ATLAS Trigger and Data Acquisition in the High Luminosity LHC era

Riccardo Scaglioni (University and INFN Pavia)

Supervisors:
Andrea Negri (University and INFN Pavia)
Gabriella Gaudio (INFN Pavia)

What is Data Acquisition?

Every experiment that uses electronic detectors needs a **Data Acquisition** system (**DAQ**).

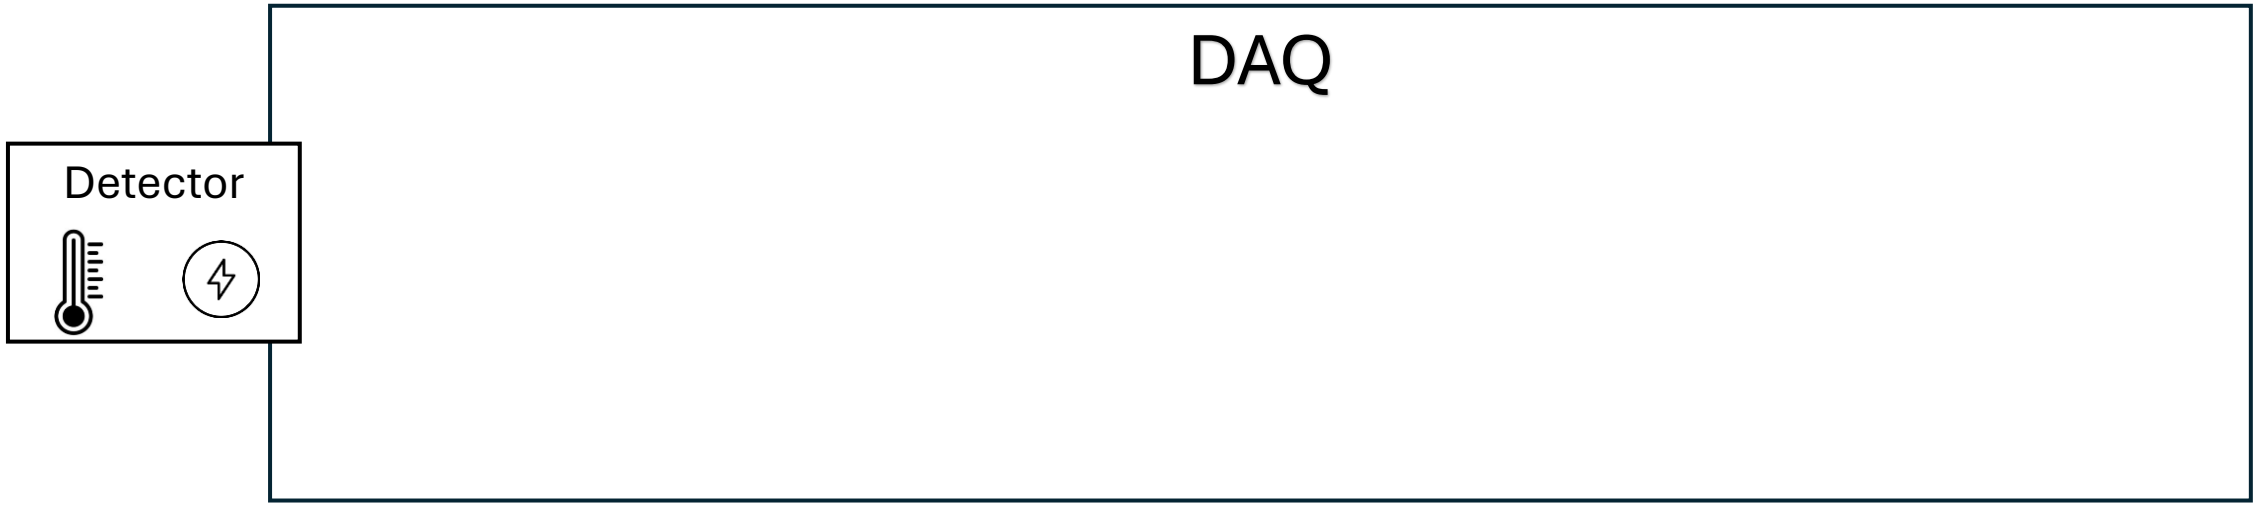

Physic Quantity

What is Data Acquisition?

Every experiment that uses electronic detectors needs a **Data Acquisition** system (**DAQ**).

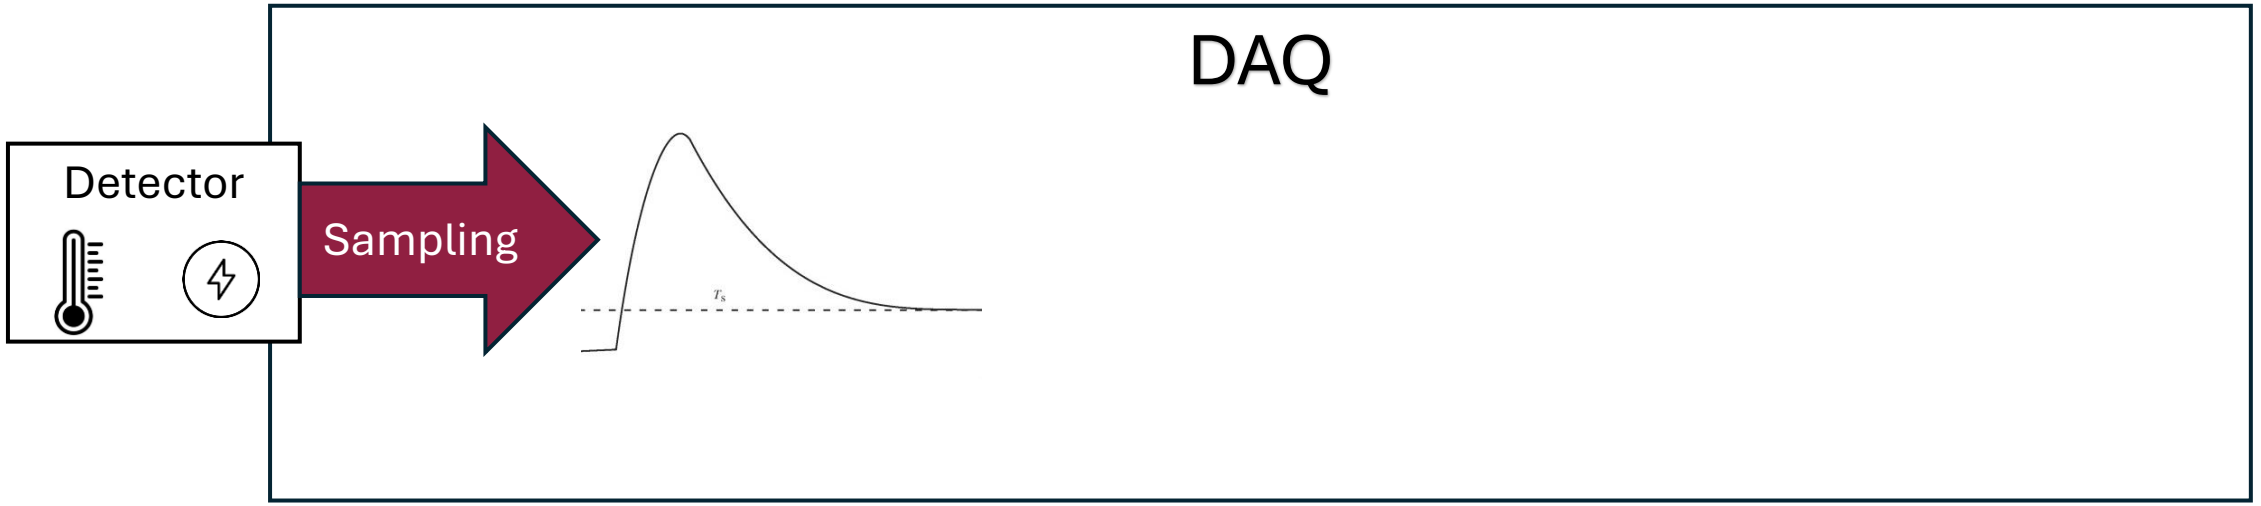

Physic Quantity

Electric Signals

What is Data Acquisition?

Every experiment that uses electronic detectors needs a **Data Acquisition** system (**DAQ**).

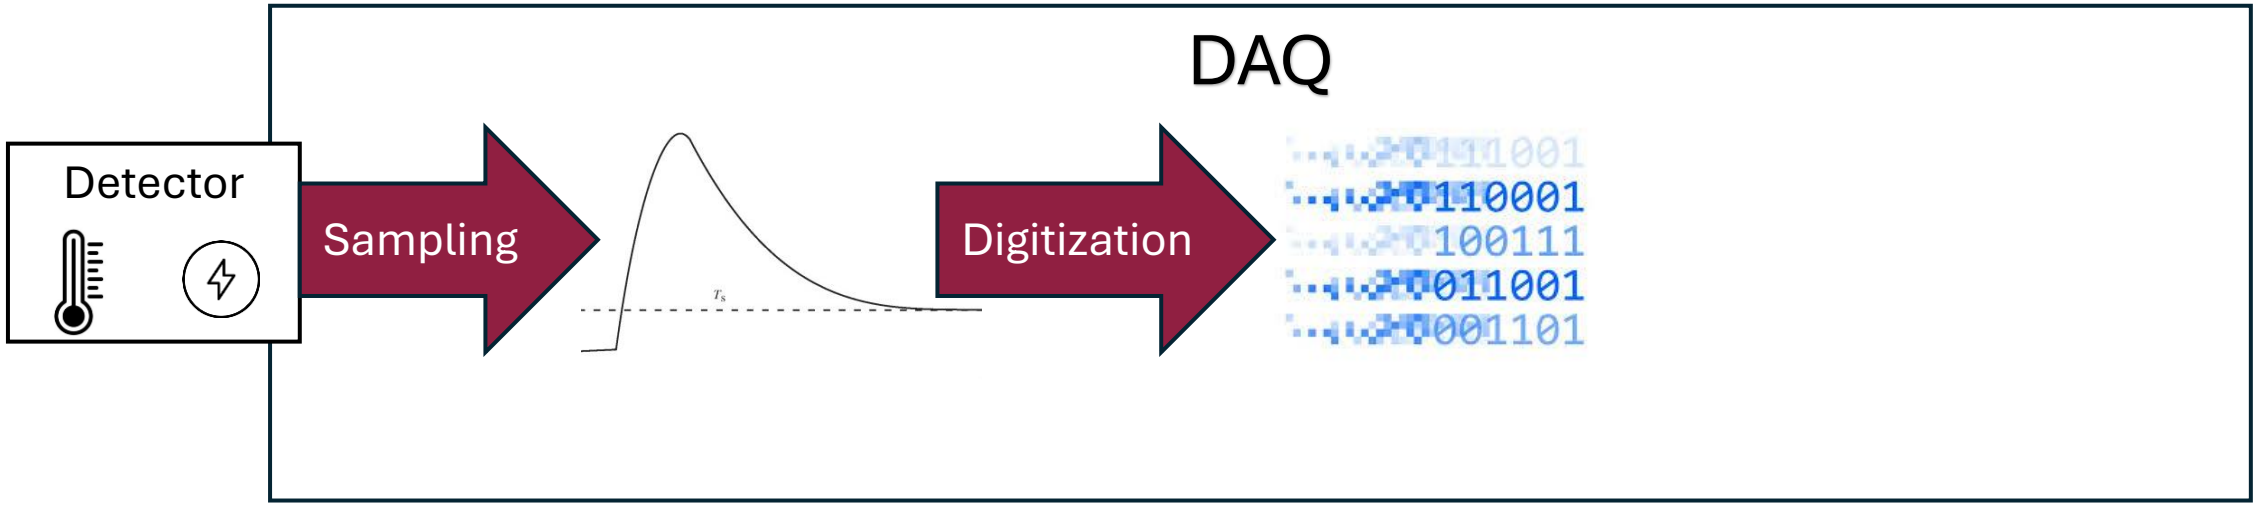

Physic Quantity

Electric Signals

Digital values

What is Data Acquisition?

Every experiment that uses electronic detectors needs a **Data Acquisition** system (**DAQ**).

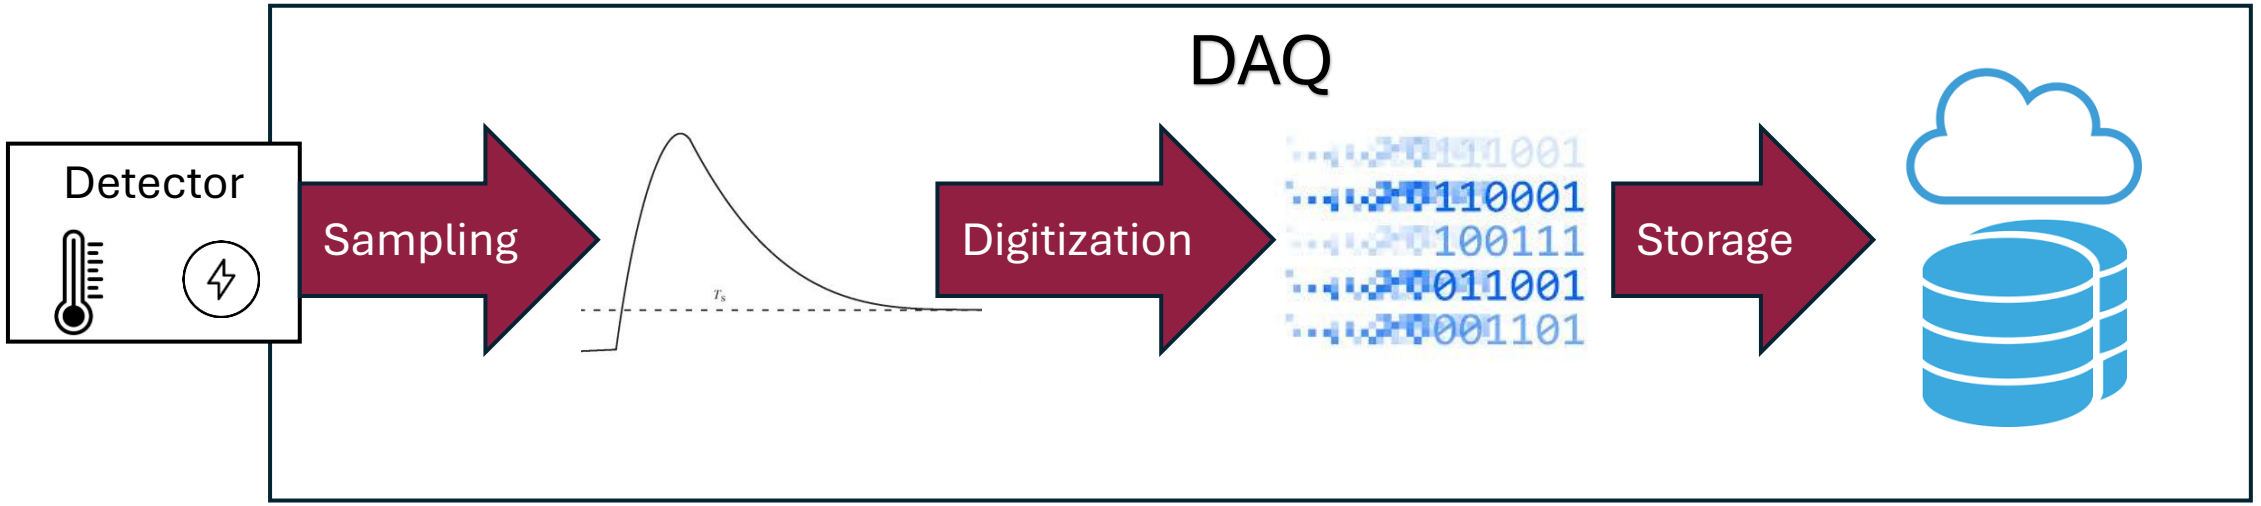

Physic Quantity

Electric Signals

Digital values

Persistent Storage

Physics at the Large Hadron Collider

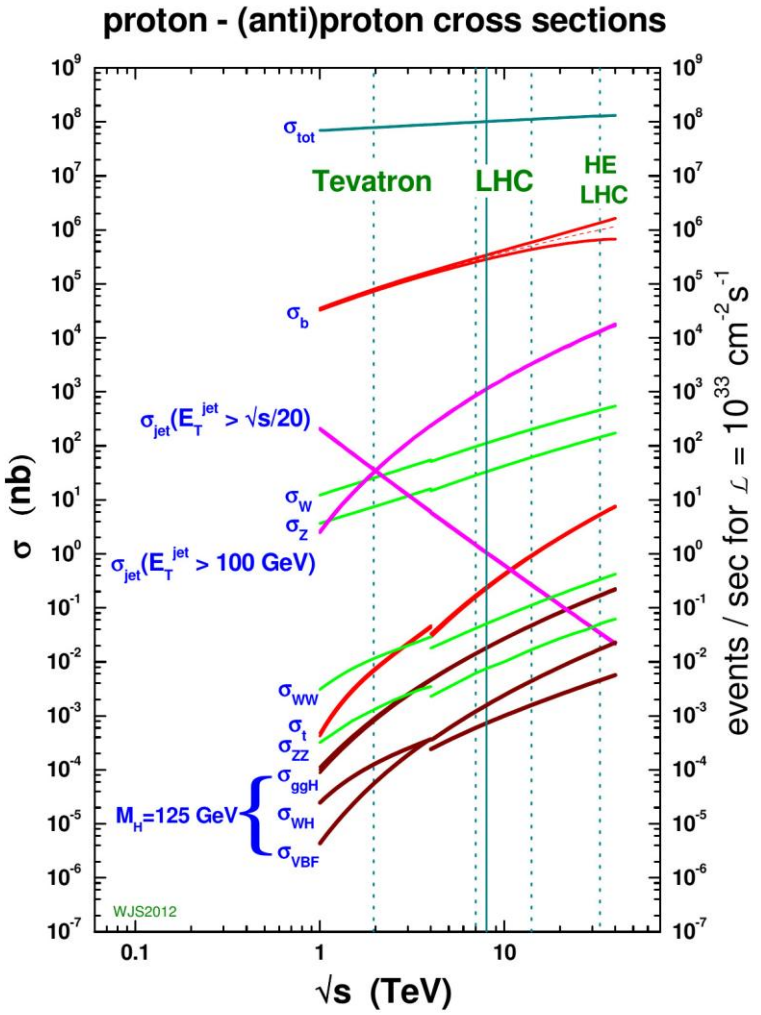

- Center of mass energy: $\sqrt{s}$ = 13.6 TeV
- High collision frequency: 40 MHz

Physics at the Large Hadron Collider

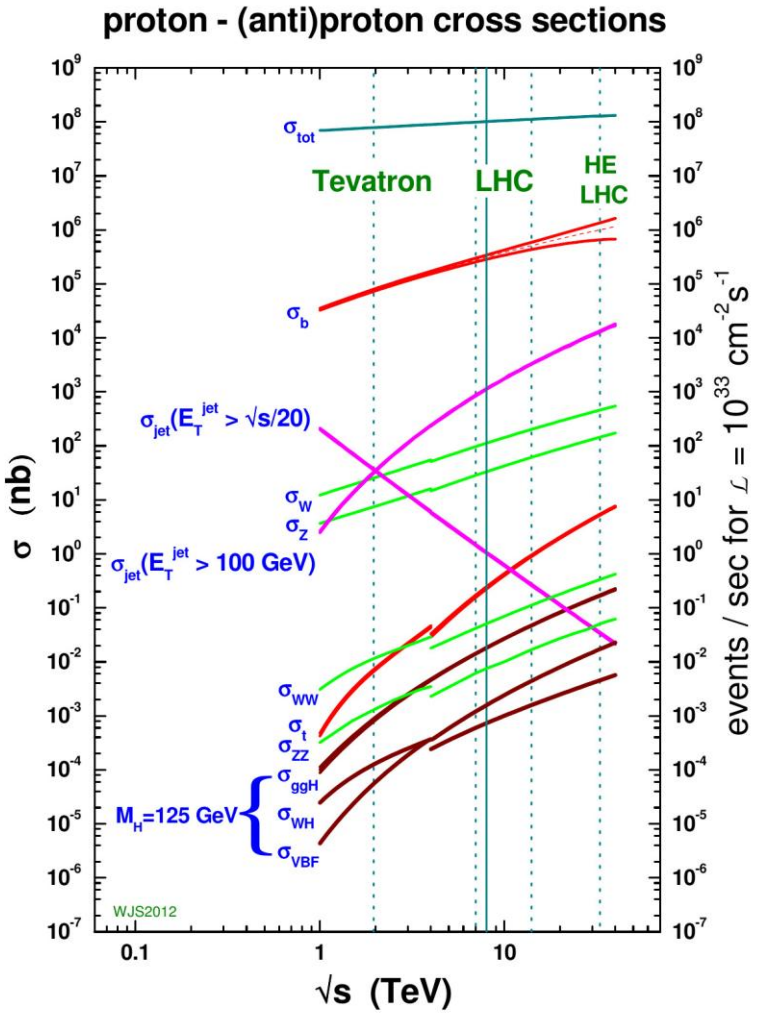

- Center of mass energy: $\sqrt{s}$ = 13.6 TeV
- High collision frequency: 40 MHz
- Luminosity: 2×10^{34} cm^{-2} s^{-1}
- Pile-up (max): 60 superposed events

