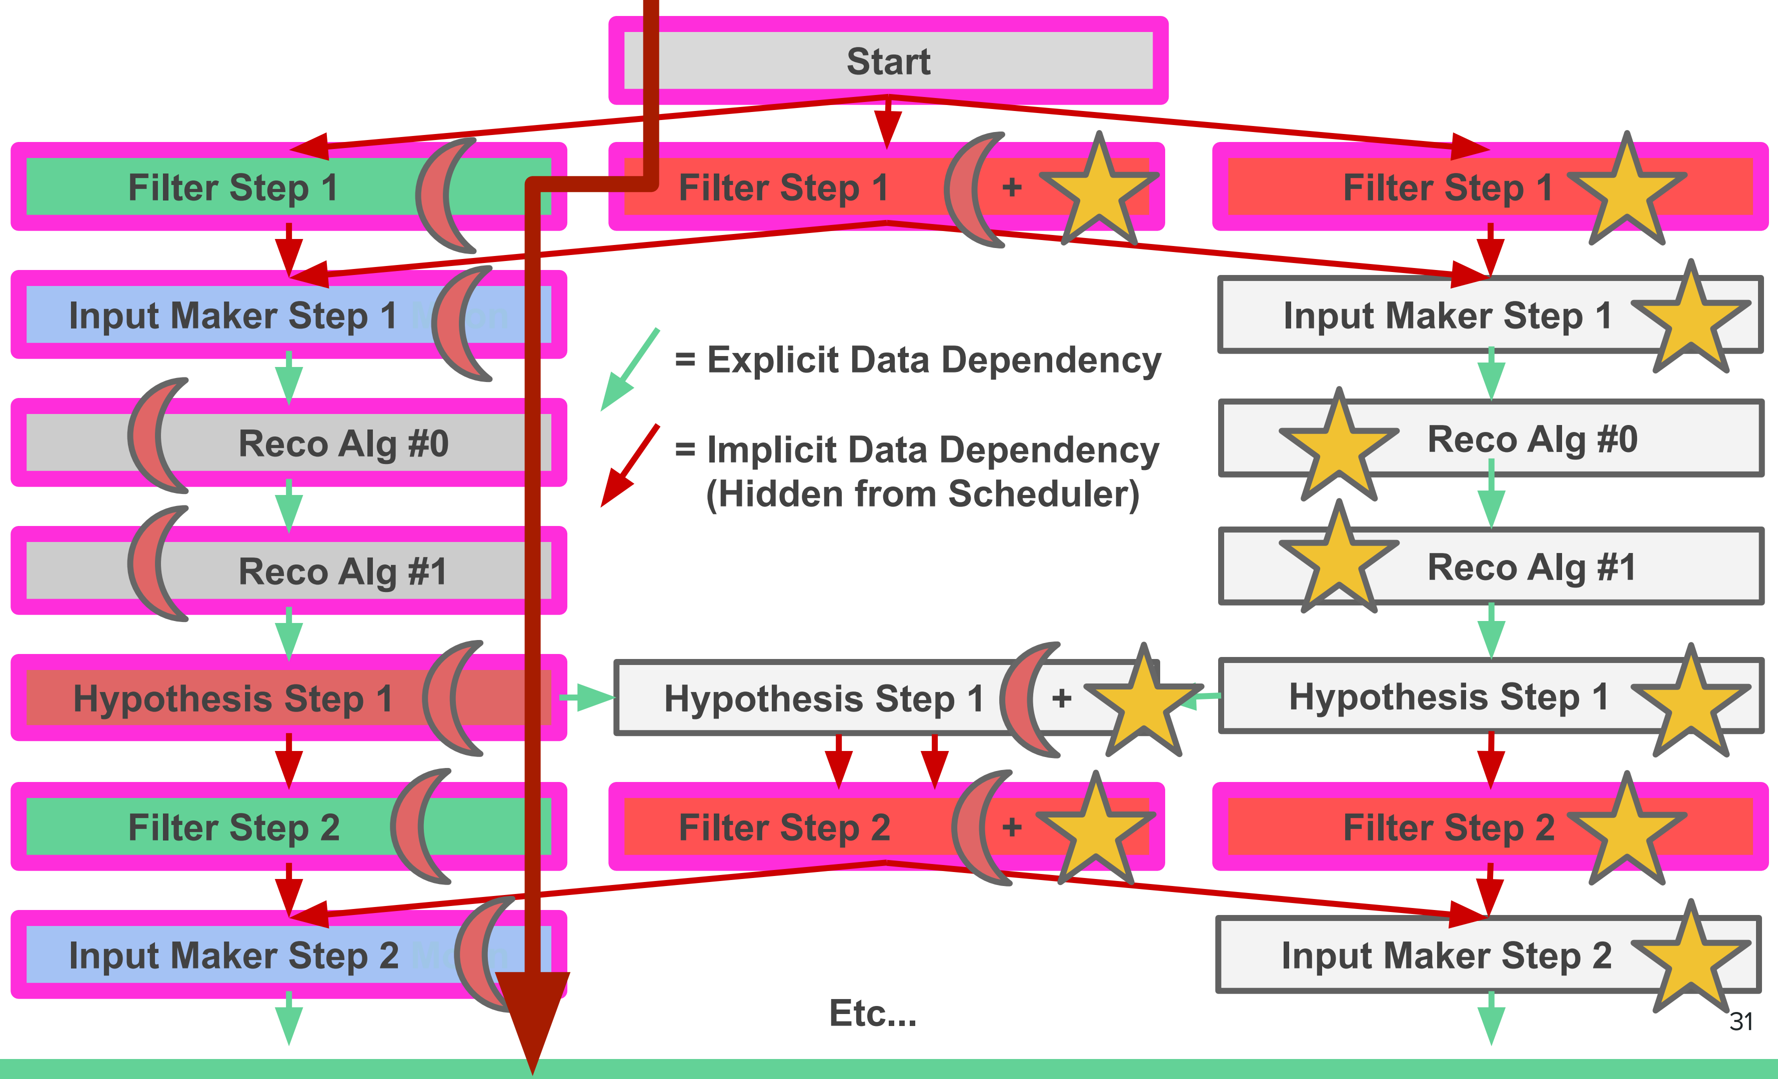

Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
Input Maker Step 1
= Explicit Data Dependency
= Implicit Data Dependency
(Hidden from Scheduler)
Reco Alg #0
Reco Alg #0
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