Physics at the Large Hadron Collider

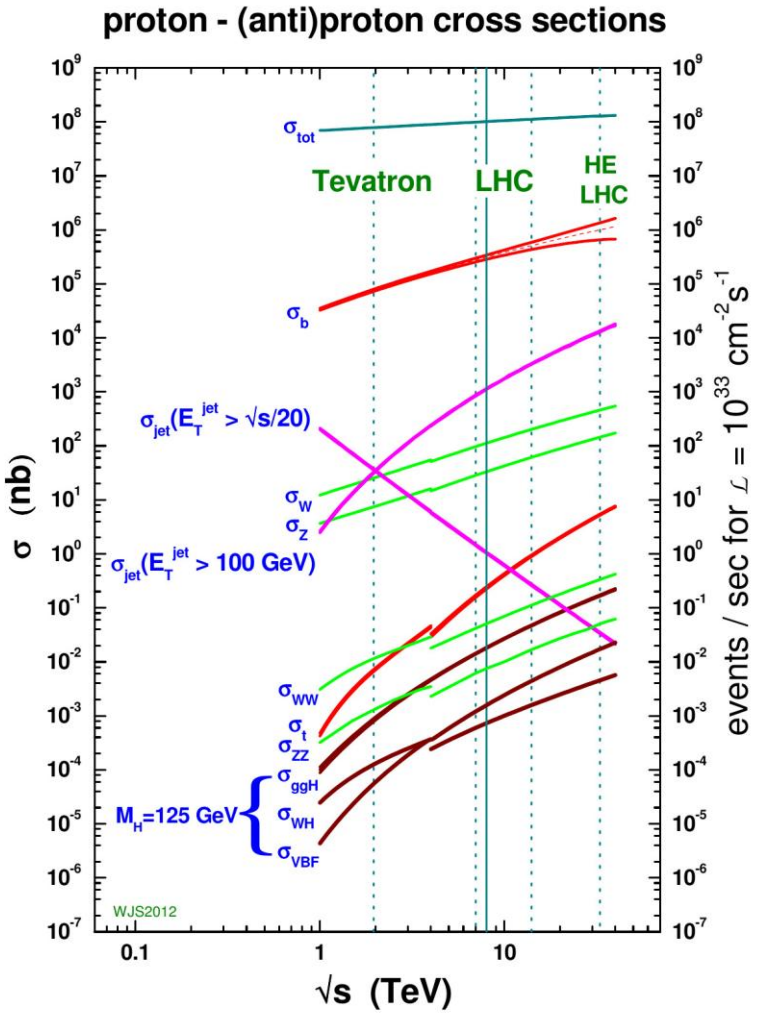

- Center of mass energy: $\sqrt{s}$ = 13.6 TeV
- High collision frequency: 40 MHz
- Luminosity: 2×10^{34} cm^{-2} s^{-1}
- Pile-up (max): 60 superposed events
- Typical event size ~ 1,5 MB
- Lot of collisions data produced: > 60 TB/s

Physics at the Large Hadron Collider

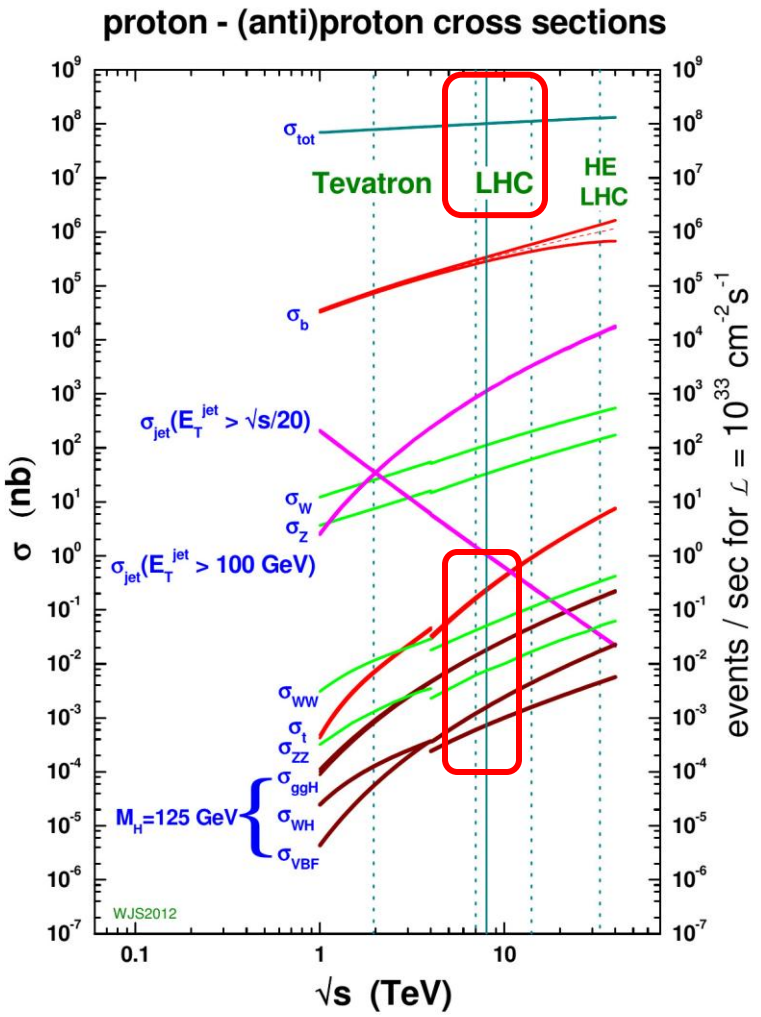

- Center of mass energy: $\sqrt{s}$ = 13.6 TeV
- High collision frequency: 40 MHz
- Luminosity: 2×10^{34} cm^{-2} s^{-1}
- Pile-up (max): 60 superposed events
- Typical event size ~ 1,5 MB
- Lot of collisions data produced: > 60 TB/s
- **Very rare interesting events (Higgs ~ 10^{-10} σ_{tot})**

Physics at the Large Hadron Collider

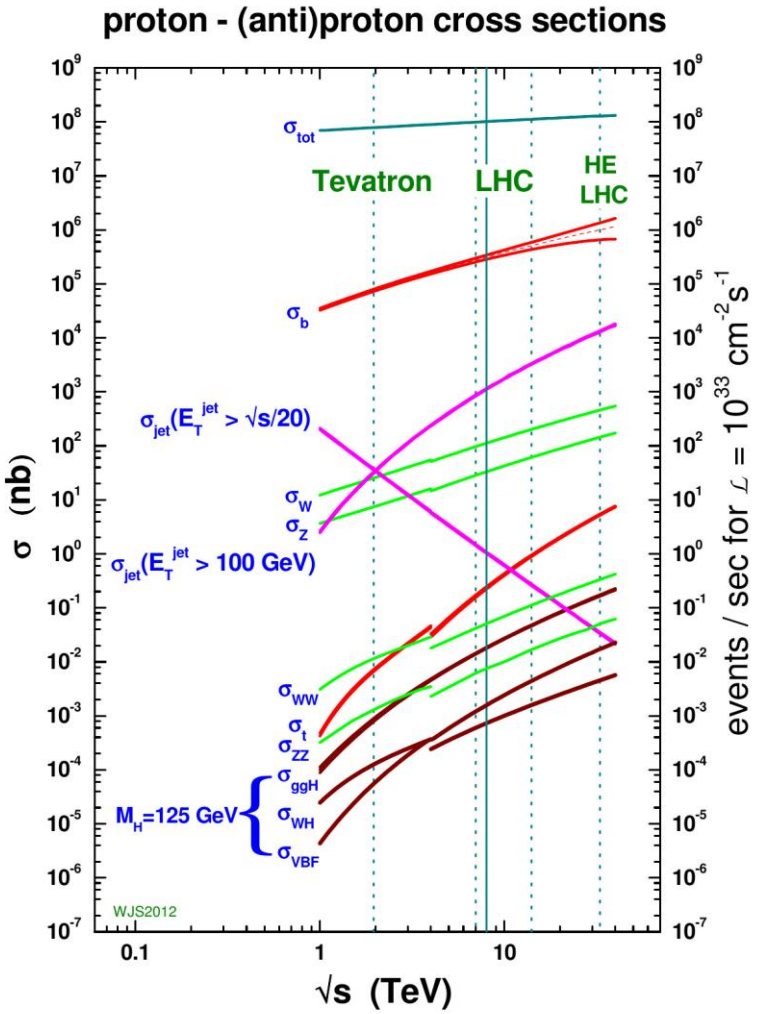

What do we with all that data?

Physics at the Large Hadron Collider

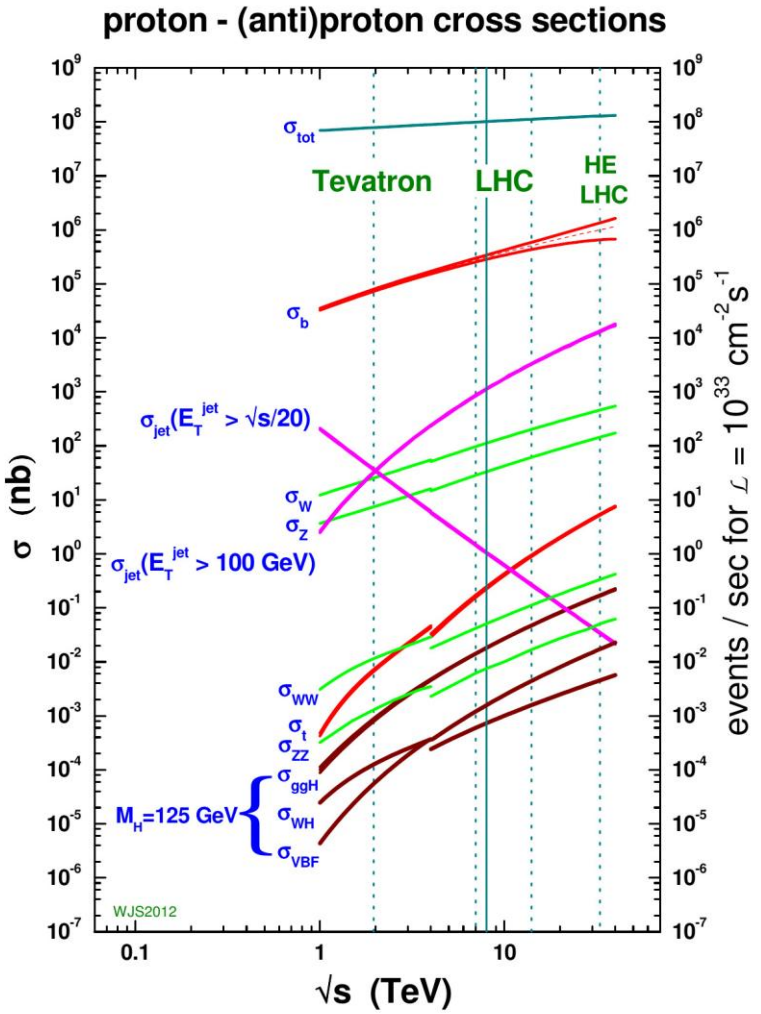

What do we with all that data?

- **Ideally:** It's all physics, KEEP IT!

Physics at the Large Hadron Collider

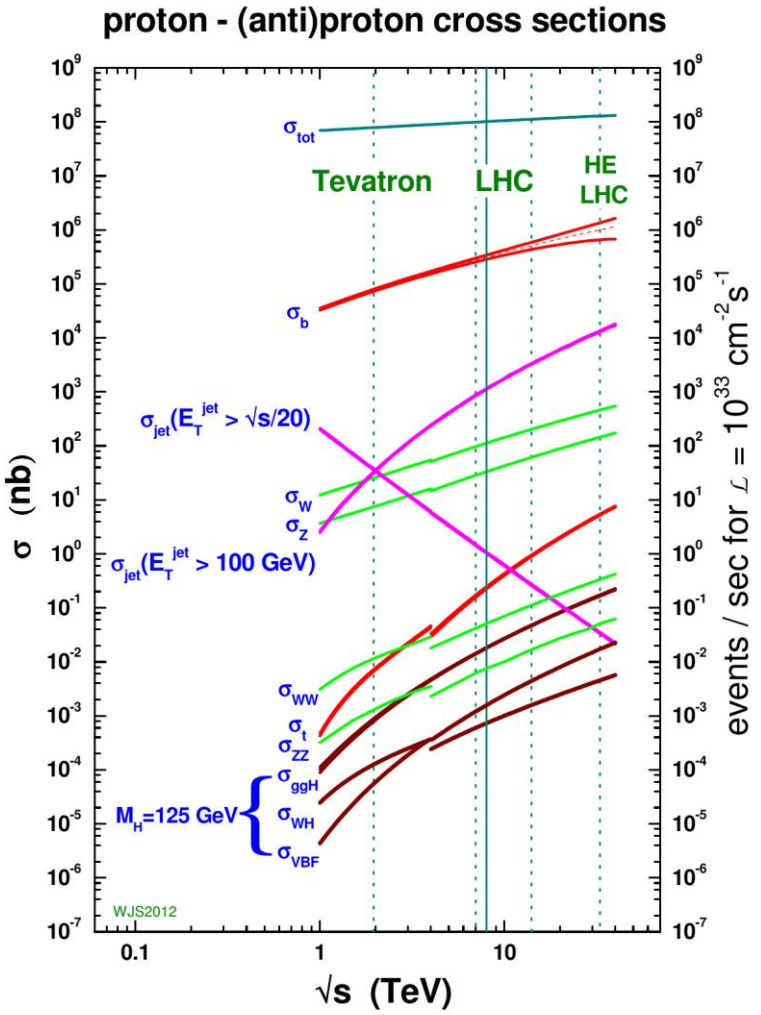

What do we with all that data?

- **Ideally:** It's all physics, KEEP IT!
- **Technically: 2 · 10^9 TB/y → ~ 1 G€/y**

Physics at the Large Hadron Collider

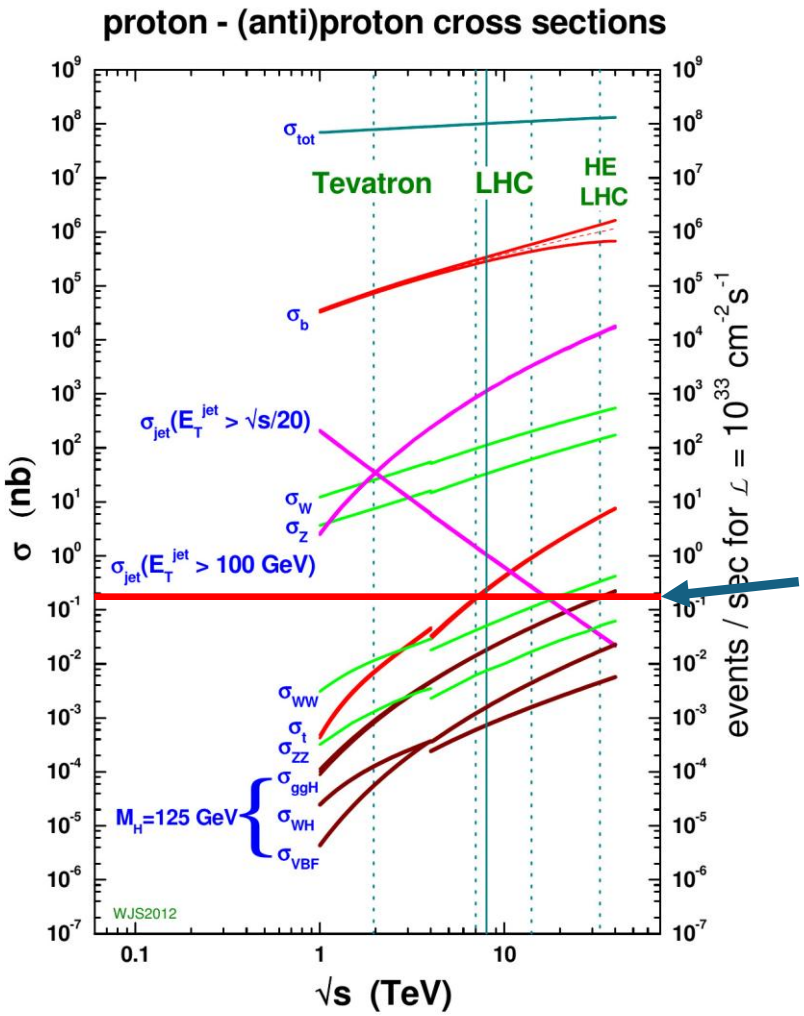

What do we with all that data?

- **Ideally:** It's all physics, KEEP IT!
- **Technically: 2 · 10^9 TB/y → ~ 1 G€/y**

- **Actually:** much lower write rate, 1,5 GB/s

Physics at the Large Hadron Collider

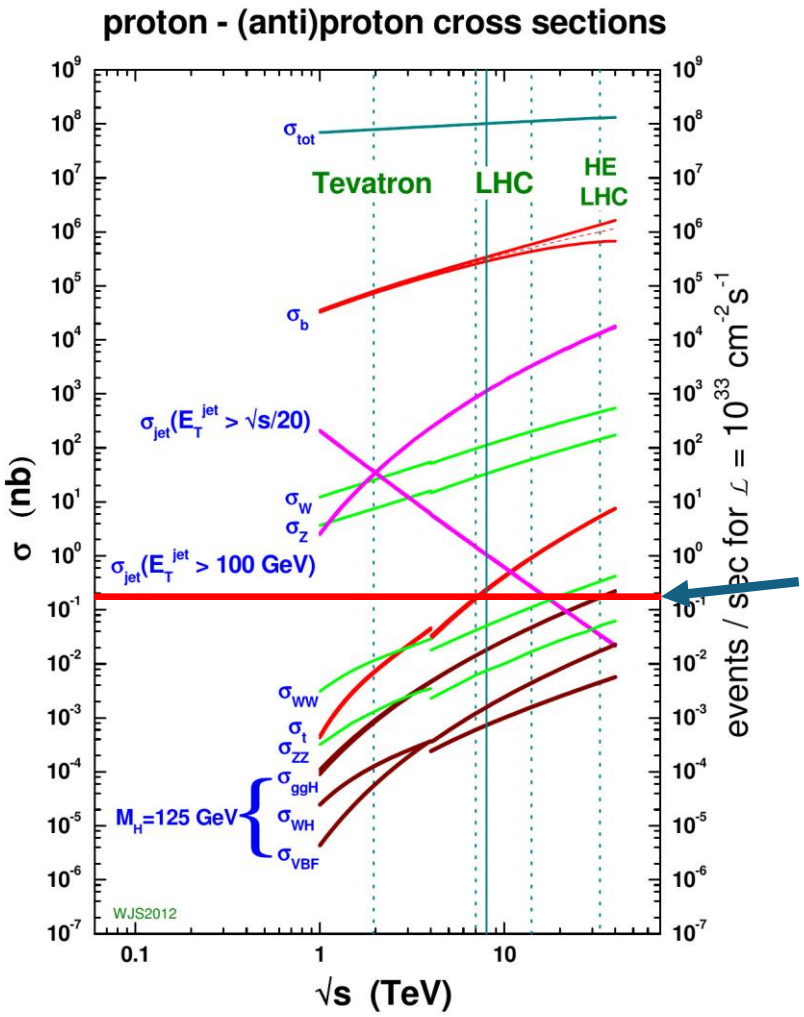

What do we with all that data?

- **Ideally:** It's all physics, KEEP IT!
- **Technically: 2 · 10^9 TB/y → ~ 1 G€/y**

- **Actually:** much lower write rate, 1,5 GB/s
- Filling one **smartphone** every two minutes

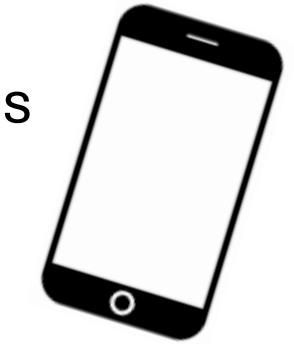

DAQ at the Large Hadron Collider

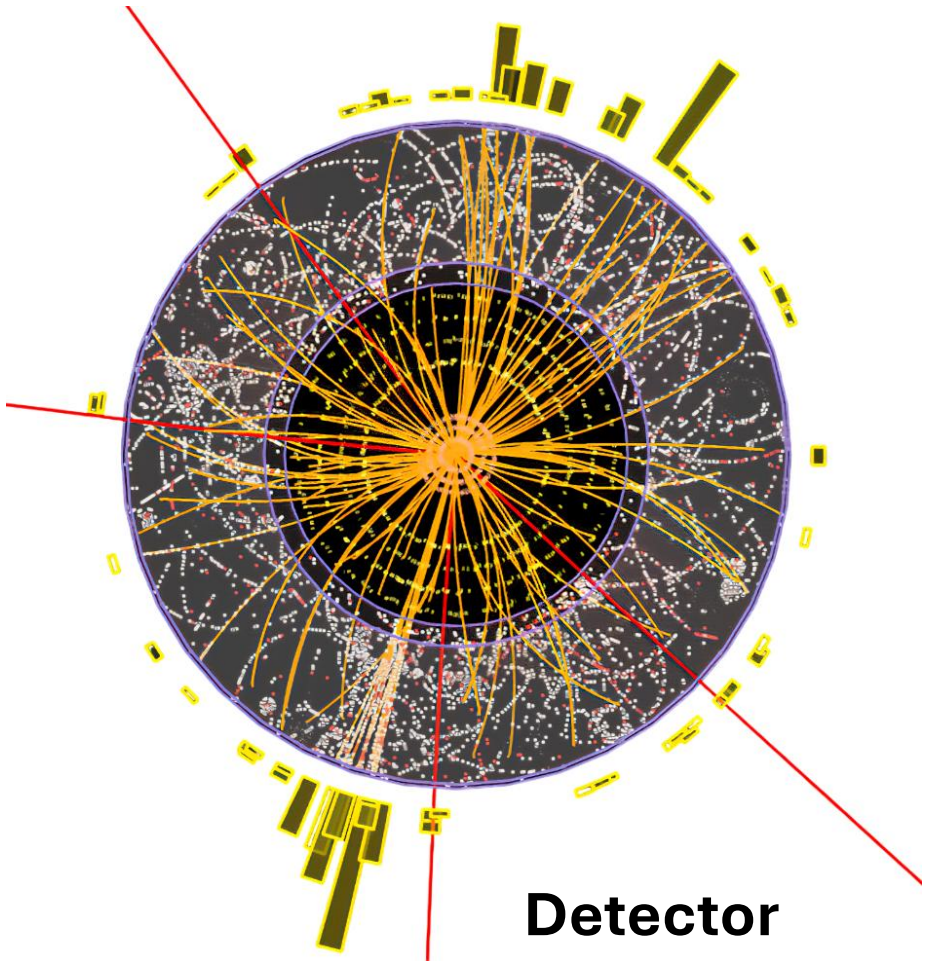

DAQ at the Large Hadron Collider

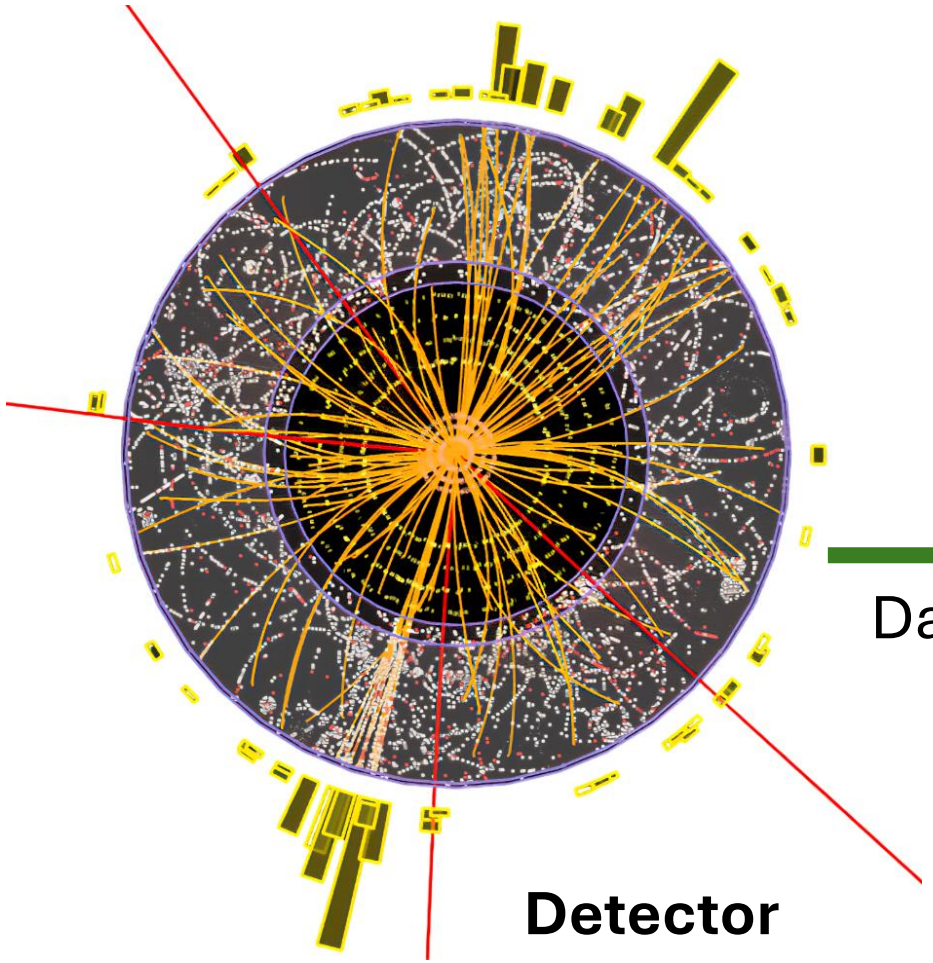

Detector

Data

DAQ

Storage

- Reconstruction
- Analysis
- Physics results

DAQ at the Large Hadron Collider

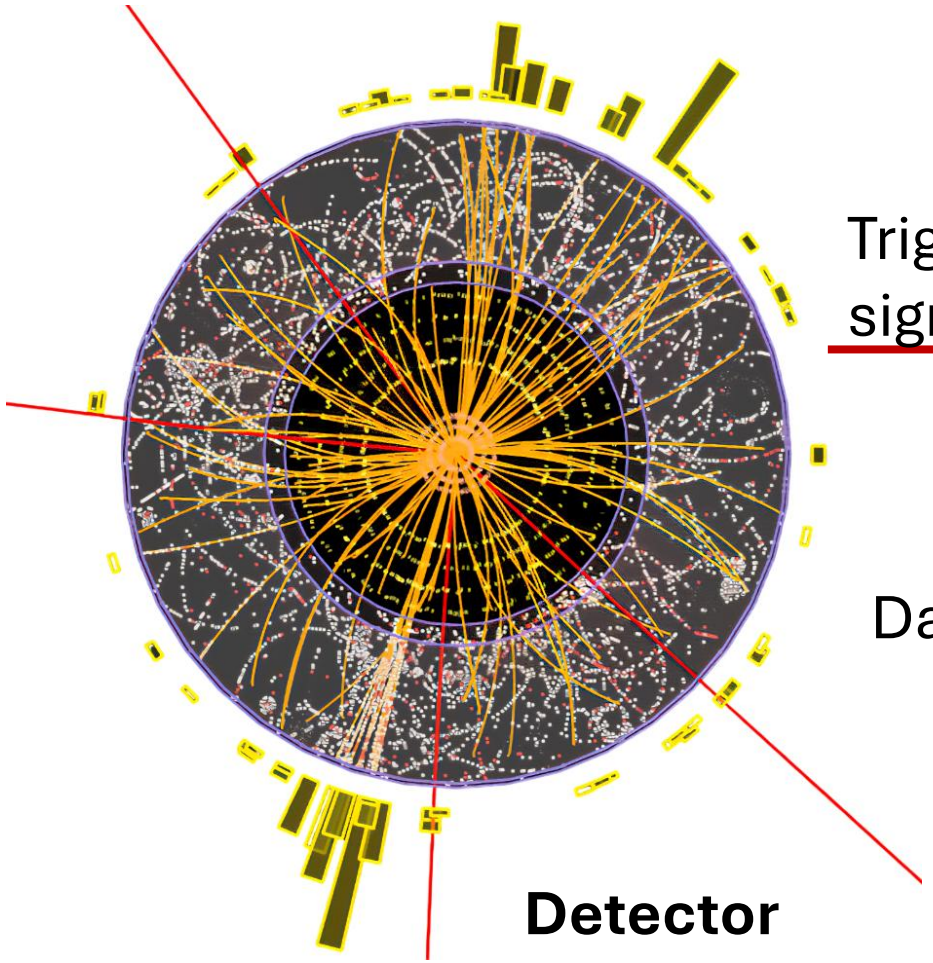

Detector

Trigger signals

Trigger

Data

DAQ

Storage

- Reconstruction
- Analysis
- Physics results

DAQ at the Large Hadron Collider

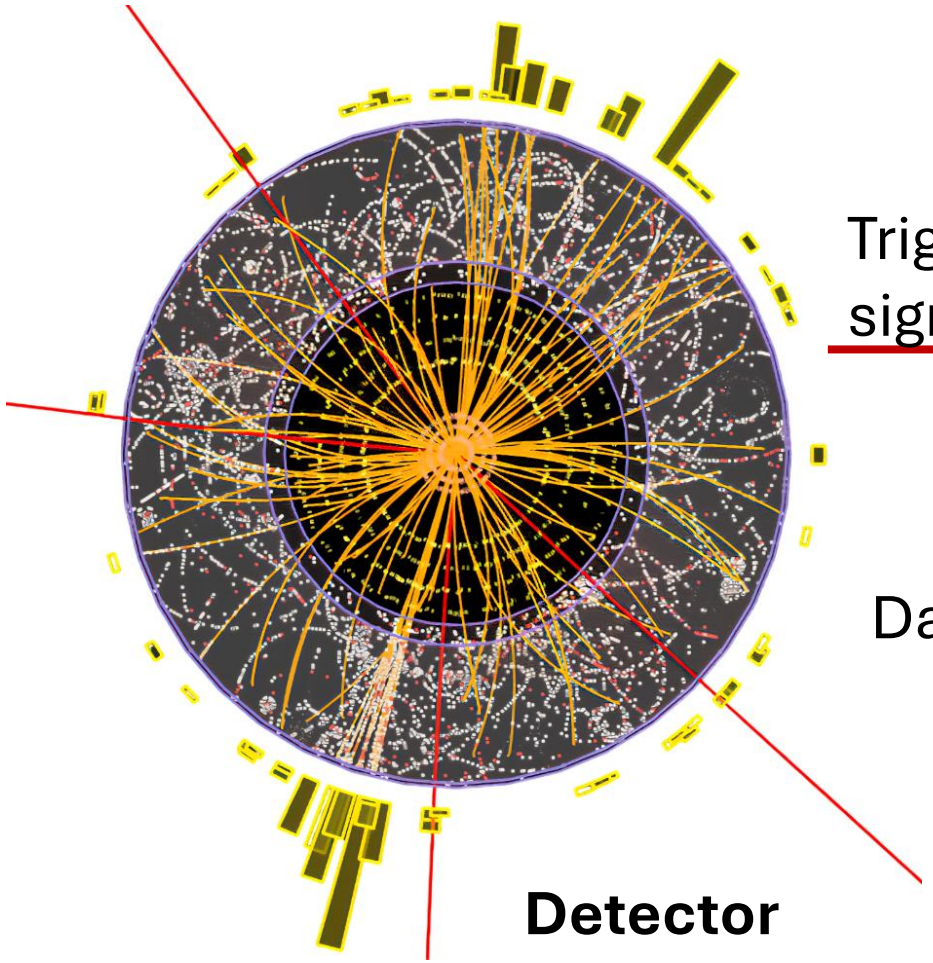

Detector

Trigger signals

Data

Trigger

Trigger Decision

DAQ

Storage

- Reconstruction
- Analysis
- Physics results