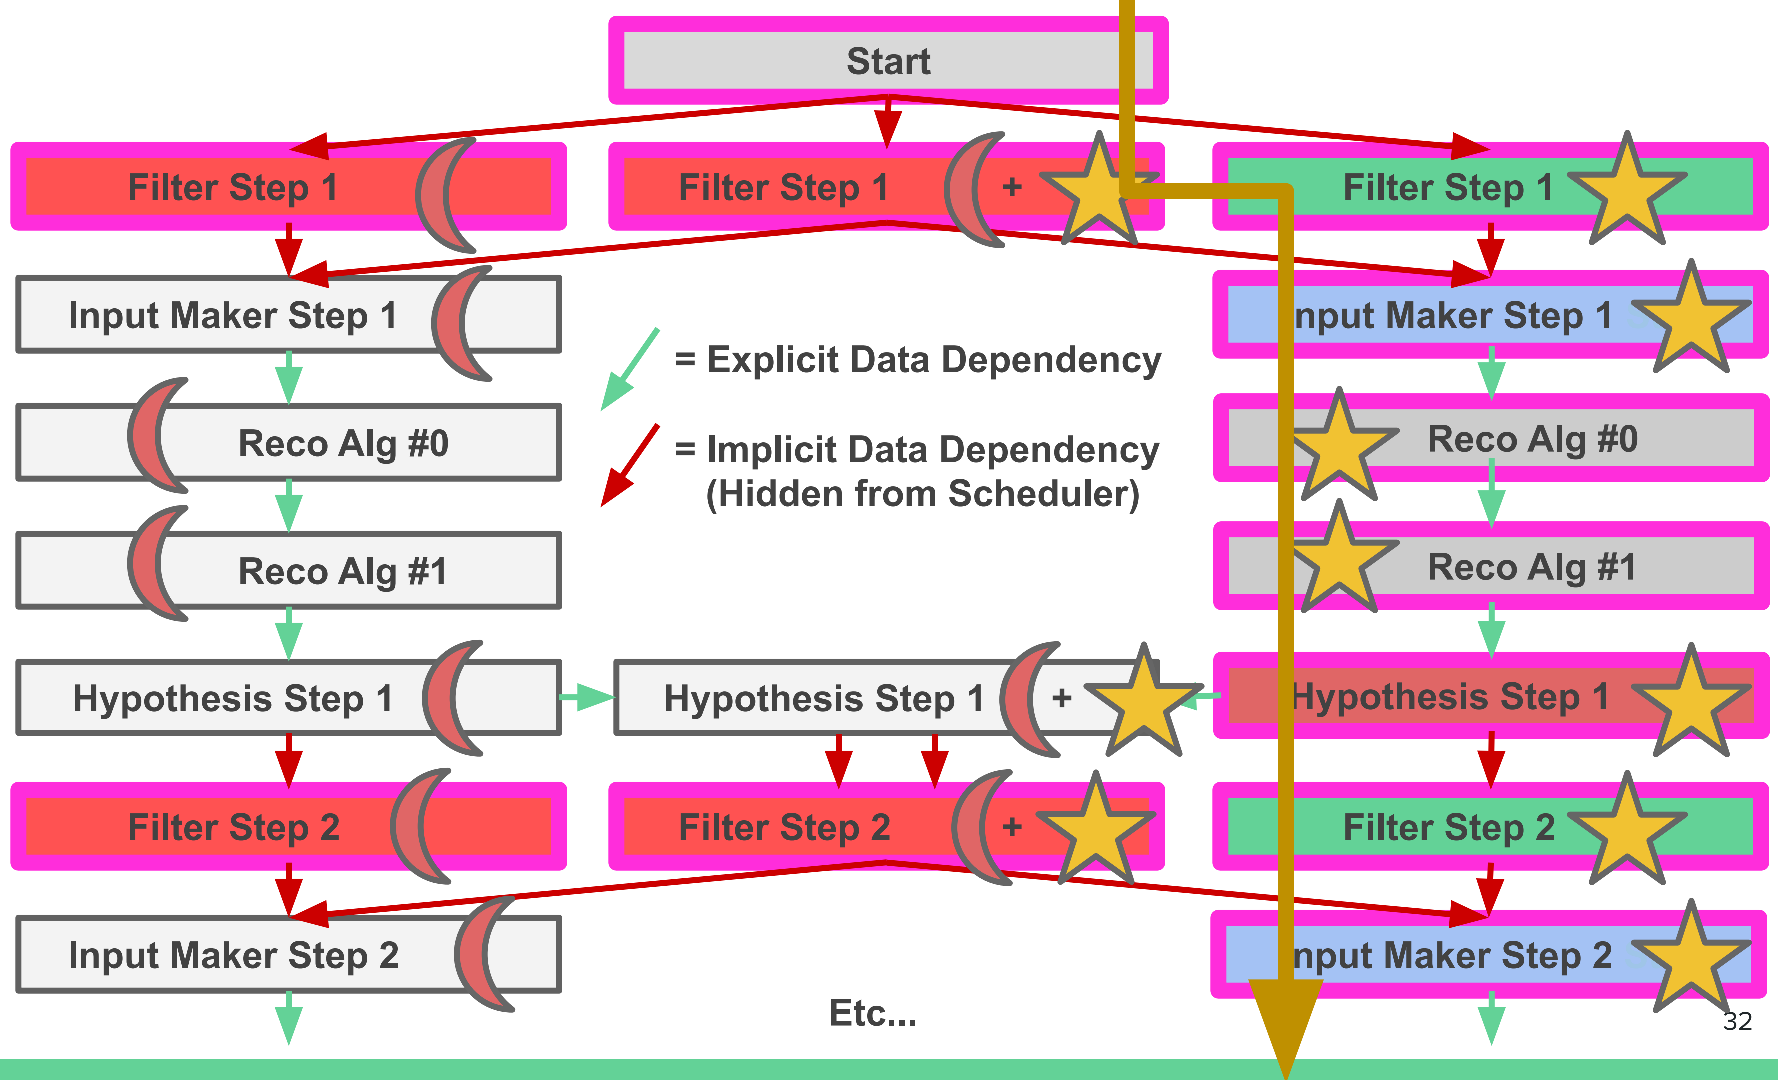
Start
Filter Step 1
Filter Step 1
+
Filter Step 1
Input Maker Step 1
nput Maker Step 1
= Explicit Data Dependency
= Implicit Data Dependency
(Hidden from Scheduler)
Reco Alg #0
Reco Alg #0
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1
+
Hypothesis Step 1
Filter Step 2
Filter Step 2
+
Filter Step 2
Input Maker Step 2
nput Maker Step 2
Etc...