DAQ at the Large Hadron Collider

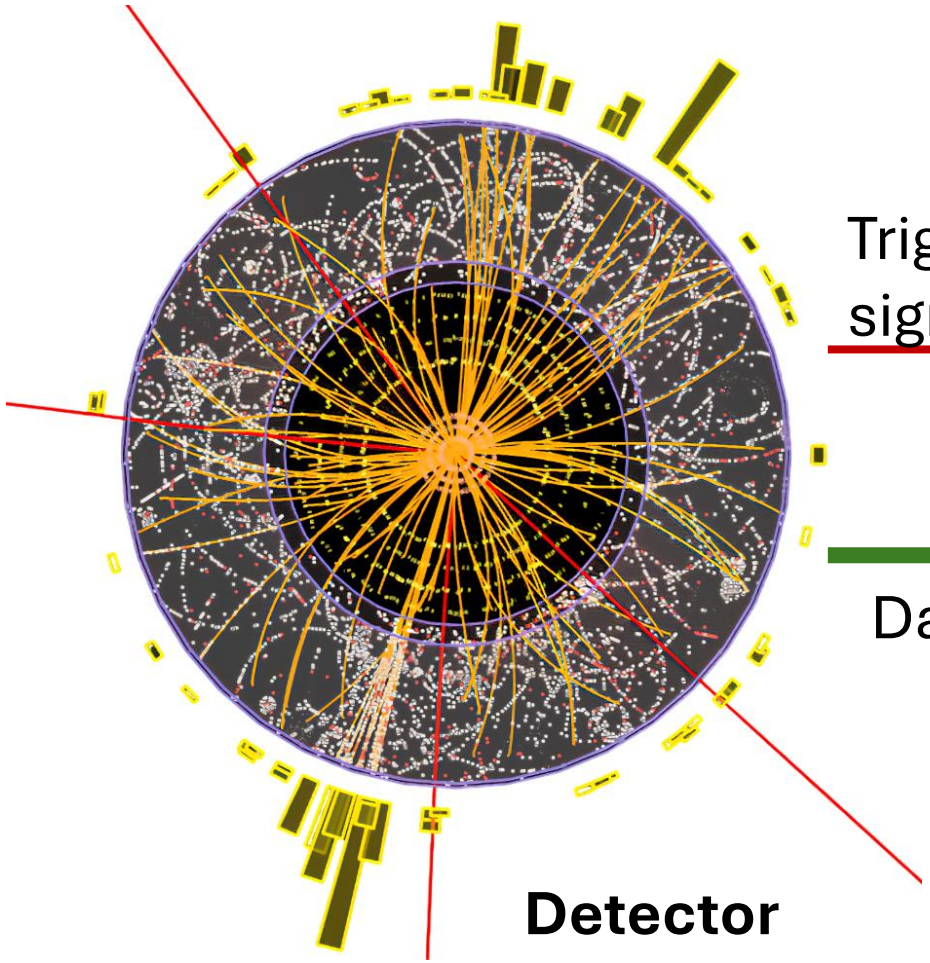

Detector

Trigger signals

Trigger

Trigger Decision

DAQ

Storage

Data

- Reconstruction
- Analysis
- Physics results

DAQ at the Large Hadron Collider

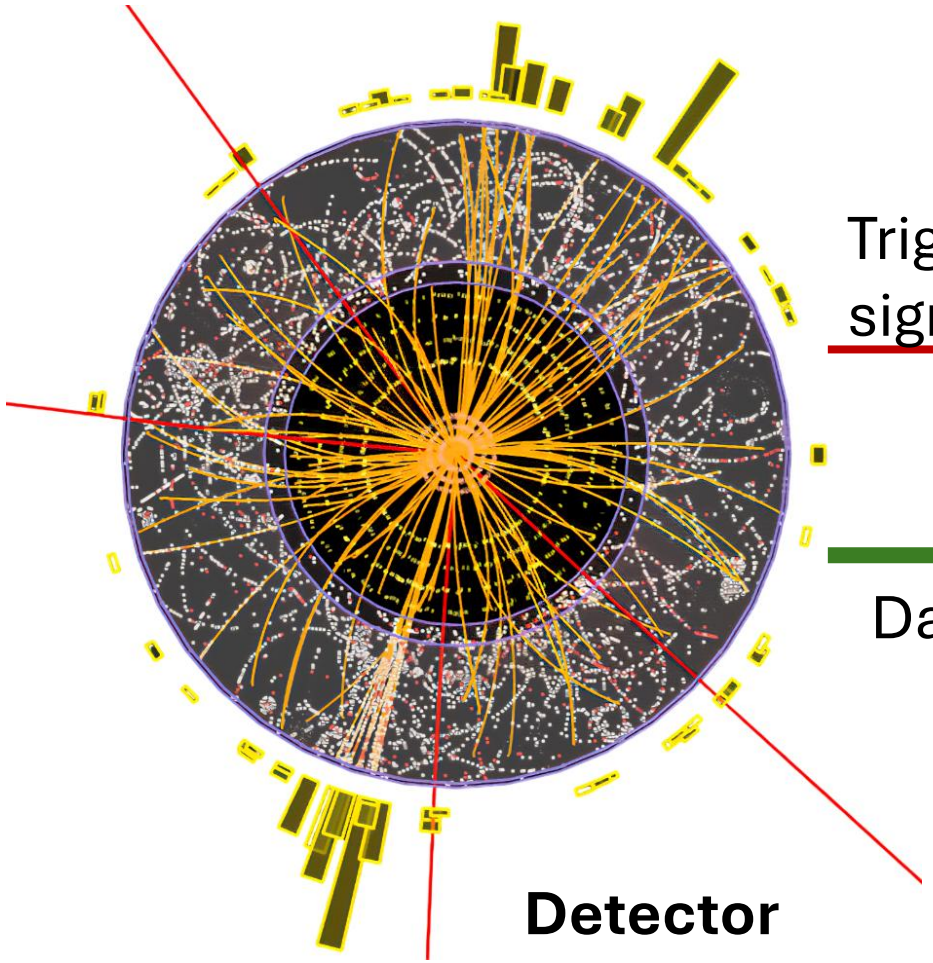

Detector

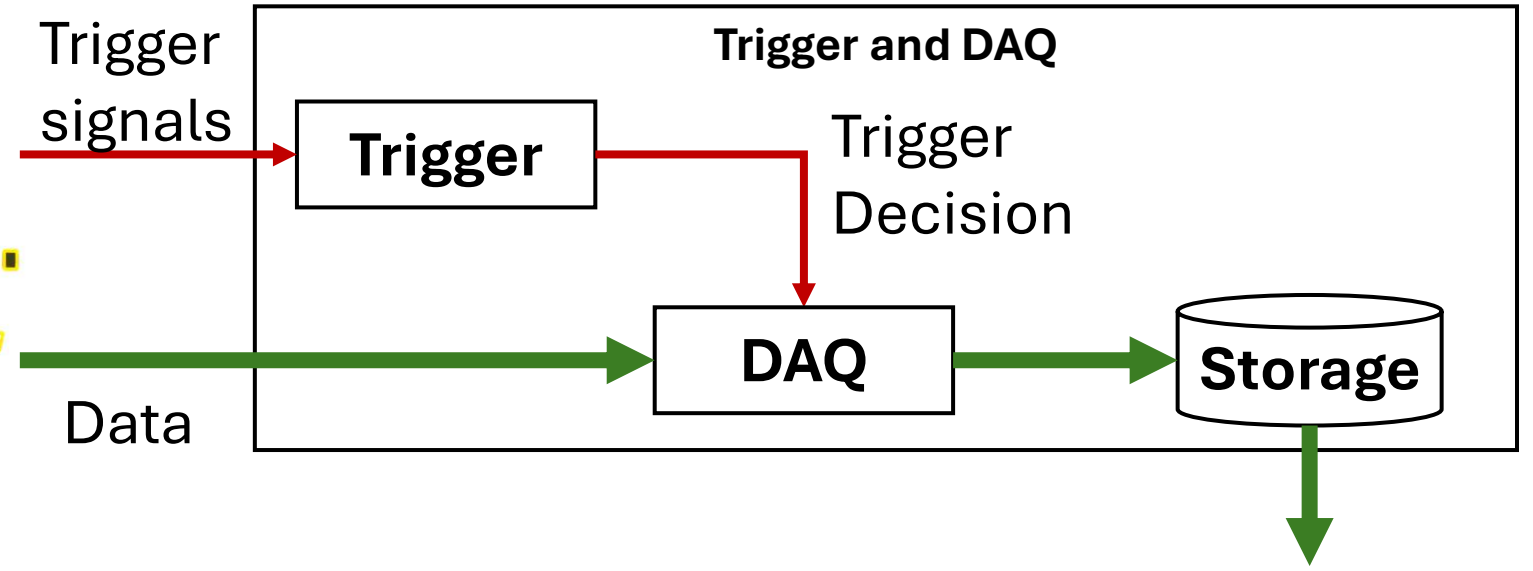

- Reconstruction
- Analysis
- Physics results

The ATLAS experiment TDAQ

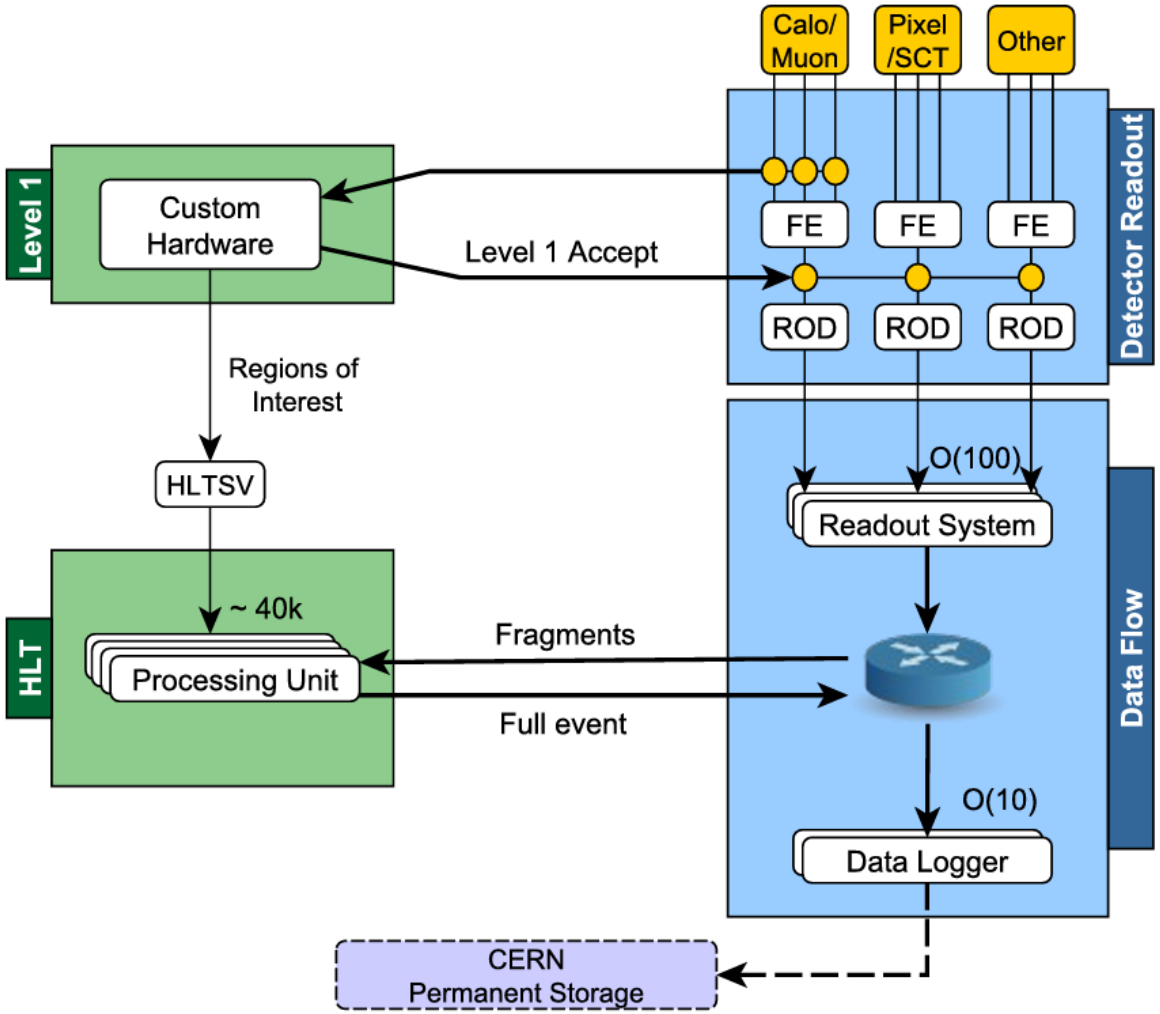

The ATLAS experiment TDAQ

Trigger

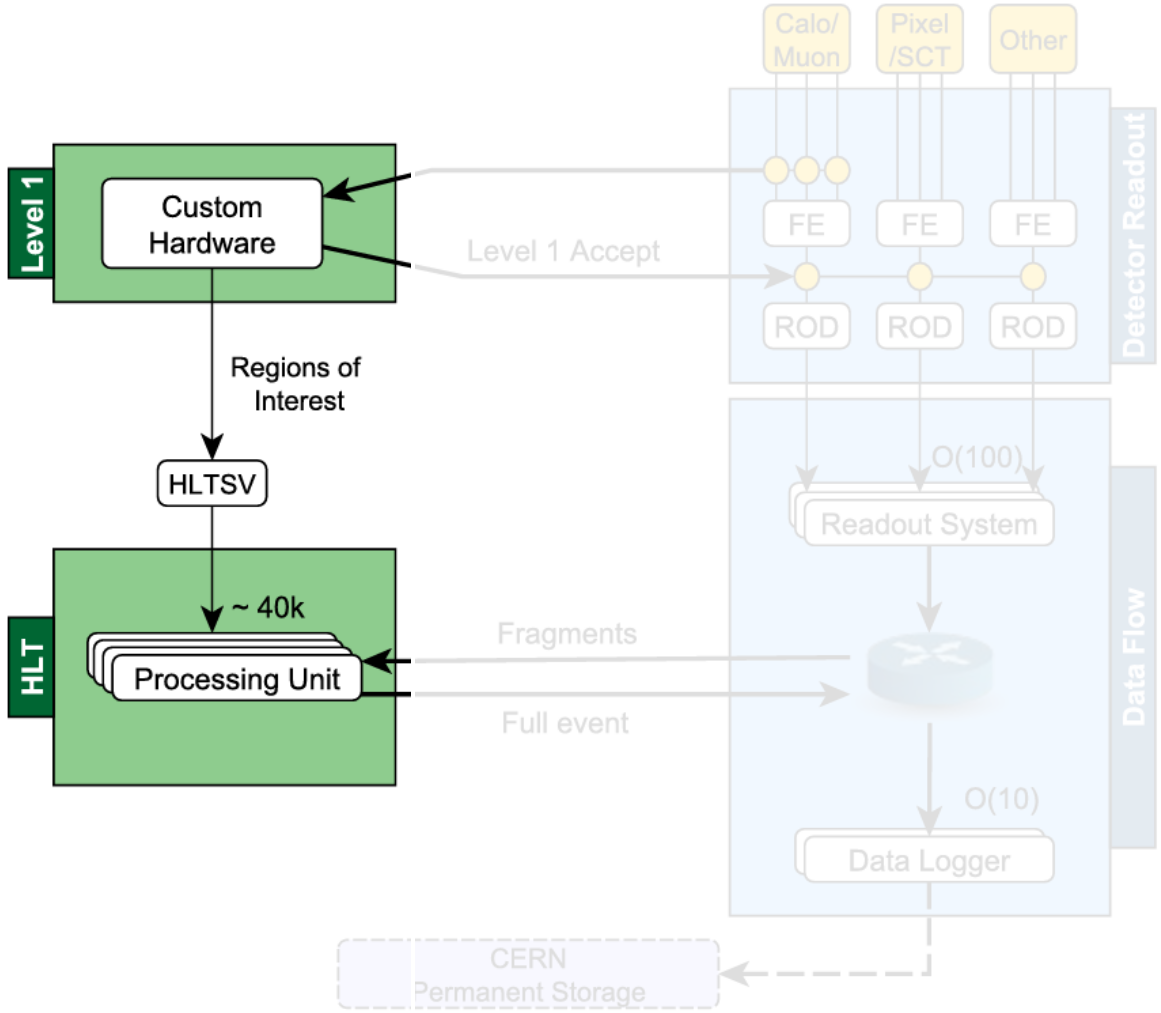

The ATLAS experiment TDAQ

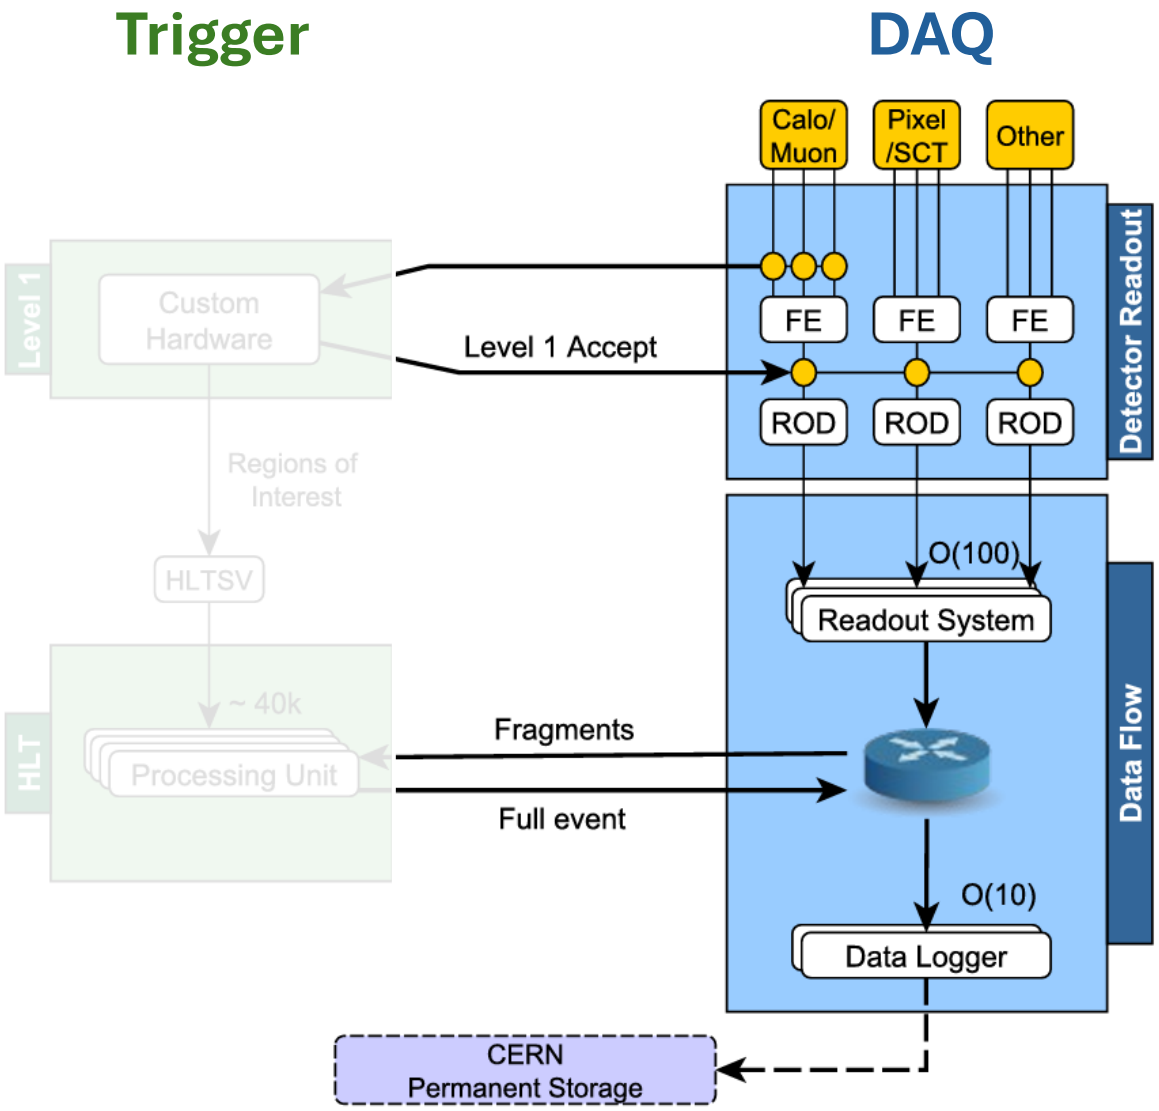

The ATLAS experiment TDAQ

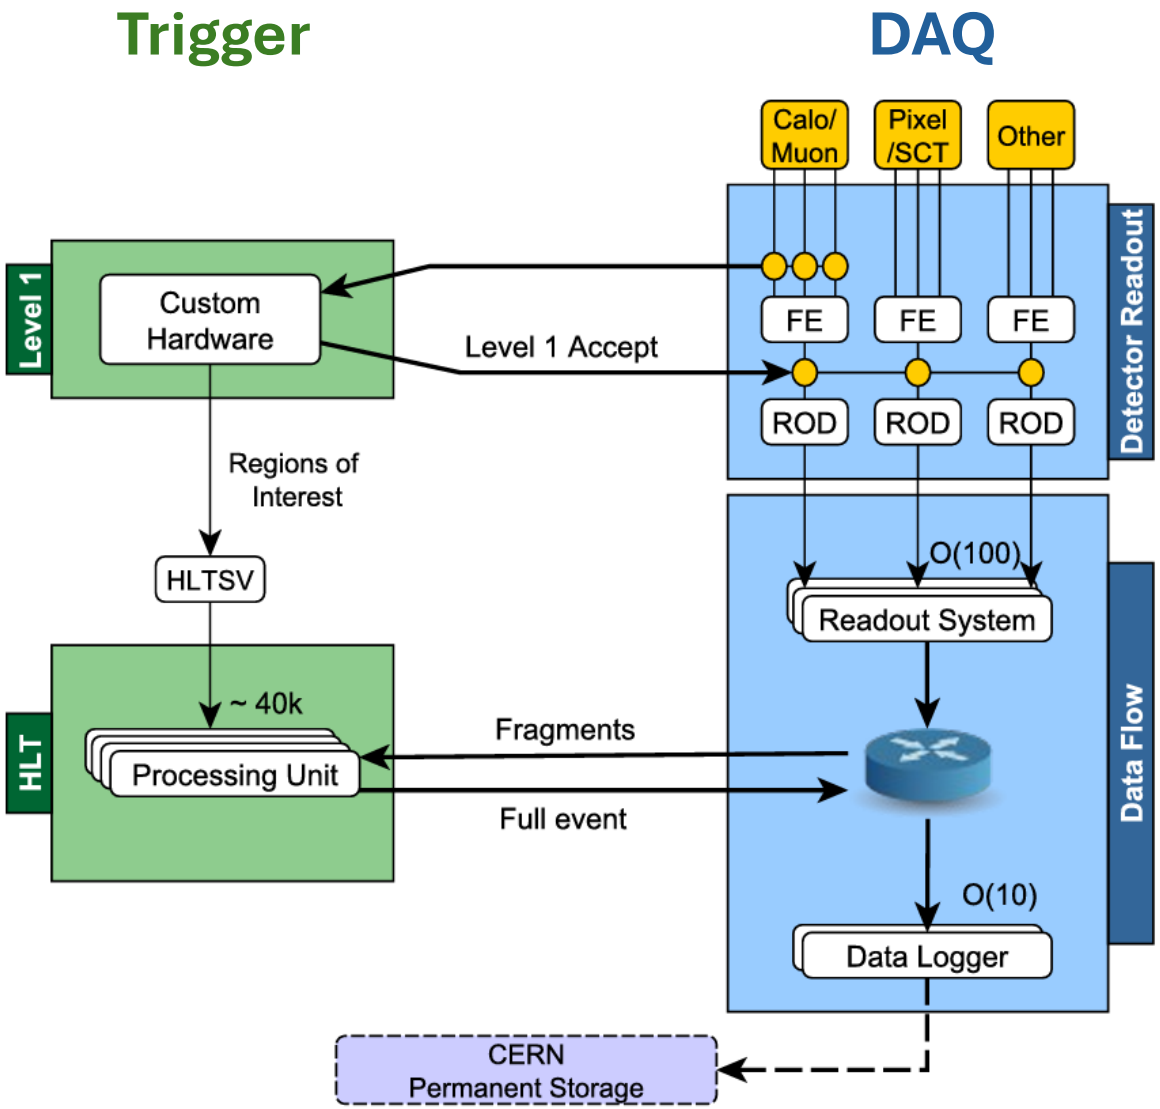

The ATLAS experiment TDAQ

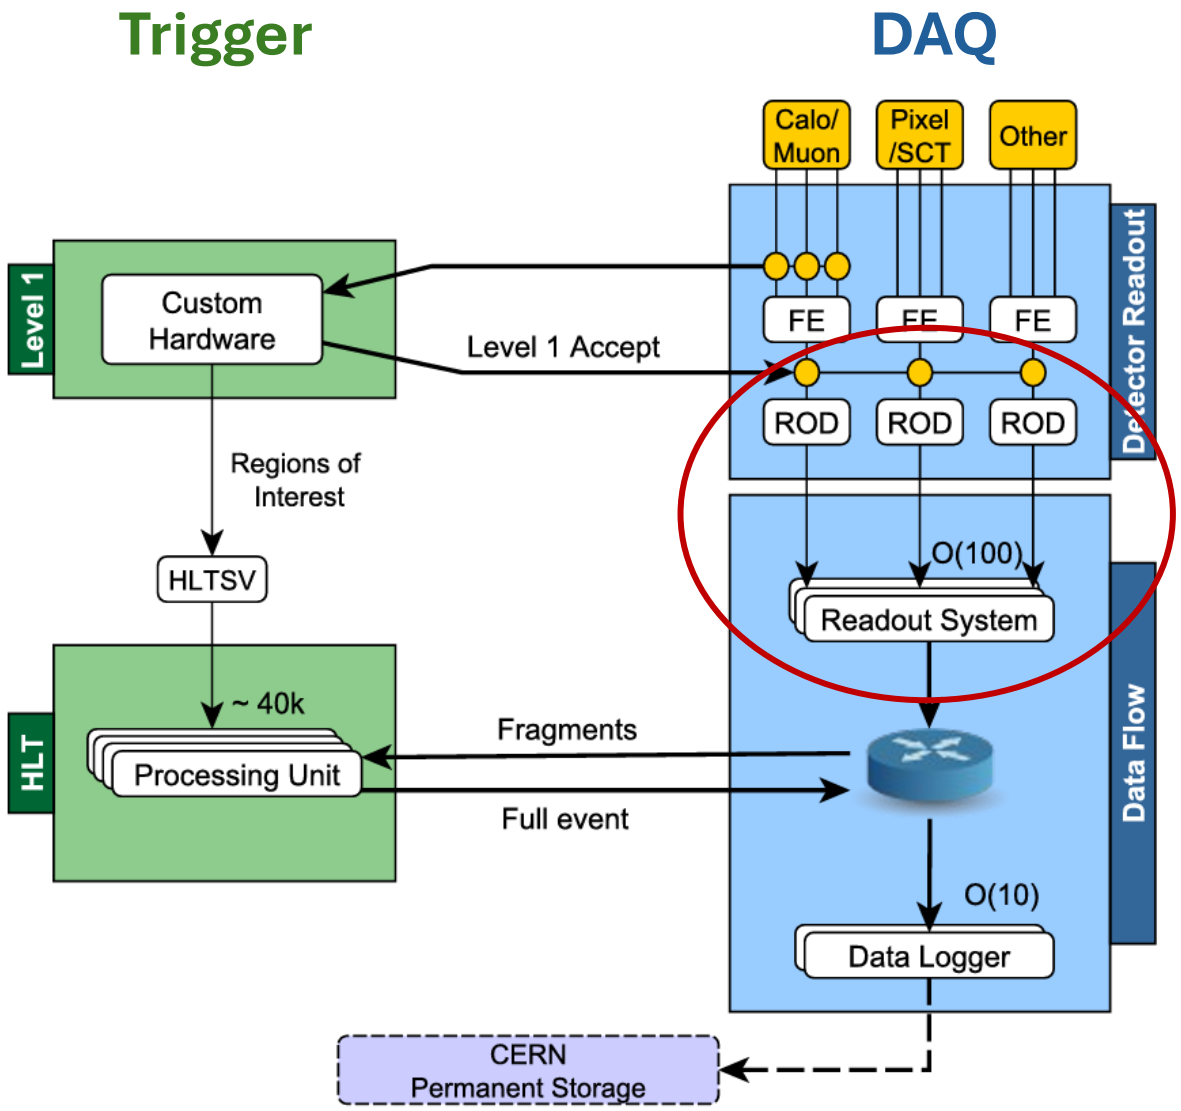

Readout
Gathers data from detectors

The ATLAS experiment TDAQ

Data Collection
Event Building

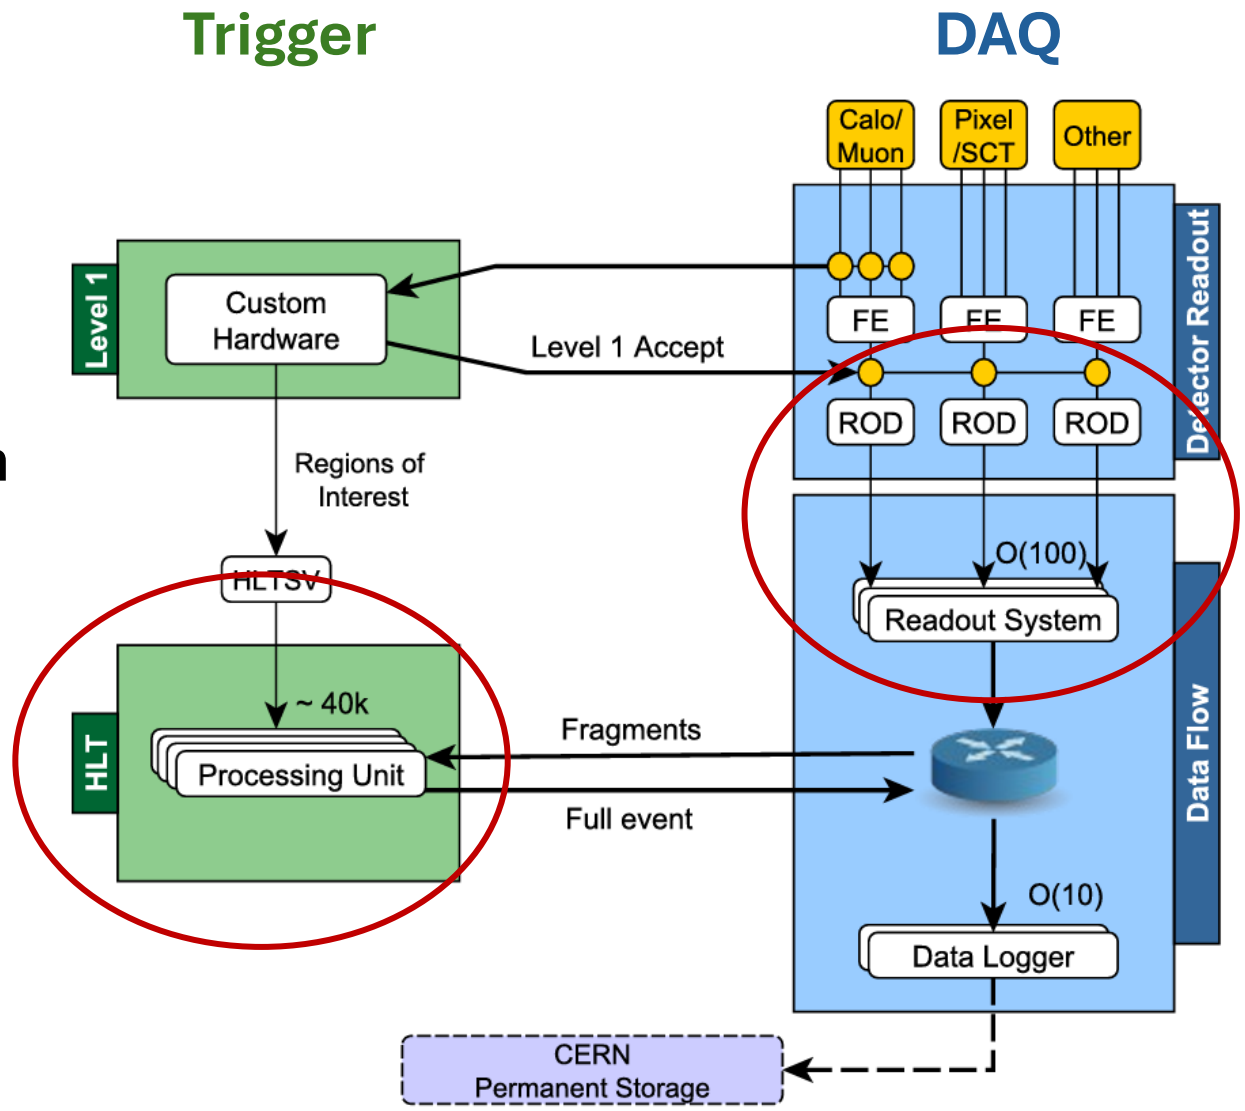