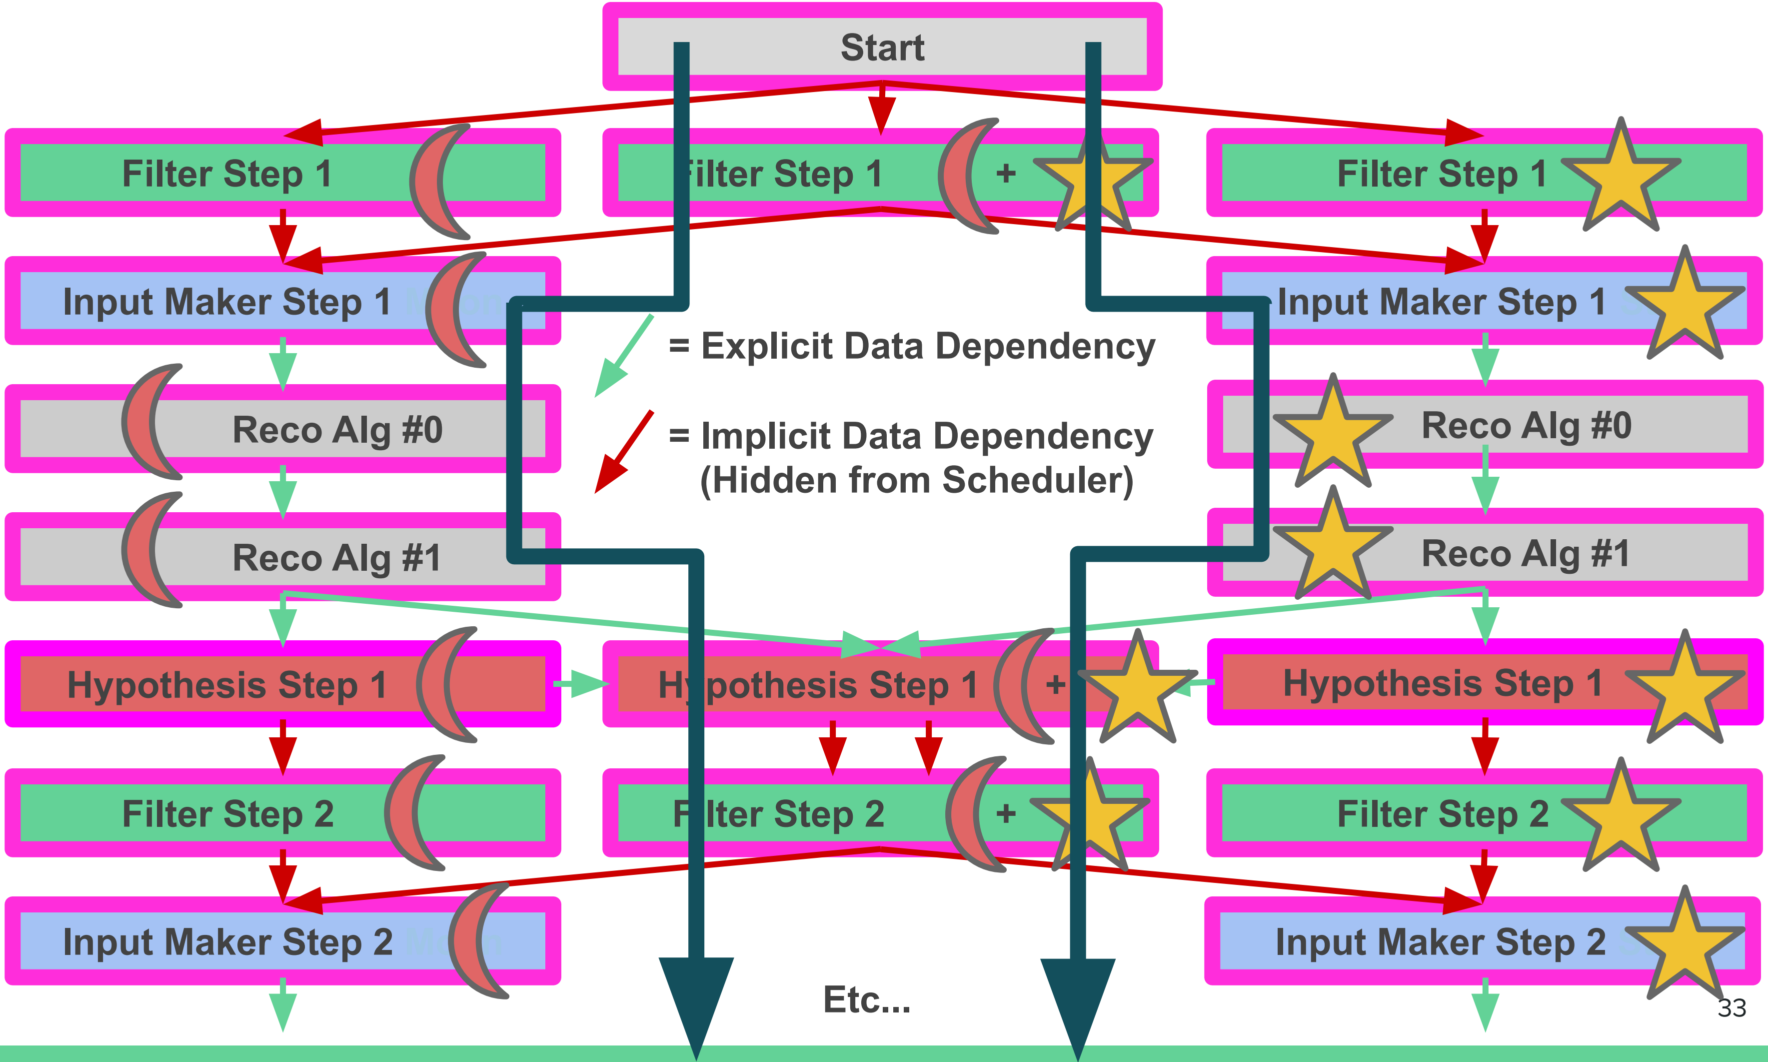
Start
Filter Step 1
Filter Step 1 +
Filter Step 1
Input Maker Step 1
Input Maker Step 1
= Explicit Data Dependency
= Implicit Data Dependency
(Hidden from Scheduler)
Reco Alg #0
Reco Alg #0
Reco Alg #1
Reco Alg #1
Hypothesis Step 1
Hypothesis Step 1 +
Hypothesis Step 1
Filter Step 2
Filter Step 2 +
Filter Step 2
Input Maker Step 2
Input Maker Step 2
Etc...

Regional Reconstruction: Event Views

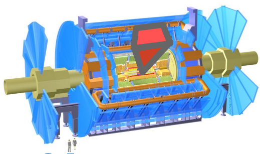

- **Gaudi Hive** will allow **each algorithm** to execute **at most once per event**.
- But **Regional Reconstruction** requires algorithms to run **once per Region of Interest**.
- **ATLAS** Extension: **Event Views**.
  - **Spawn** one **Event View** per **Region of Interest**, Schedule algorithms **per View**.
  - **Event View** Implements the **Event Store** interface.
  - Completely **transparent** to the **algorithm**. It does not know it's in a view. (...but it can find out)
- Applies name-mangling, "`HLT_Muon`" becomes "`MuView_0_HLT_Muon`"

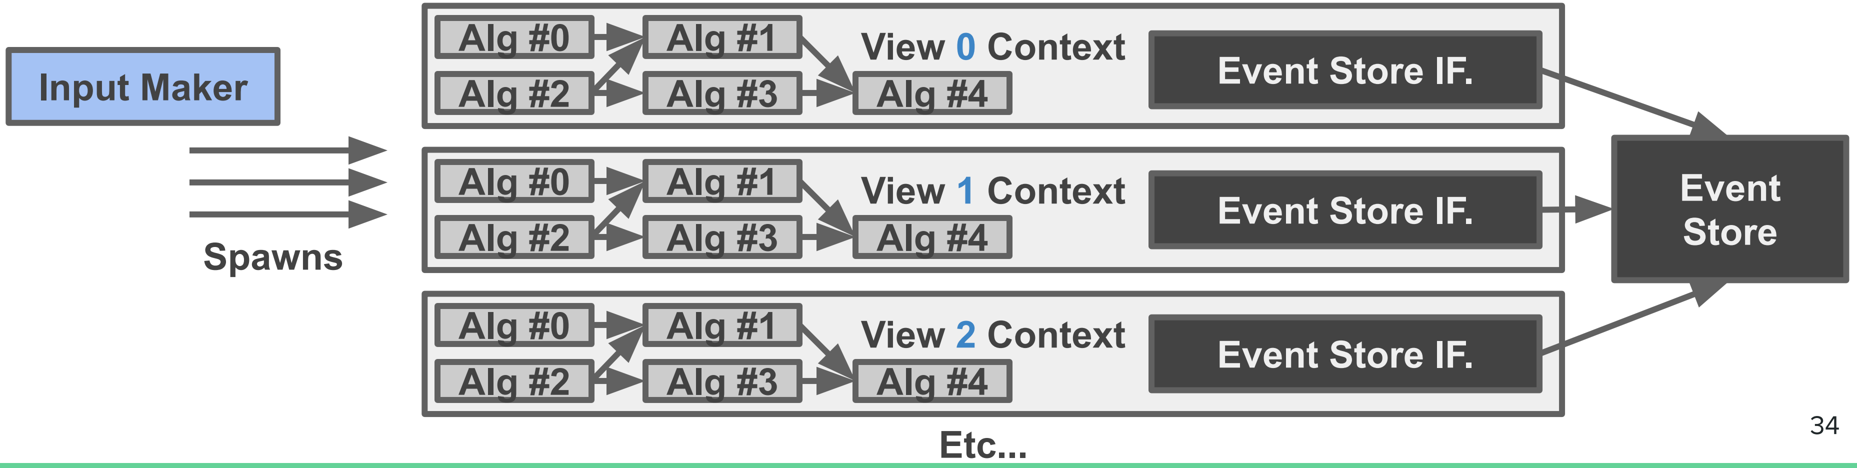

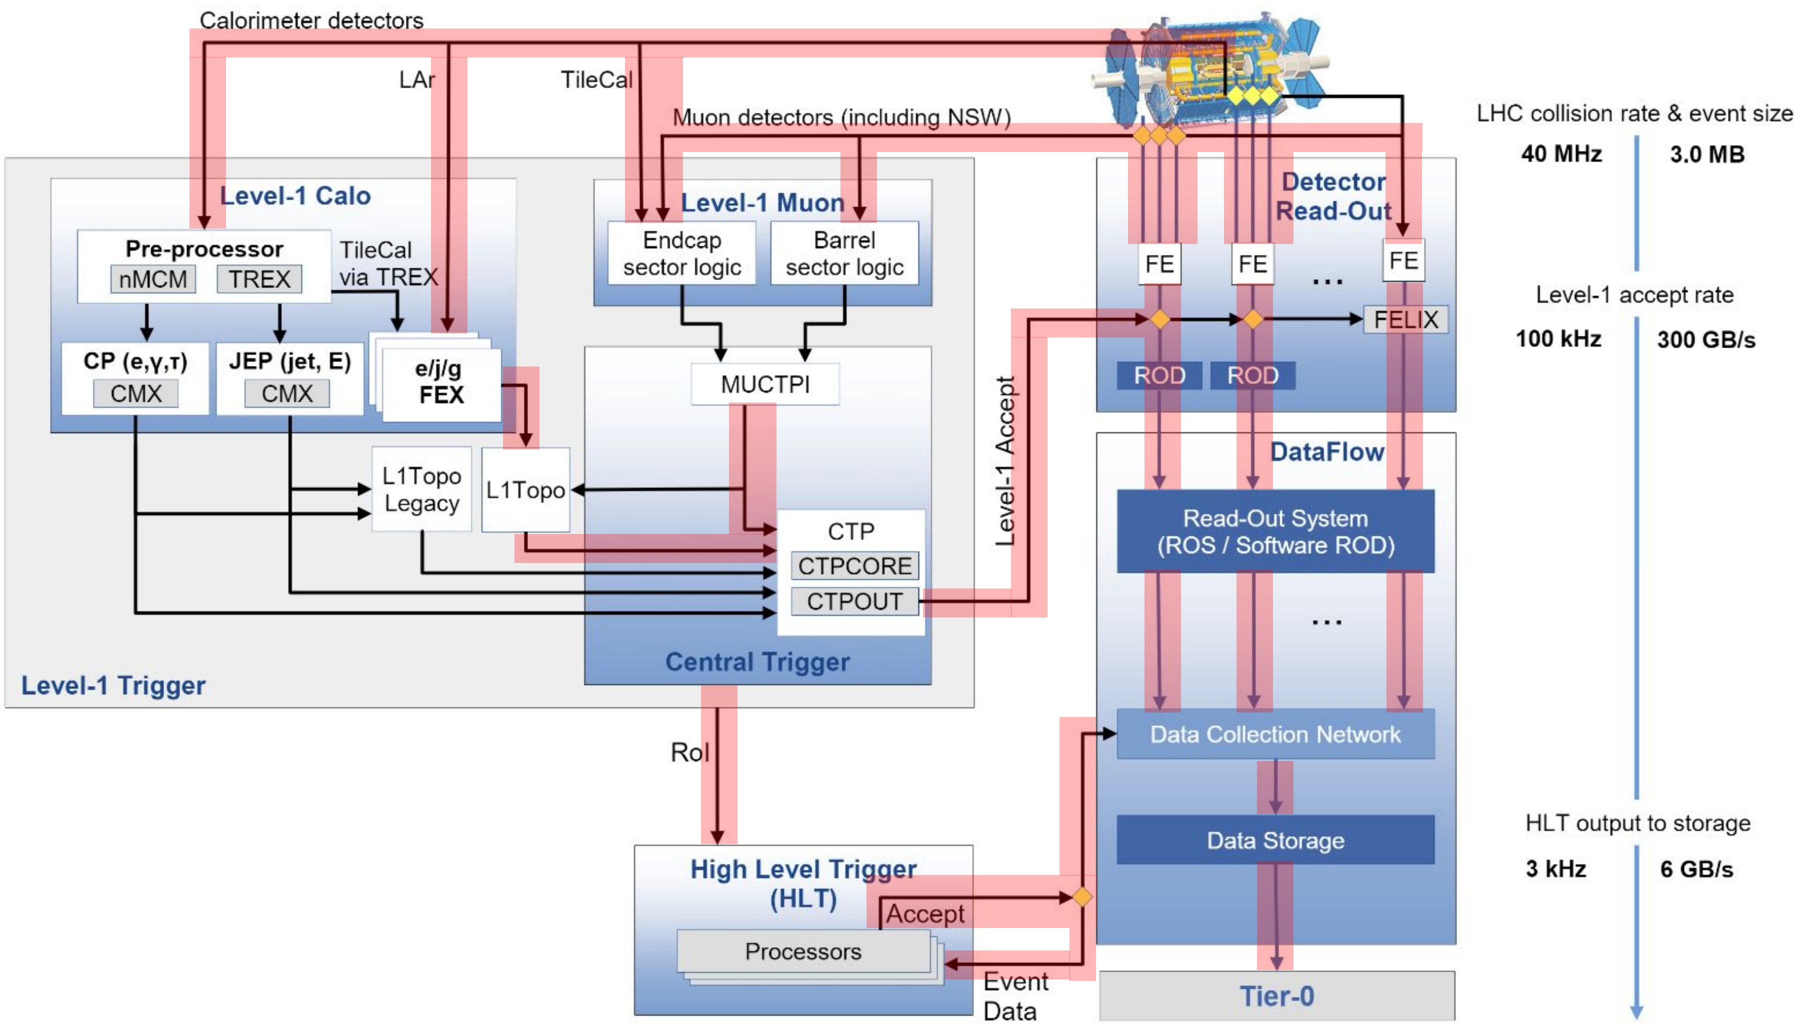

Calorimeter detectors
LAr
TileCal
Muon detectors (including NSW)
Level-1 Calo
Pre-processor
nMCM
TREX
TileCal via TREX
CP (e,γ,τ)
CMX
JEP (jet, E)
CMX
e/j/g FEX
L1Topo Legacy
L1Topo
Level-1 Muon
Endcap sector logic
Barrel sector logic
MUCTPI
CTP
CTPCORE
CTPOUT
Central Trigger
Level-1 Trigger
RoI
High Level Trigger (HLT)
Accept
Processors
Event Data
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
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