Readout
Gathers data from detectors

The ATLAS experiment TDAQ

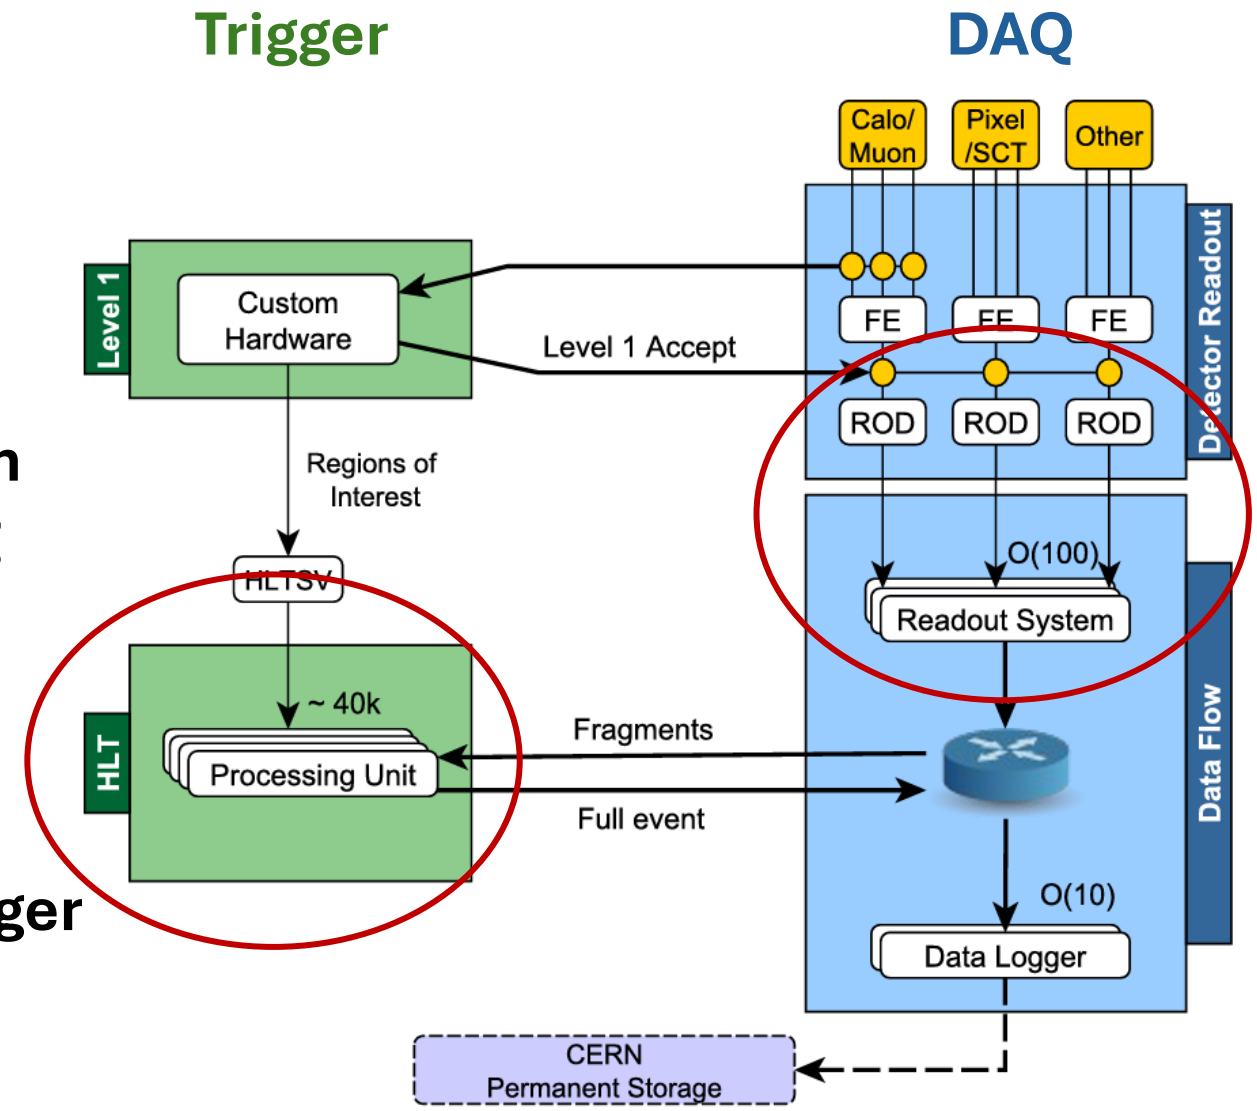

Data Collection Event Building

High Level Trigger

Readout
Gathers data from detectors

The ATLAS experiment TDAQ

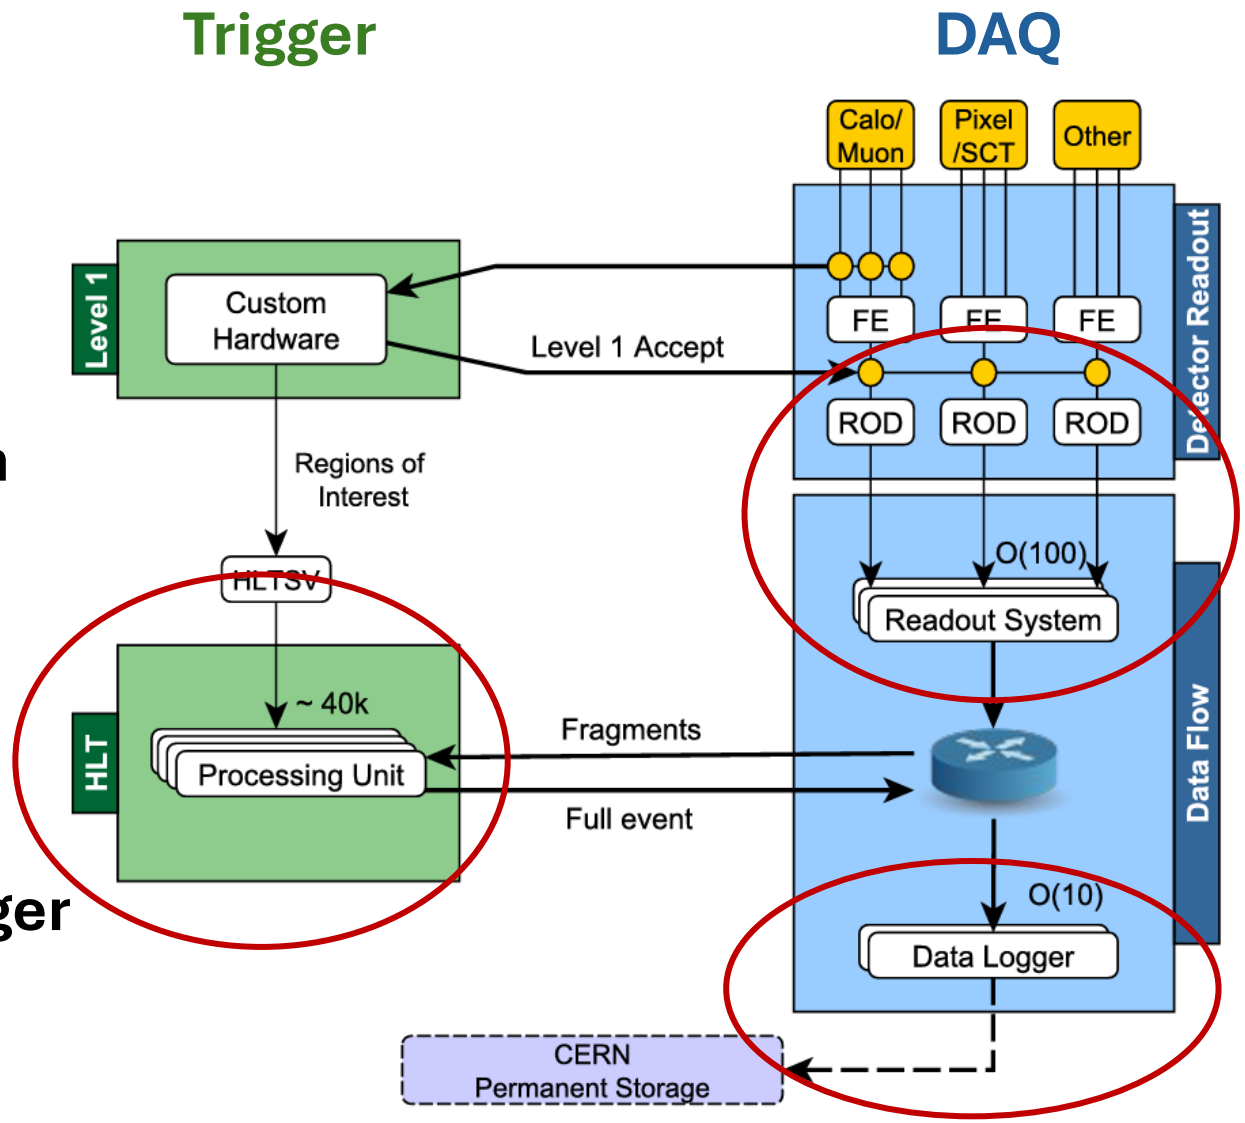

Data Collection Event Building

High Level Trigger

Readout
Gathers data from detectors

Data logging
Store accepted events data

The ATLAS experiment TDAQ

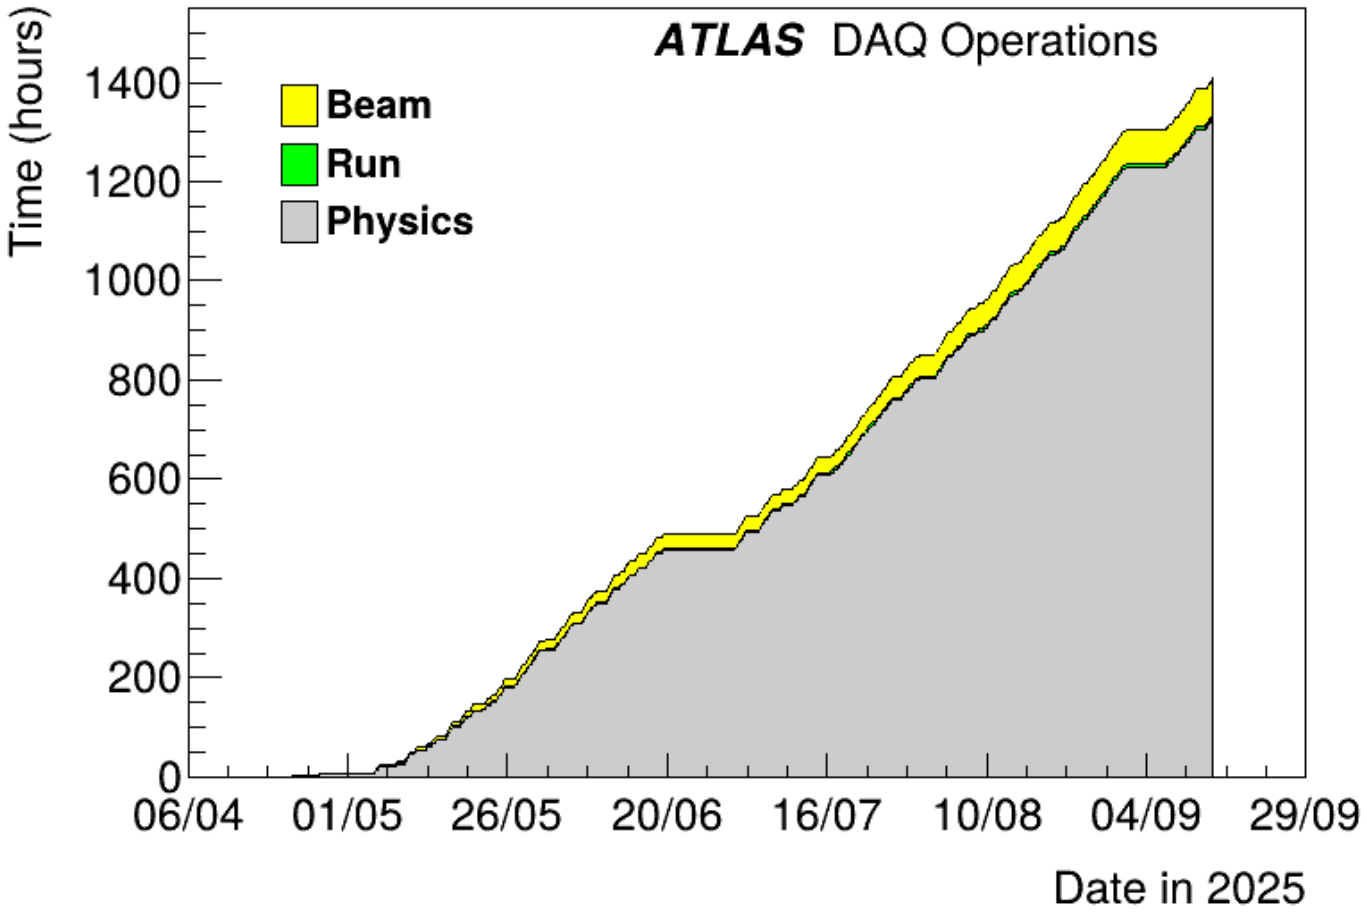

DAQ must maximize **efficiency**.

The ATLAS experiment TDAQ

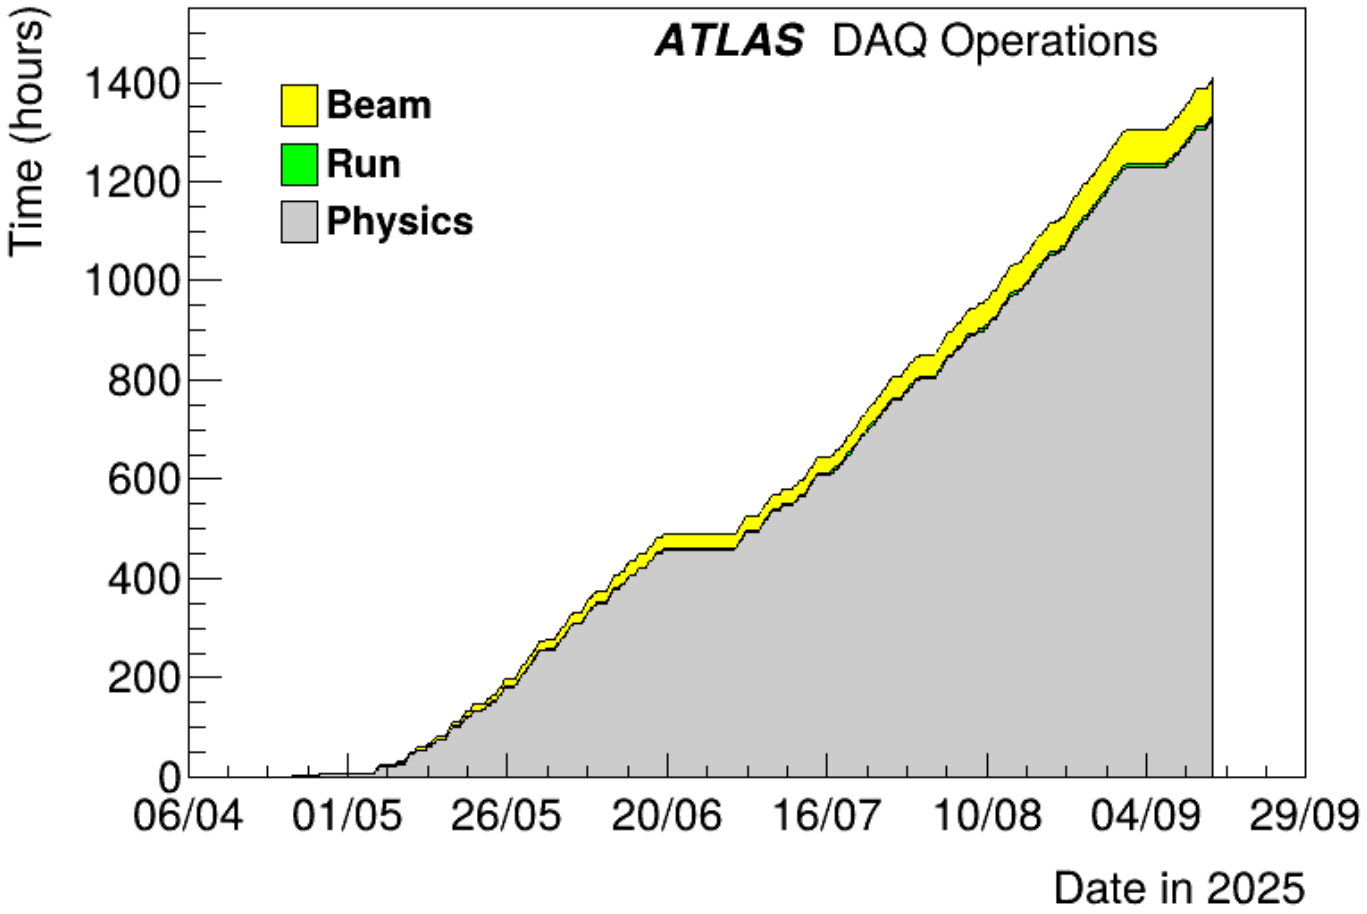

DAQ must maximize **efficiency**.

In previous Runs:

- Very small data loss from DAQ: **93%** all time avg efficiency
- Dataflow efficiency: **100%**

The ATLAS experiment TDAQ

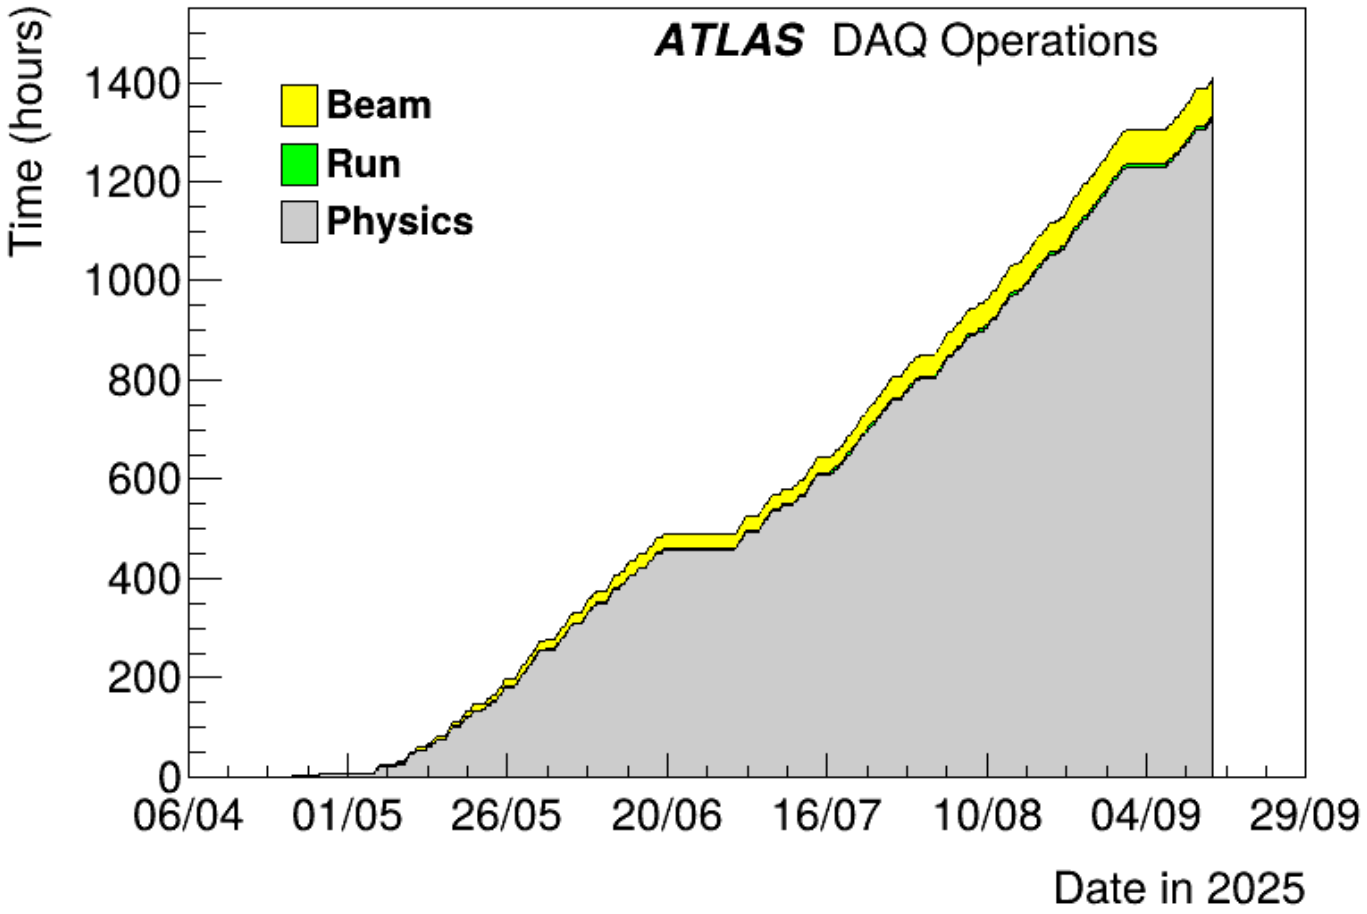

DAQ must maximize **efficiency**.

In previous Runs:

- Very small data loss from DAQ: **93%** all time avg efficiency
- Dataflow efficiency: **100%**

What happens if the Dataflow system fails?

The ATLAS experiment TDAQ

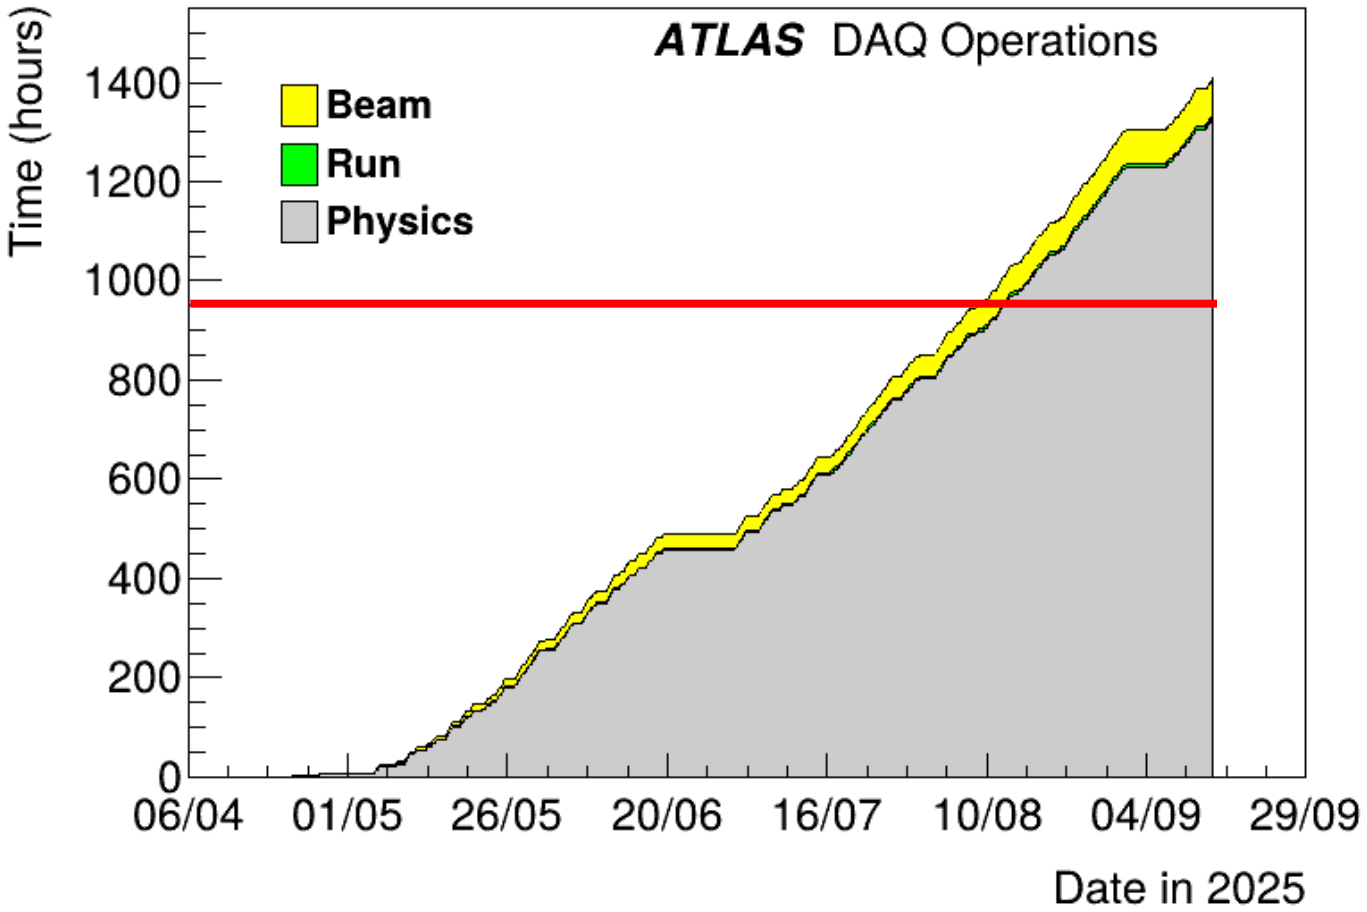

DAQ must maximize **efficiency**.

In previous Runs:

- Very small data loss from DAQ: **93%** all time avg efficiency
- Dataflow efficiency: **100%**

What happens if the Dataflow system fails?

With 80% efficiency:

- DAQ efficiency would drop to 74,4%
- **Severe impact** on rare event detection and measurement

The ATLAS experiment TDAQ

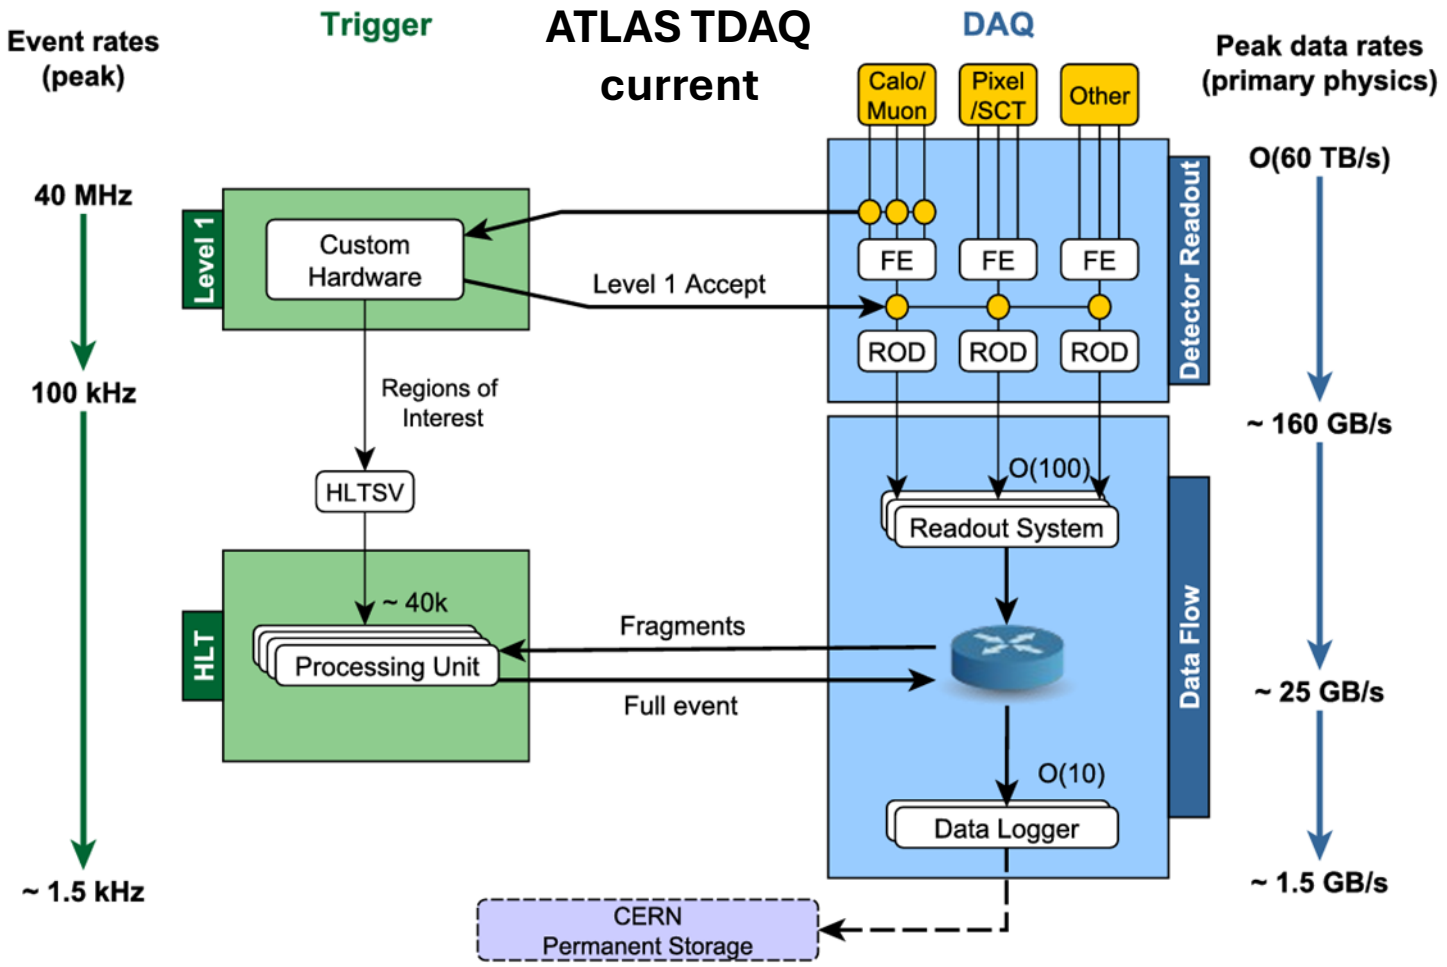

Dataflow is the backbone of DAQ:

The ATLAS experiment TDAQ

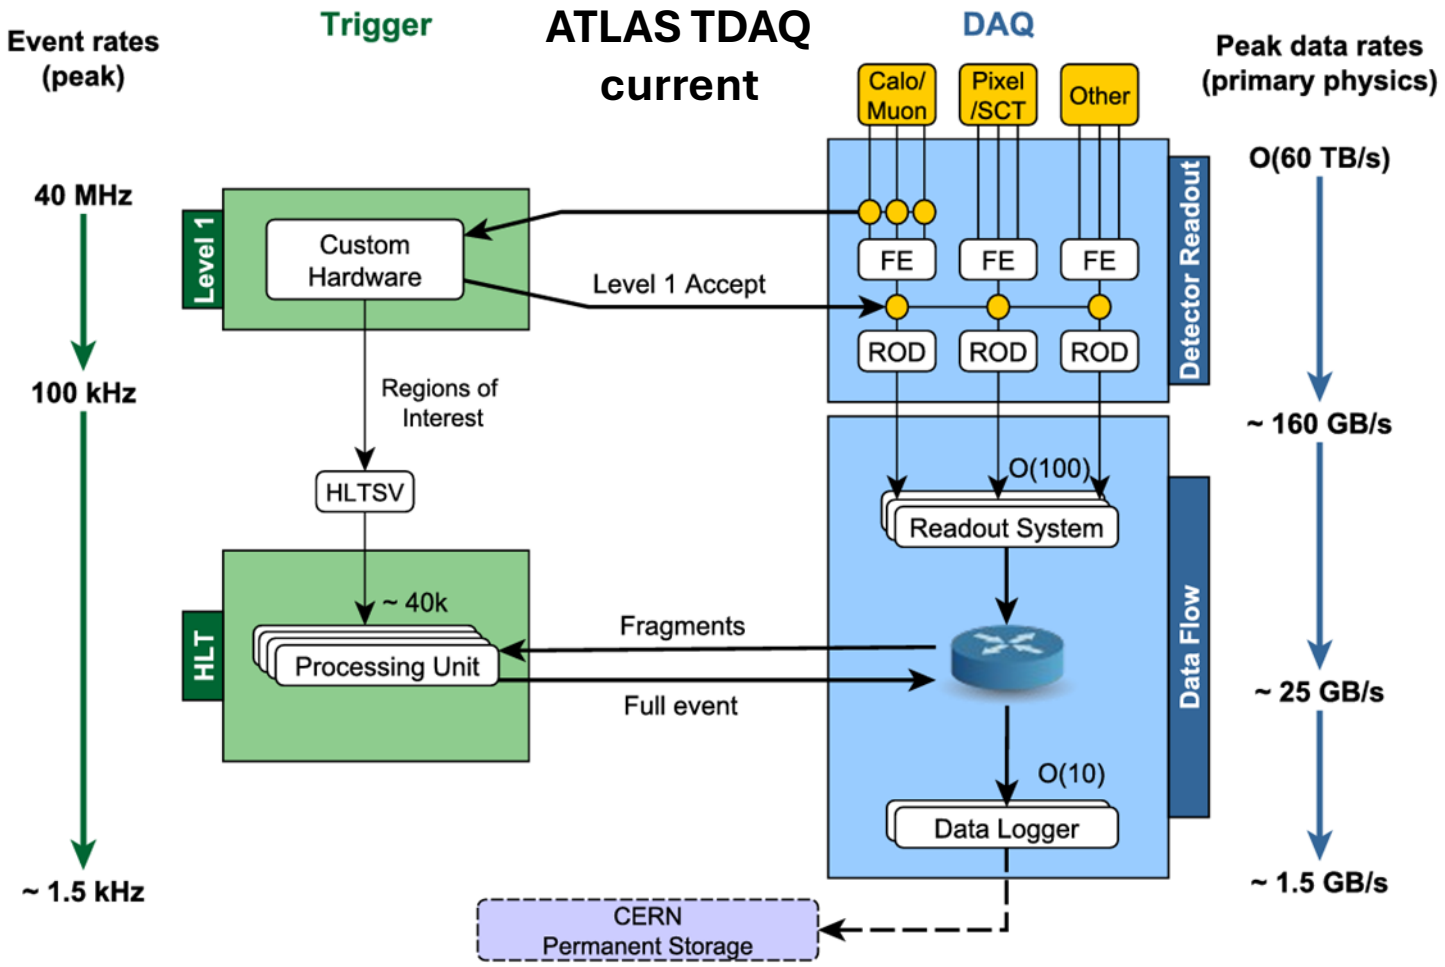

Dataflow is the backbone of DAQ:

- Receives events data at **100 kHz** and 160 GB/s from detectors **Readout**
- **Builds events** for reconstruction and filtering at **25 GB/s**
- **Stores** 1,5 kHz of selected events with 1.5 GB/s data rate

The ATLAS experiment TDAQ

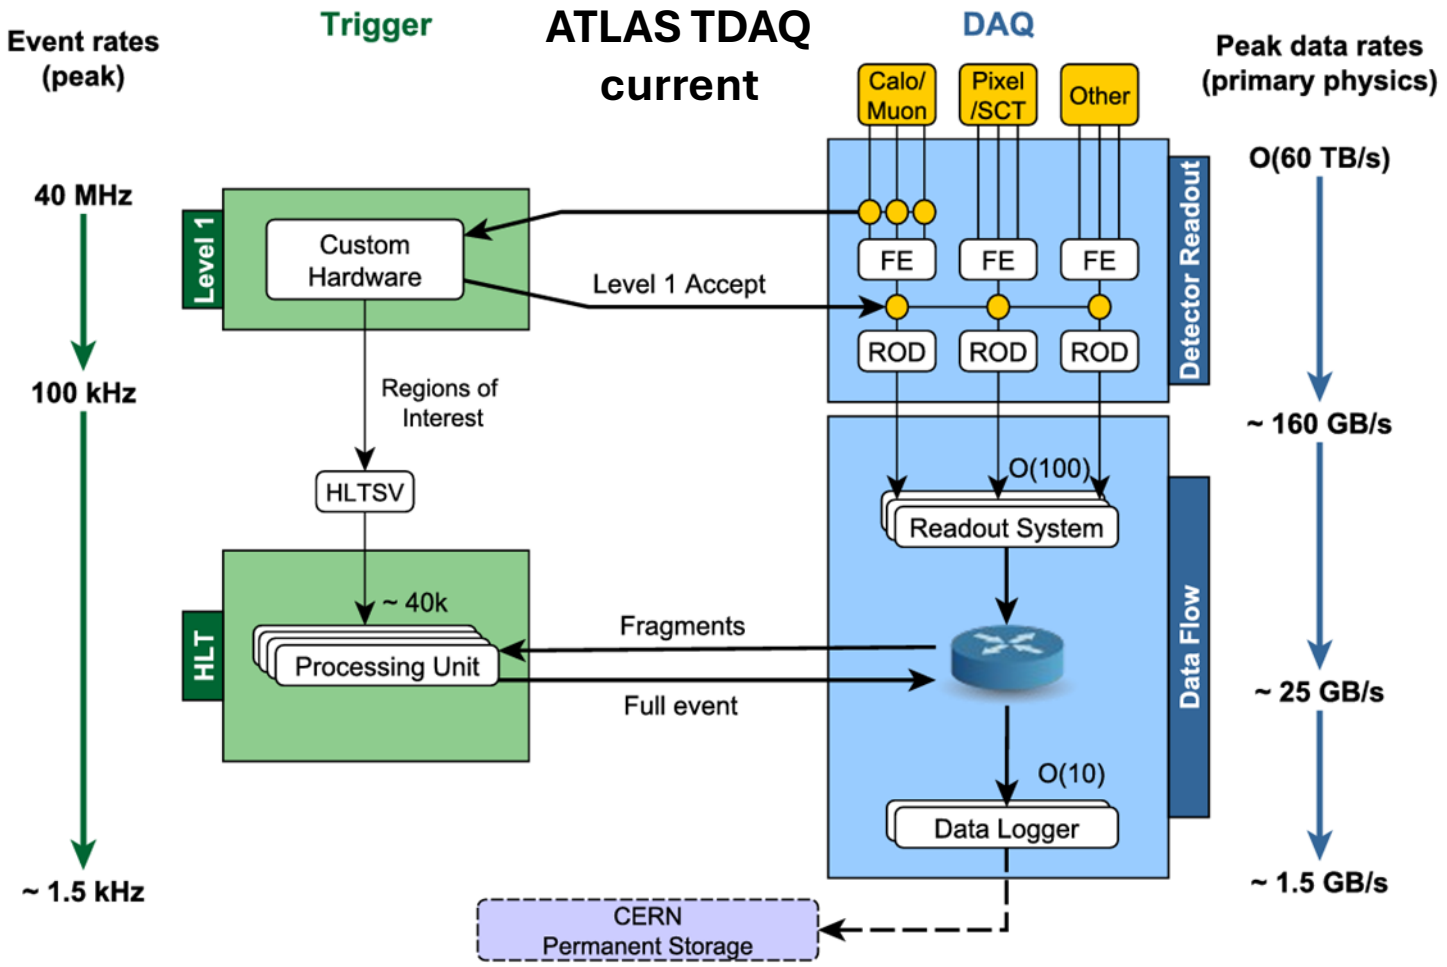

Dataflow is the backbone of DAQ:

- Receives events data at **100 kHz** and 160 GB/s from detectors **Readout**
- **Builds events** for reconstruction and filtering at **25 GB/s**
- **Stores** 1,5 kHz of selected events with 1.5 GB/s data rate
- Always maintains maximum efficiency.
  If it stops, **everything** stops