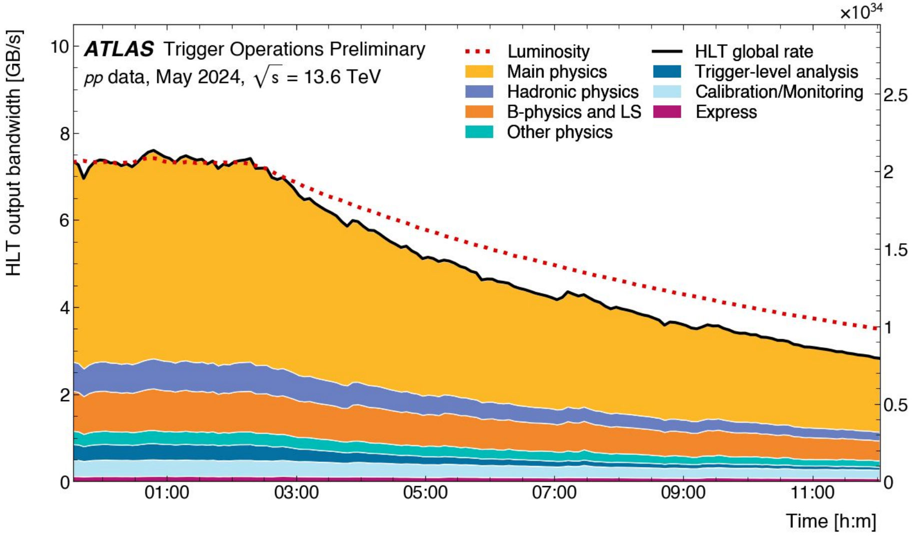

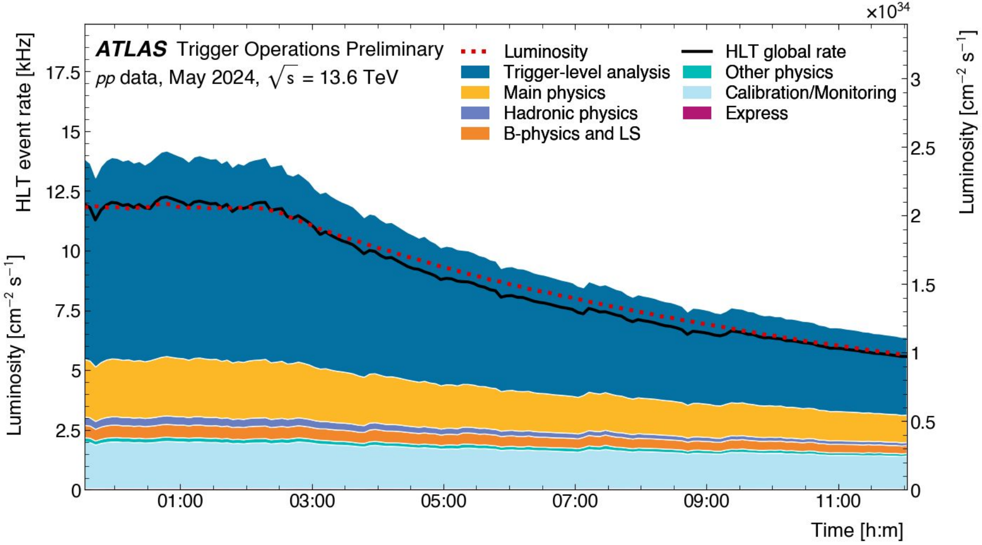

- ATLAS' HLT Trigger Menu records around **3 kHz full-events to disk** (x32 rejection).
- Record **7 GB/s**, including **extra events** for b-physics and b-jet analyses to be reconstructed when resources are available, **small events** with only the HLT result, and **hybrid approaches** too.

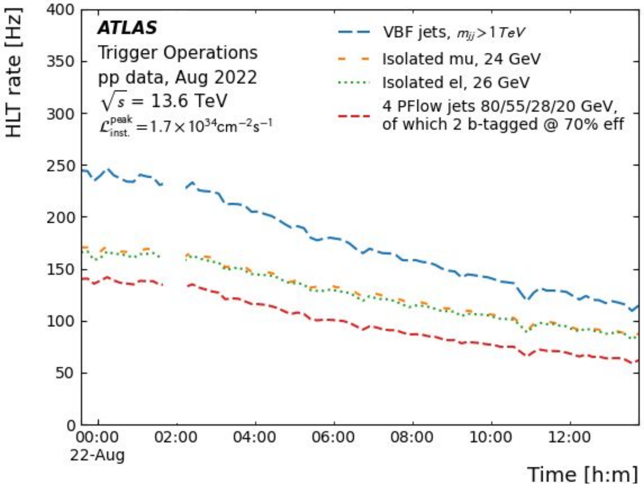
ATLAS
Trigger Operations
pp data, Aug 2022
√s = 13.6 TeV
$\mathcal{L}_{inst.}^{peak} = 1.7 \times 10^{34} cm^{-2}s^{-1}$
VBF jets, $m_{jj} > 1 TeV$
Isolated mu, 24 GeV
Isolated el, 26 GeV
4 PFlow jets 80/55/28/20 GeV, of which 2 b-tagged @ 70% eff
HLT rate [Hz]
400
350
300
250
200
150
100
50
0
00:00 22-Aug
02:00
04:00
06:00
08:00
10:00
12:00
Time [h:m]

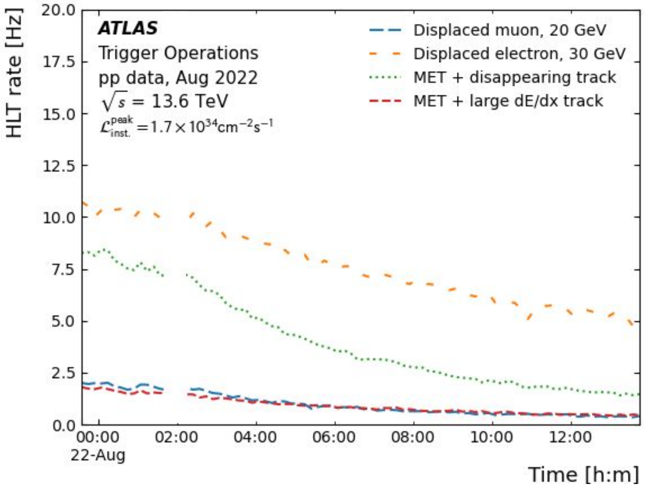
ATLAS
Trigger Operations
pp data, Aug 2022
√s = 13.6 TeV
$\mathcal{L}_{inst.}^{peak} = 1.7 \times 10^{34} cm^{-2}s^{-1}$
Displaced muon, 20 GeV
Displaced electron, 30 GeV
MET + disappearing track
MET + large dE/dx track
HLT rate [Hz]
20.0
17.5
15.0
12.5
10.0
7.5
5.0
2.5
0.0
00:00 22-Aug
02:00
04:00
06:00
08:00
10:00
12:00
Time [h:m]

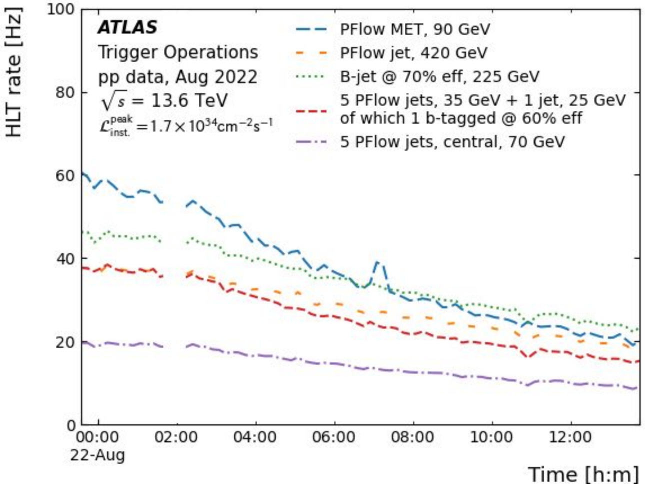
ATLAS
Trigger Operations
pp data, Aug 2022
√s = 13.6 TeV
$\mathcal{L}_{inst.}^{peak} = 1.7 \times 10^{34} cm^{-2}s^{-1}$
PFlow MET, 90 GeV
PFlow jet, 420 GeV
B-jet @ 70% eff, 225 GeV
5 PFlow jets, 35 GeV + 1 jet, 25 GeV of which 1 b-tagged @ 60% eff
5 PFlow jets, central, 70 GeV
HLT rate [Hz]
100
80
60
40
20
0
00:00 22-Aug
02:00
04:00
06:00
08:00
10:00
12:00
Time [h:m]