The High Luminosity Upgrade

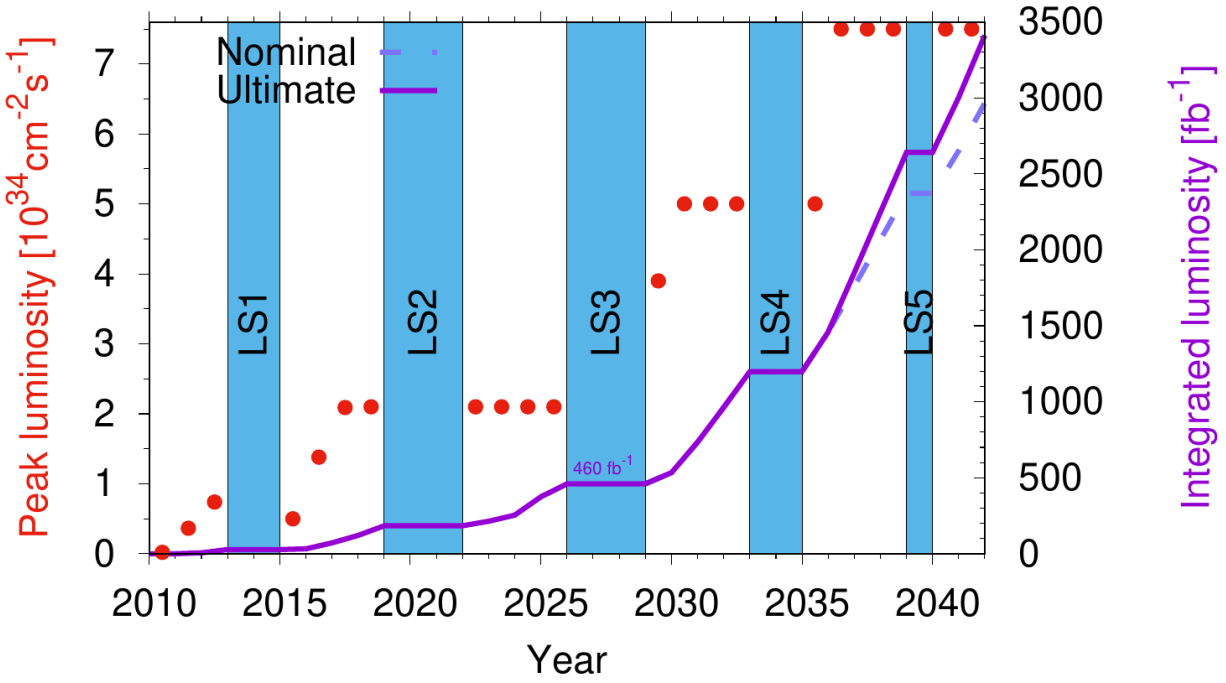

The High Luminosity Upgrade

UNIVERSITÀ DI PAVIA

You are here

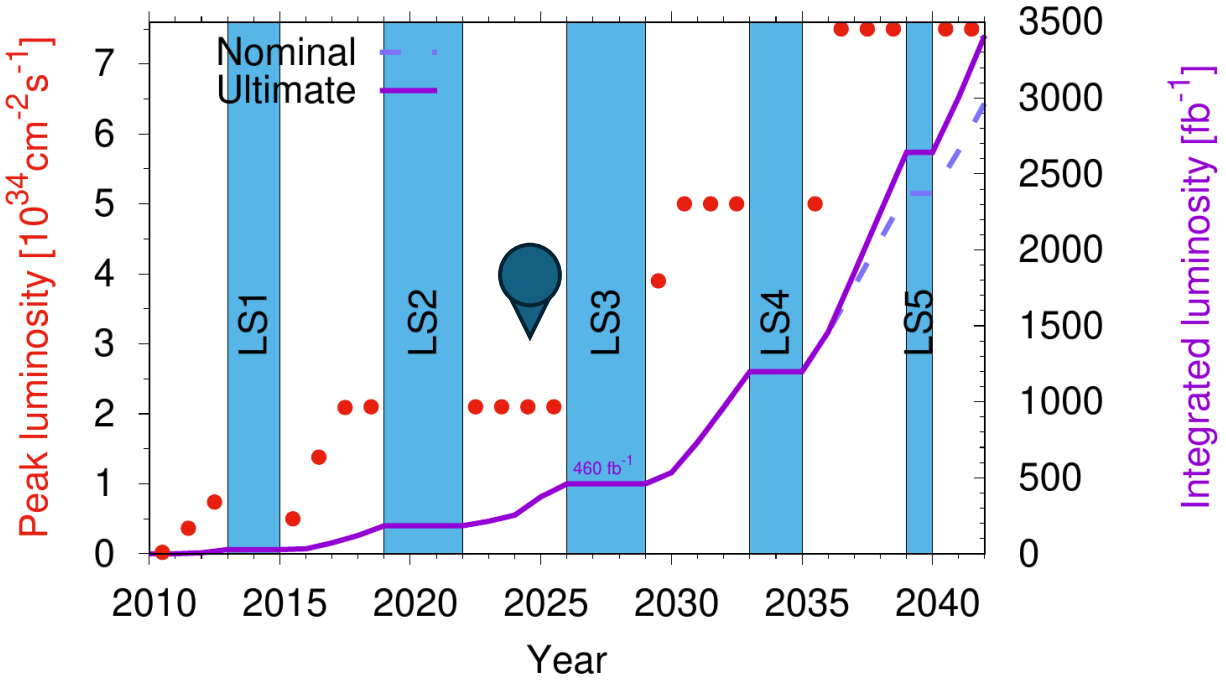

The High Luminosity Upgrade

UNIVERSITÀ DI PAVIA

You are here

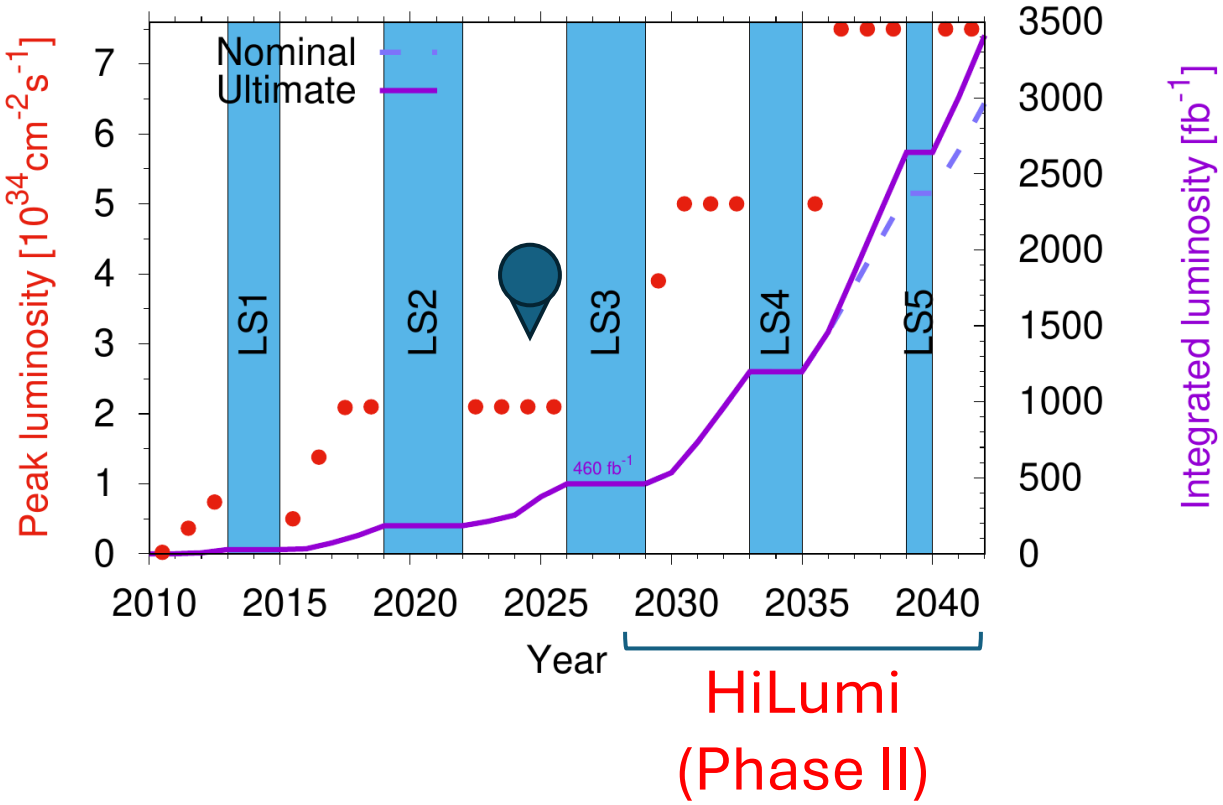

Major LHC upgrade starting from late 2026: **High Luminosity Upgrade**

The High Luminosity Upgrade

You are here

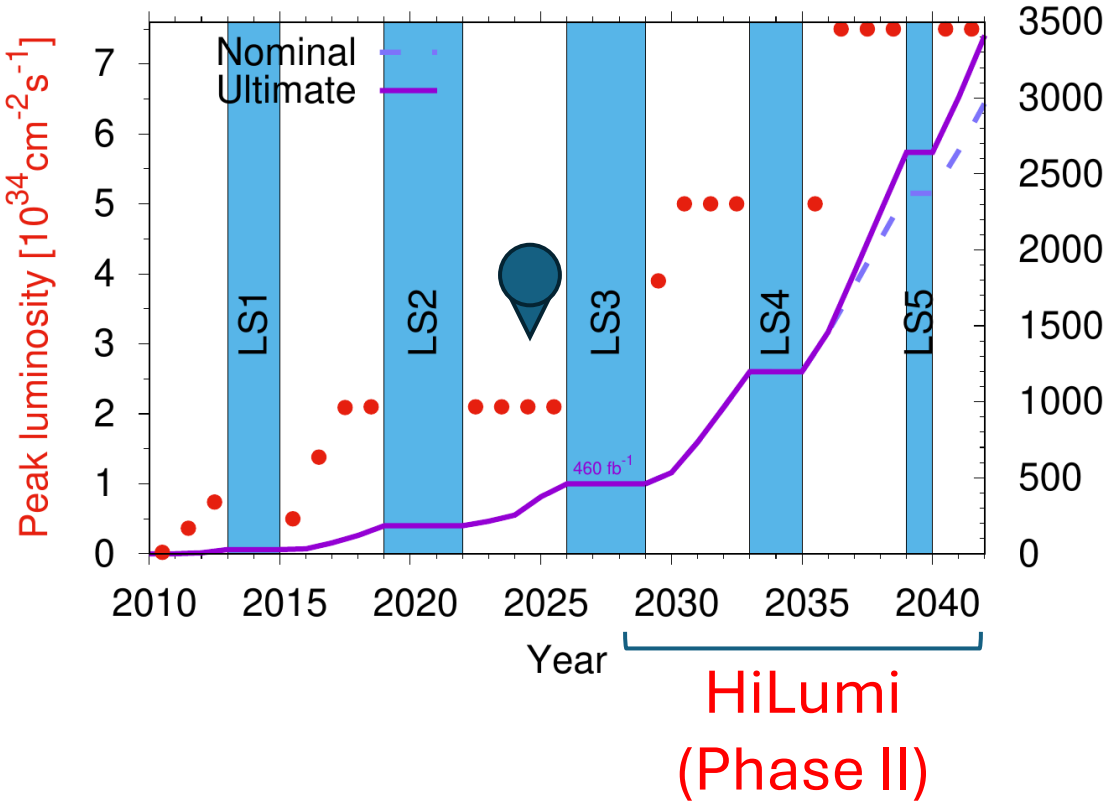

Major LHC upgrade starting from late 2026: **High Luminosity Upgrade**

- Top priority of the *European Strategy for Particle Physics*
- **4x** Peak luminosity increase
- Total integrated luminosity: **10x** LHC design value

The High Luminosity Upgrade

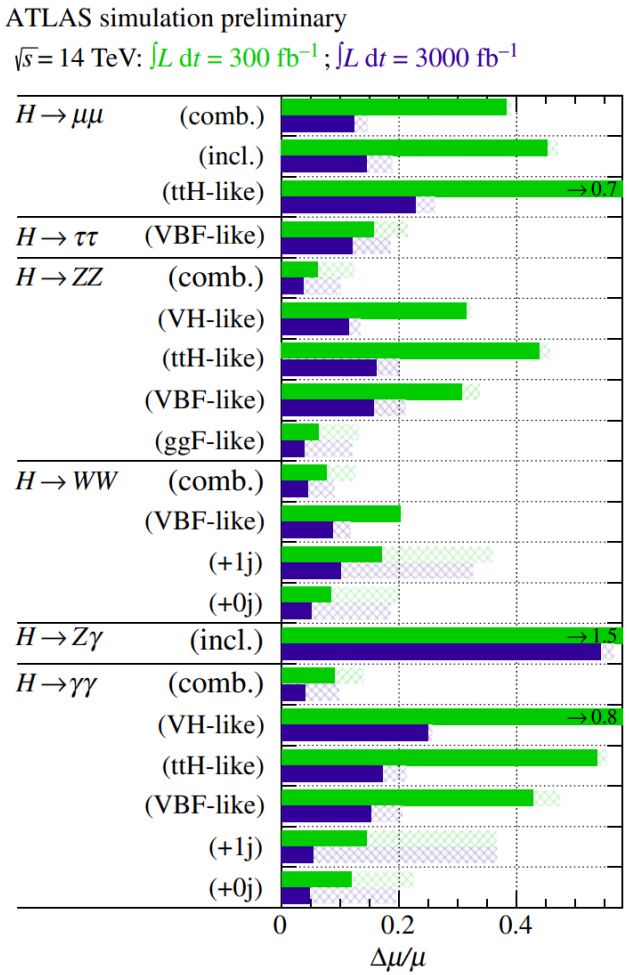

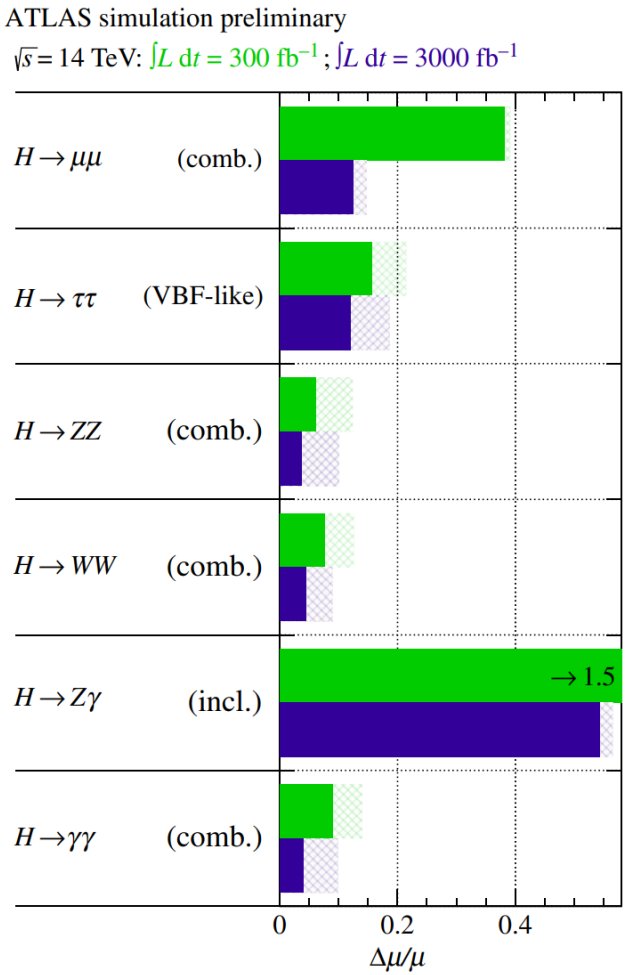

The **High Luminosity Upgrade** aim is:

- Switch from discovery to **precision measurements**
- **Extend** the accelerator **lifetime** collecting as much data as possible
- Significantly **improve precision** for the Higgs sector
- Observe **very rare** new phenomena

The High Luminosity Upgrade

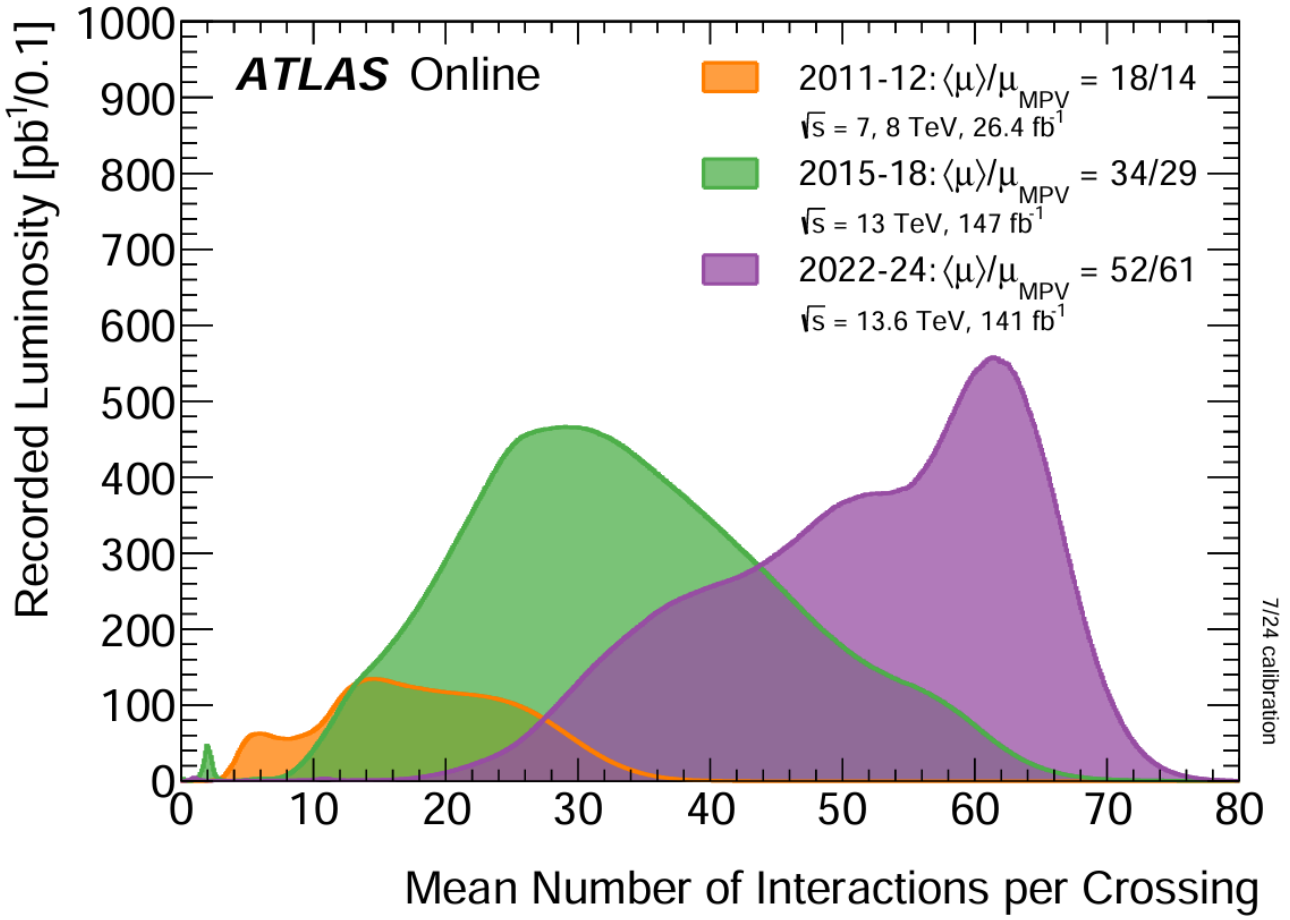

The upgrade will bring new technical **challenges**:

- Higher **Pile-up** is necessary to increase the luminosity
- Average of 60 during current Run
- Will reach up to **200** in Phase II