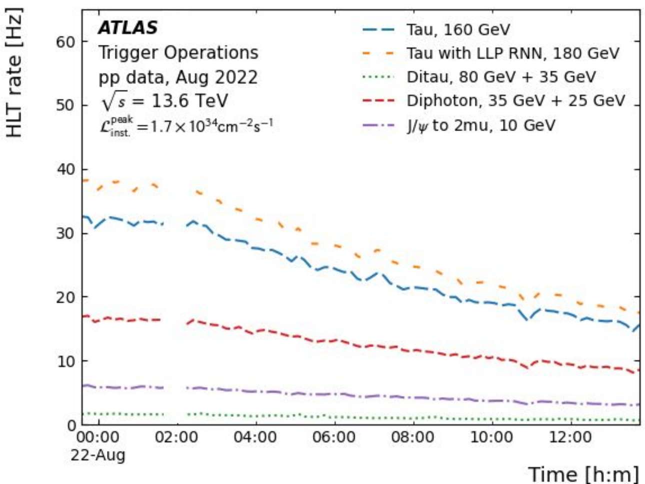
ATLAS
Trigger Operations
pp data, Aug 2022
√s = 13.6 TeV
$\mathcal{L}_{inst.}^{peak} = 1.7 \times 10^{34} cm^{-2}s^{-1}$
Tau, 160 GeV
Tau with LLP RNN, 180 GeV
Ditau, 80 GeV + 35 GeV
Diphoton, 35 GeV + 25 GeV
J/ψ to 2mu, 10 GeV
HLT rate [Hz]
60
50
40
30
20
10
0
00:00 22-Aug
02:00
04:00
06:00
08:00
10:00
12:00
Time [h:m]

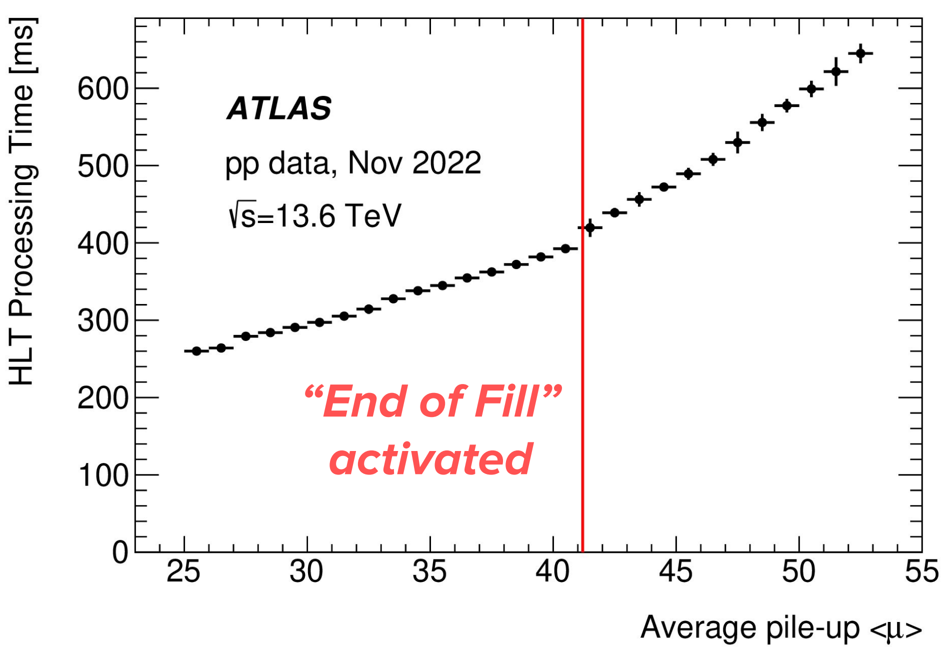

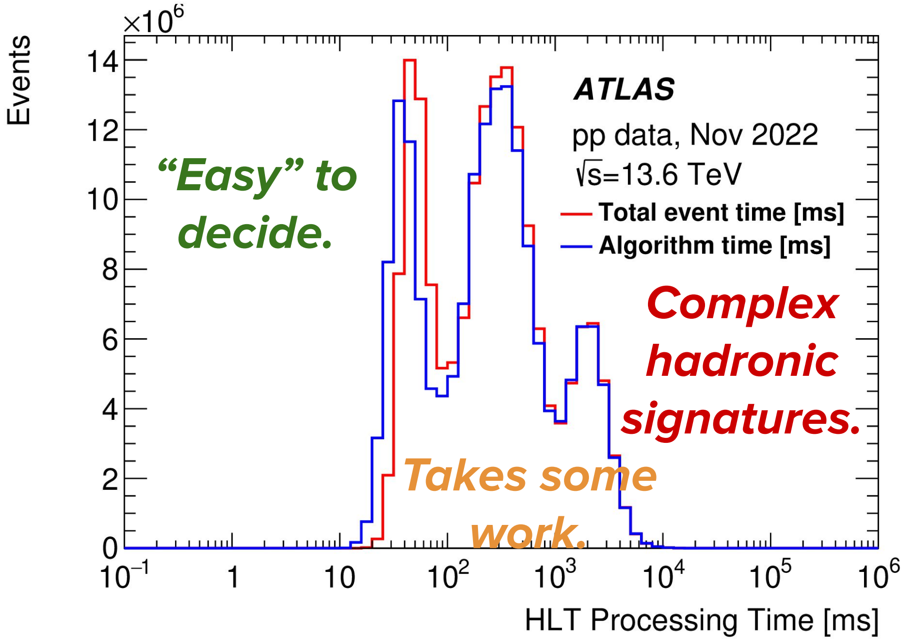

- Dual-CPU machines, **128 GB RAM**, **2x AMD EPYC 7302**. Total of **32 physical** and **64 physical+virtual** (hyper-threading) threads per machine.
- Total of around **120,000 threads** in the farm.

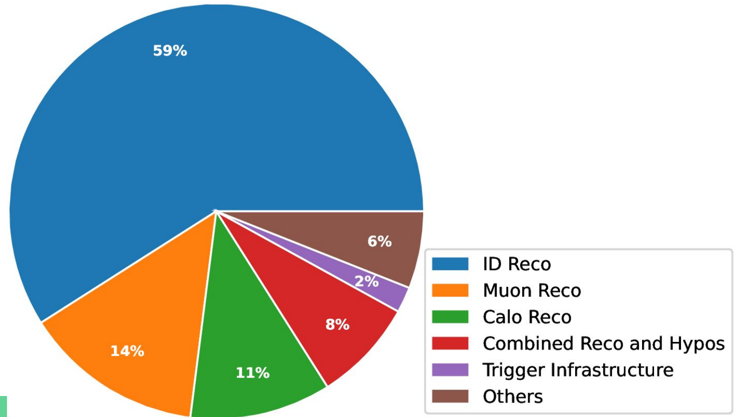

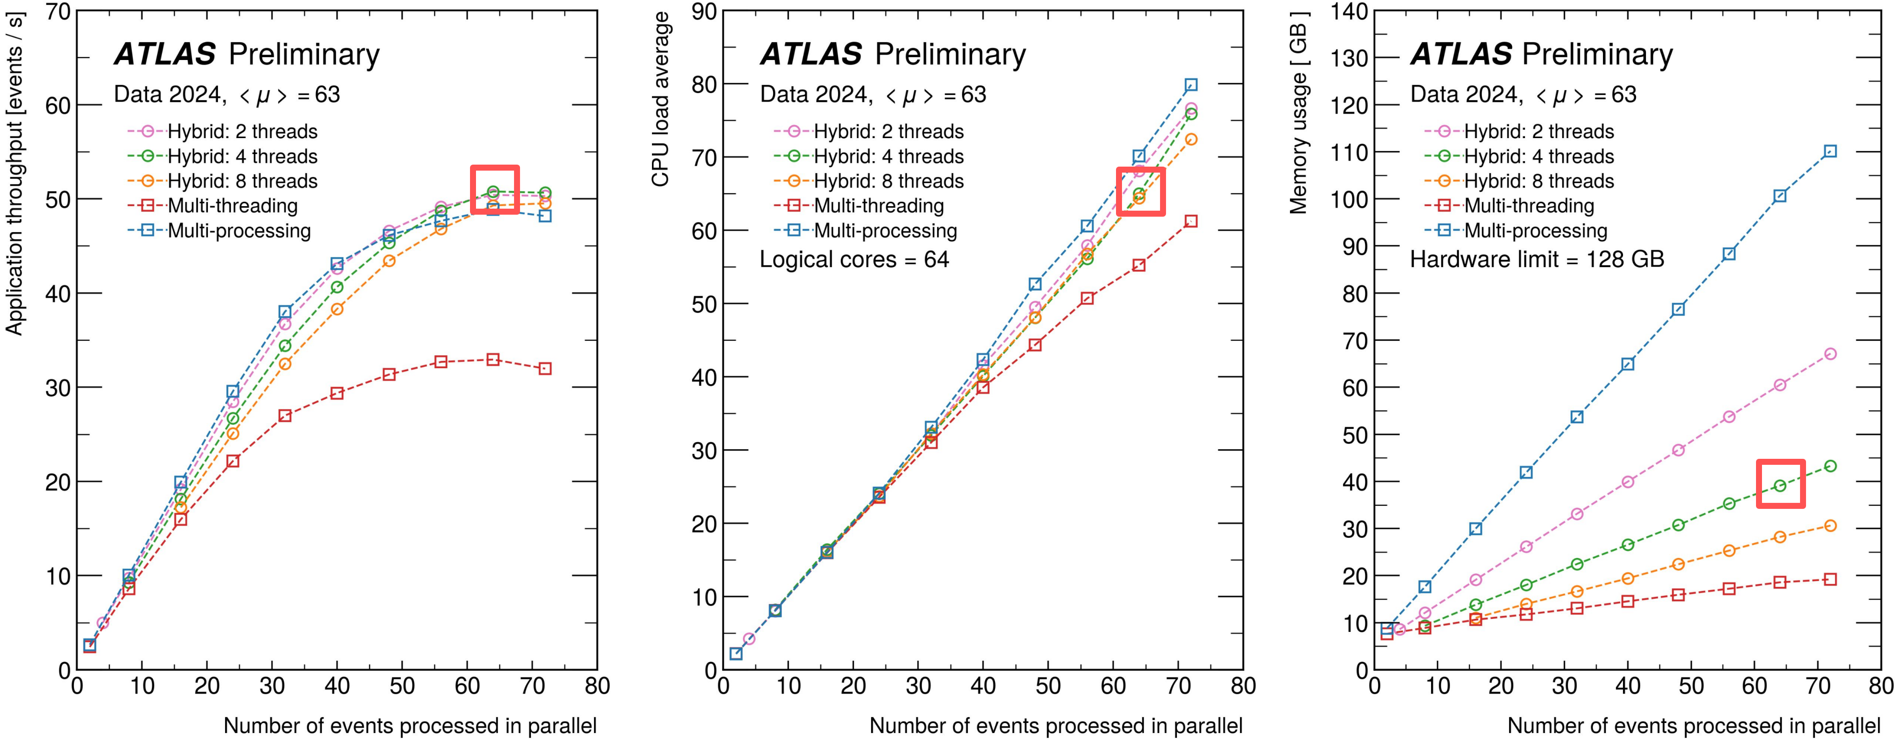

- With our MT-system we can fine tune the number of ***forks*** (multi-process) and the number of **threads** *(= concurrent events)* (multi-threading) used by each process.
- We currently **maximise throughput** by forking our process **16 times**, each processing **4 concurrent events** over **4 threads**.

To The Future! (Phase-II)

- **L1** becomes **L0**, **HLT** becomes **EF**.
- **Front End Link eXchange (FELIX)** readout is rolled out to all sub detectors.
- New **time multiplexed Global FPGA** trigger at L0 with access to calorimeter **cells**.
- **x10** and **x~3** increases in **L0** and **EF** output bandwidth.
- EF processing nodes have access to **all event data fragments**.
- EF farm may integrate **GPU or FPGA based accelerators**.
- ***The UK and the RAL-ATLAS group have long histories of involvement in the ATLAS trigger and the Phase-II upgrade.***