The High Luminosity Upgrade

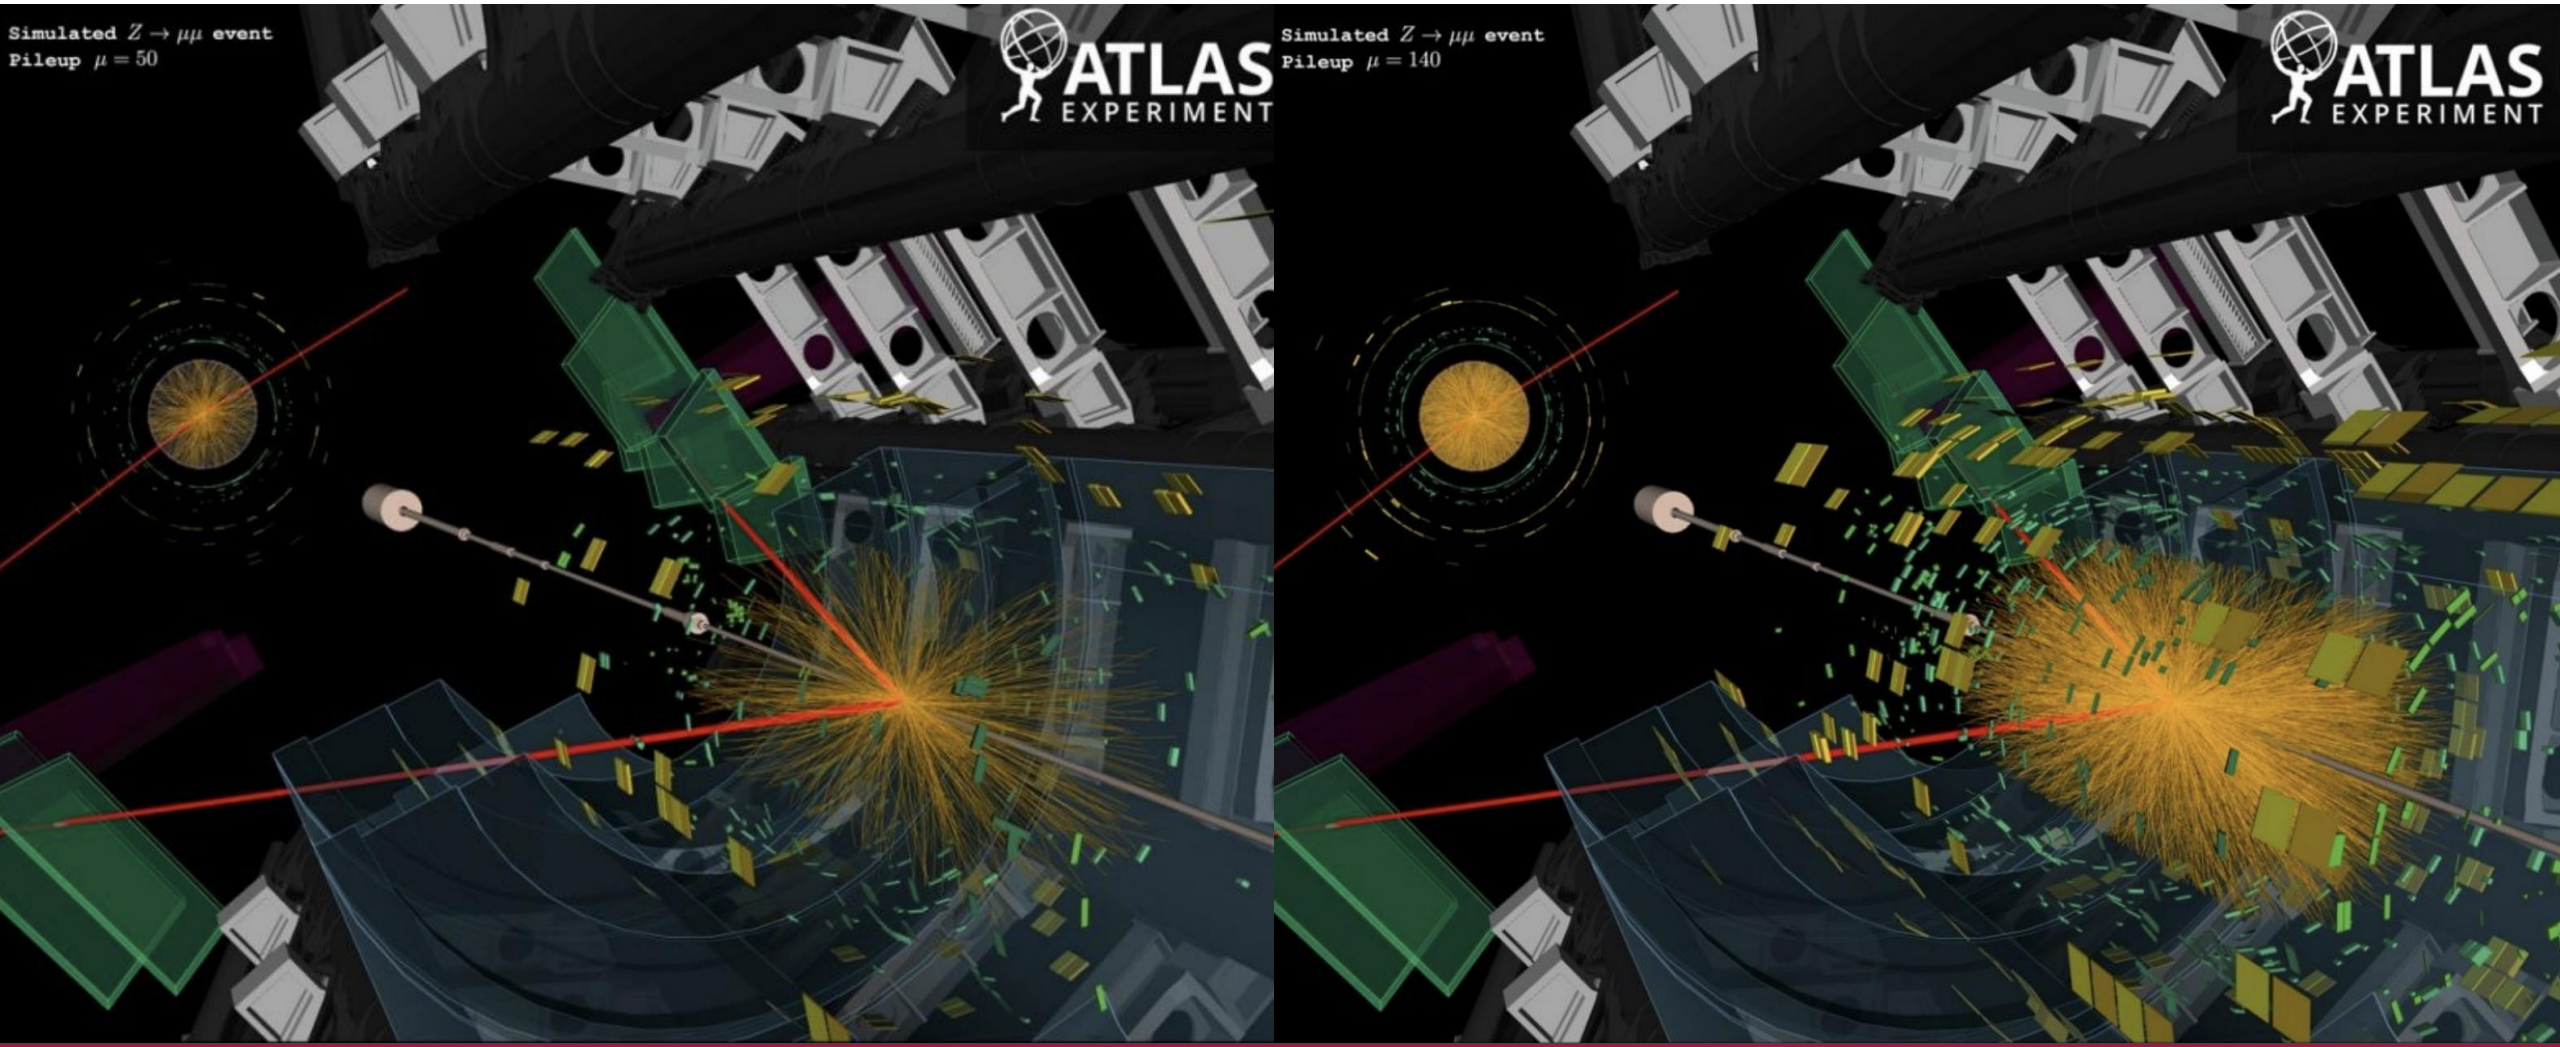

The High Luminosity Upgrade

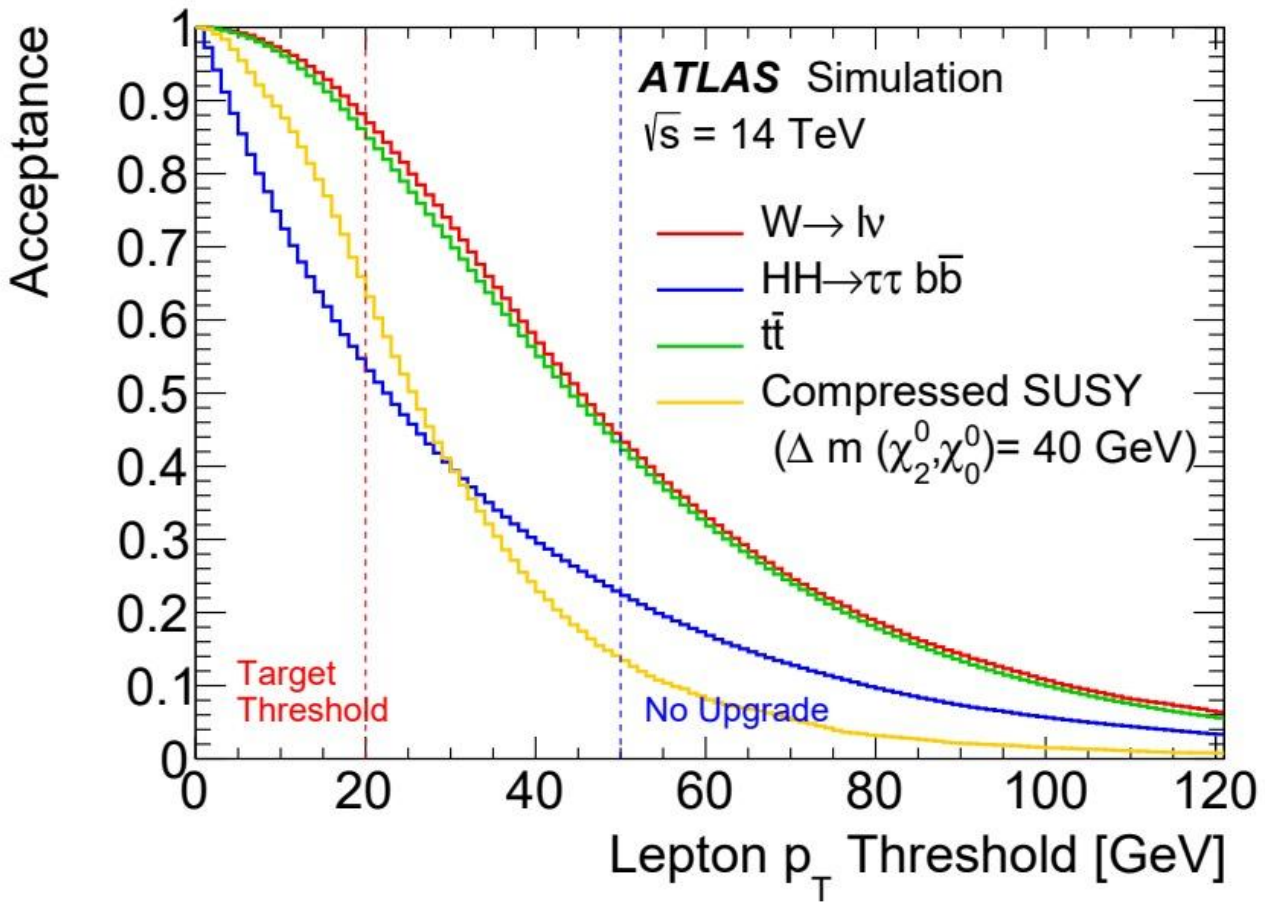

To solve the issue, one can raise the **thresholds, but**

- severely reduce acceptance

The High Luminosity Upgrade

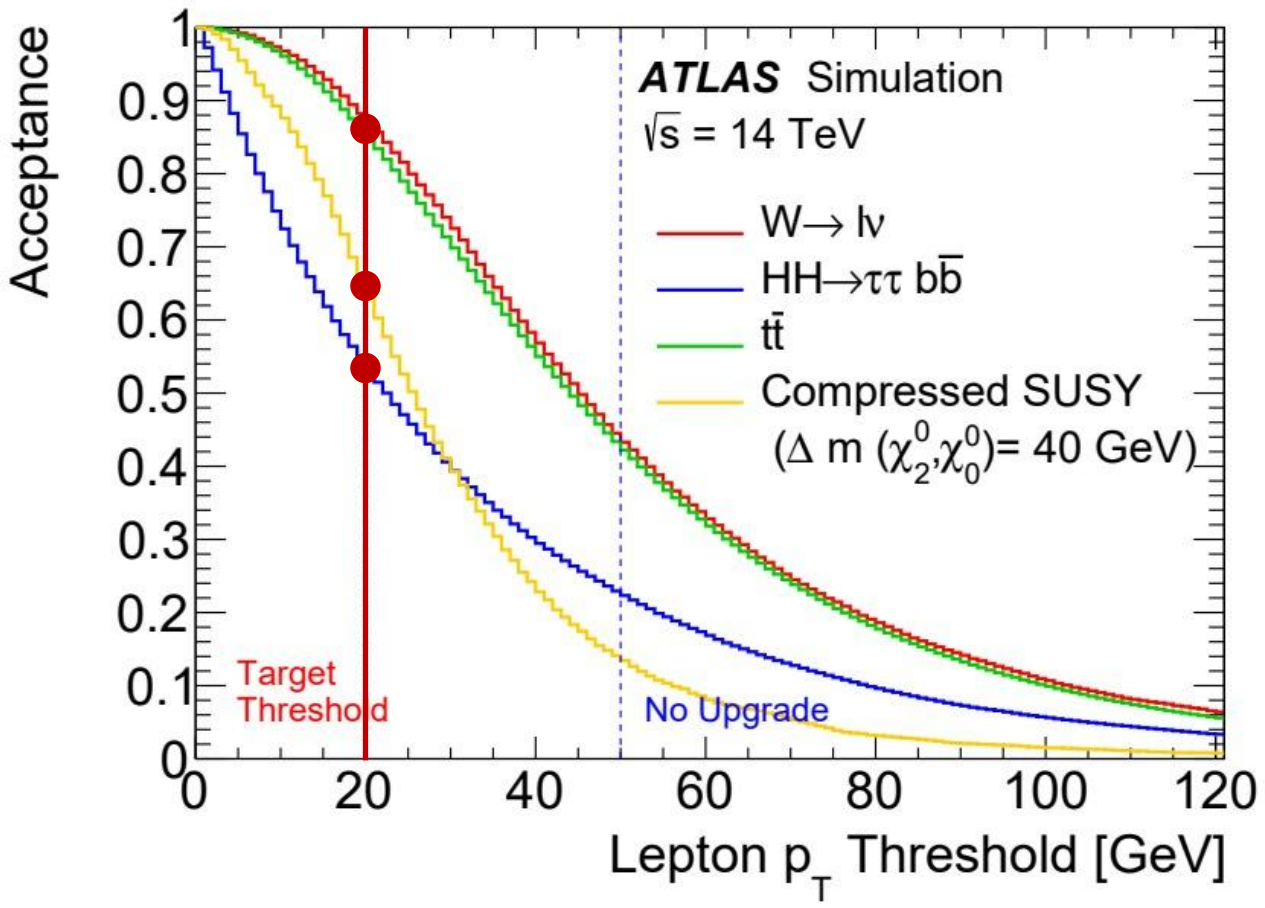

To solve the issue, one can raise the **thresholds, but**

- severely reduce acceptance

With the current detector this means

The High Luminosity Upgrade

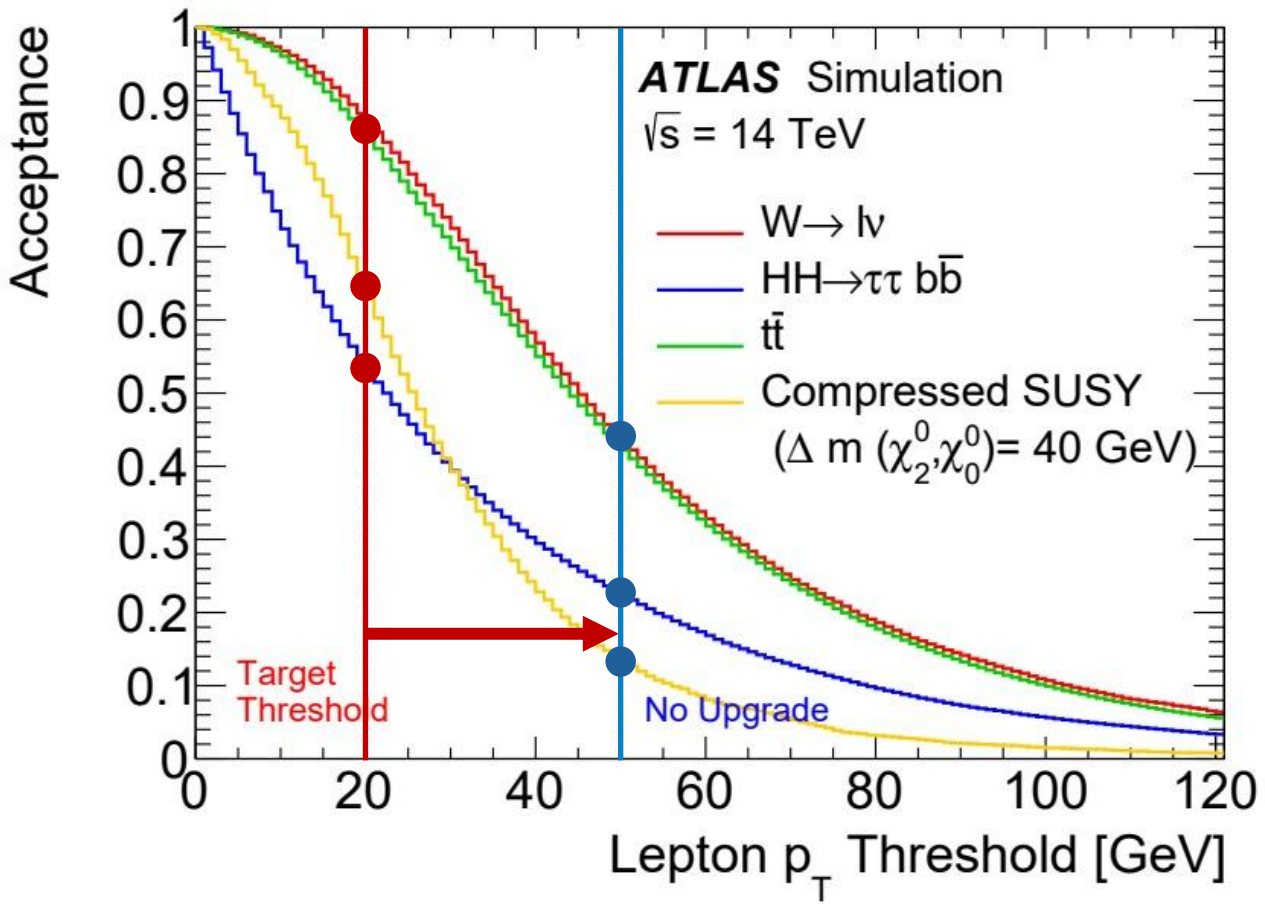

To solve the issue, one can raise the **thresholds, but**

- severely reduce acceptance

With the current detector this means

- Raising single μ, e^- thresholds from 20 GeV to 50 GeV
- Efficiency dropping by almost 50%

Trigger rate must be increased

ATLAS Phase-II Upgrade Requirements

Luminosity and pile-up increase will provide a significant challenge for the ATLAS TDAQ.

Both data rate and throughput will increase significantly.

ATLAS Phase-II Upgrade Requirements

Luminosity and pile-up increase will provide a significant challenge for the ATLAS TDAQ.

Both data rate and throughput will increase significantly.

- ~ 10x increase in input event rate
- ~ 25x increase in data throughput

ATLAS Phase-II Upgrade Requirements

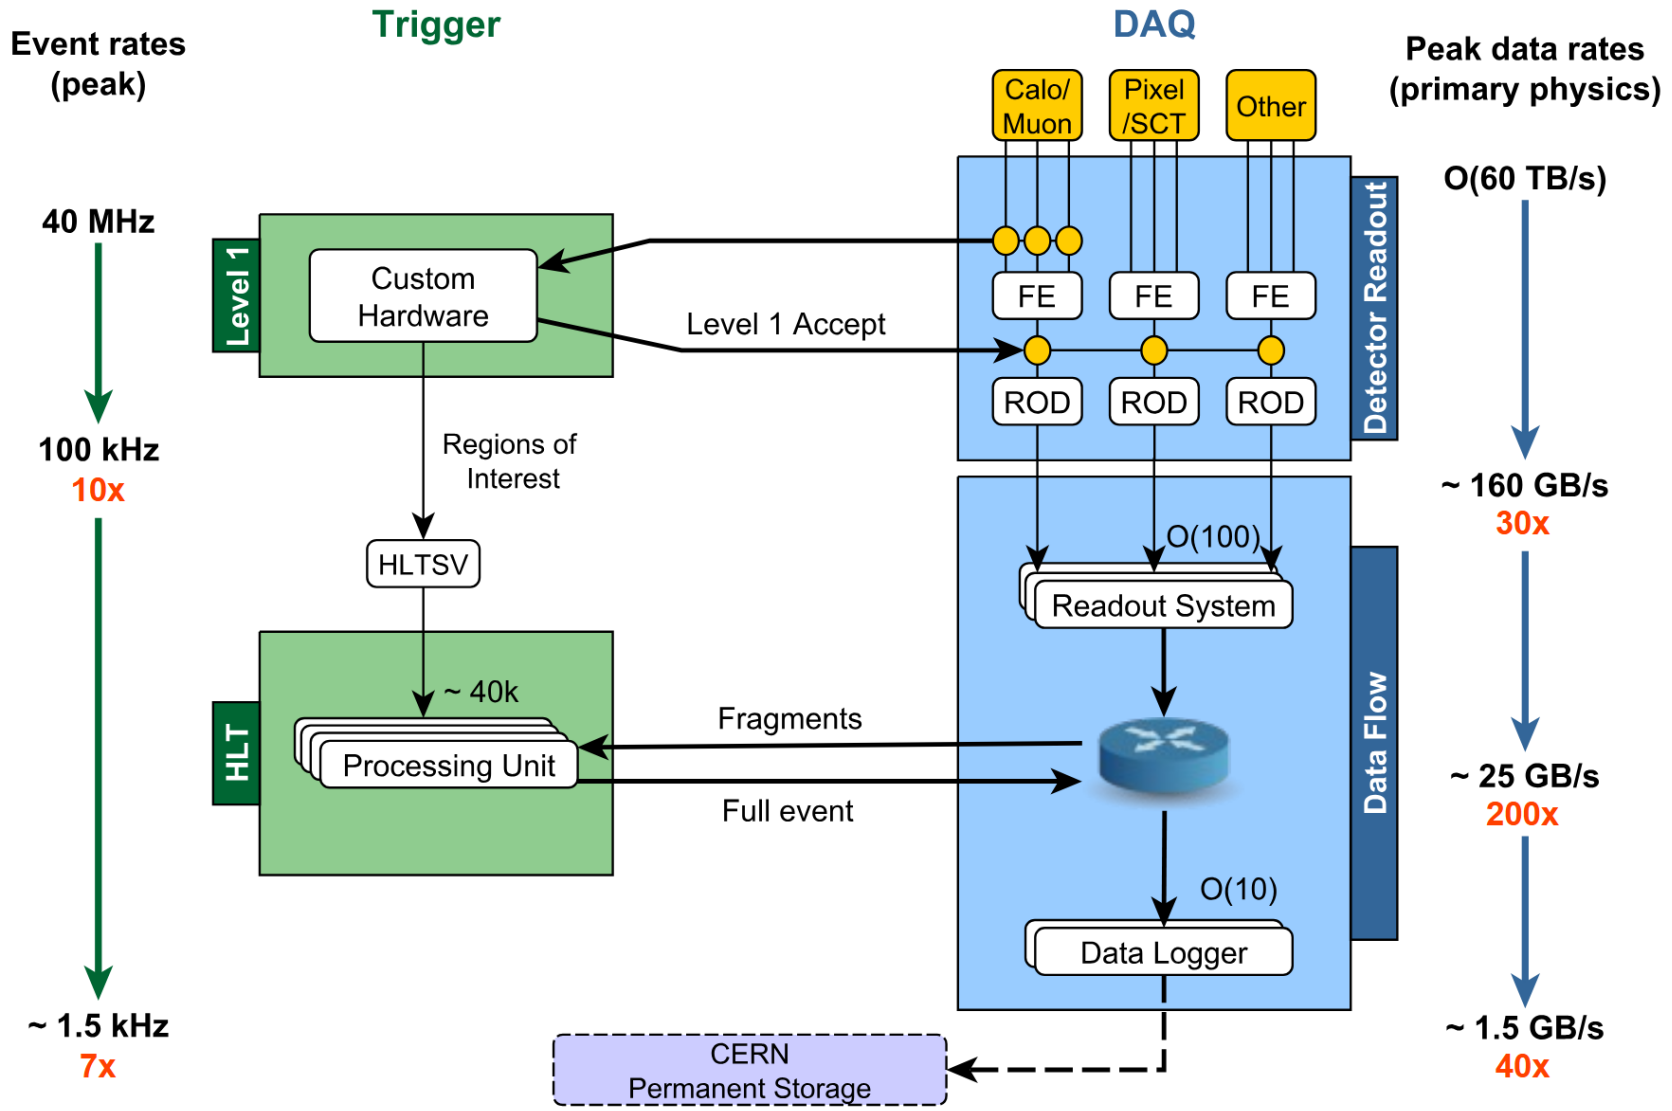

ATLAS Phase-II Upgrade Requirements

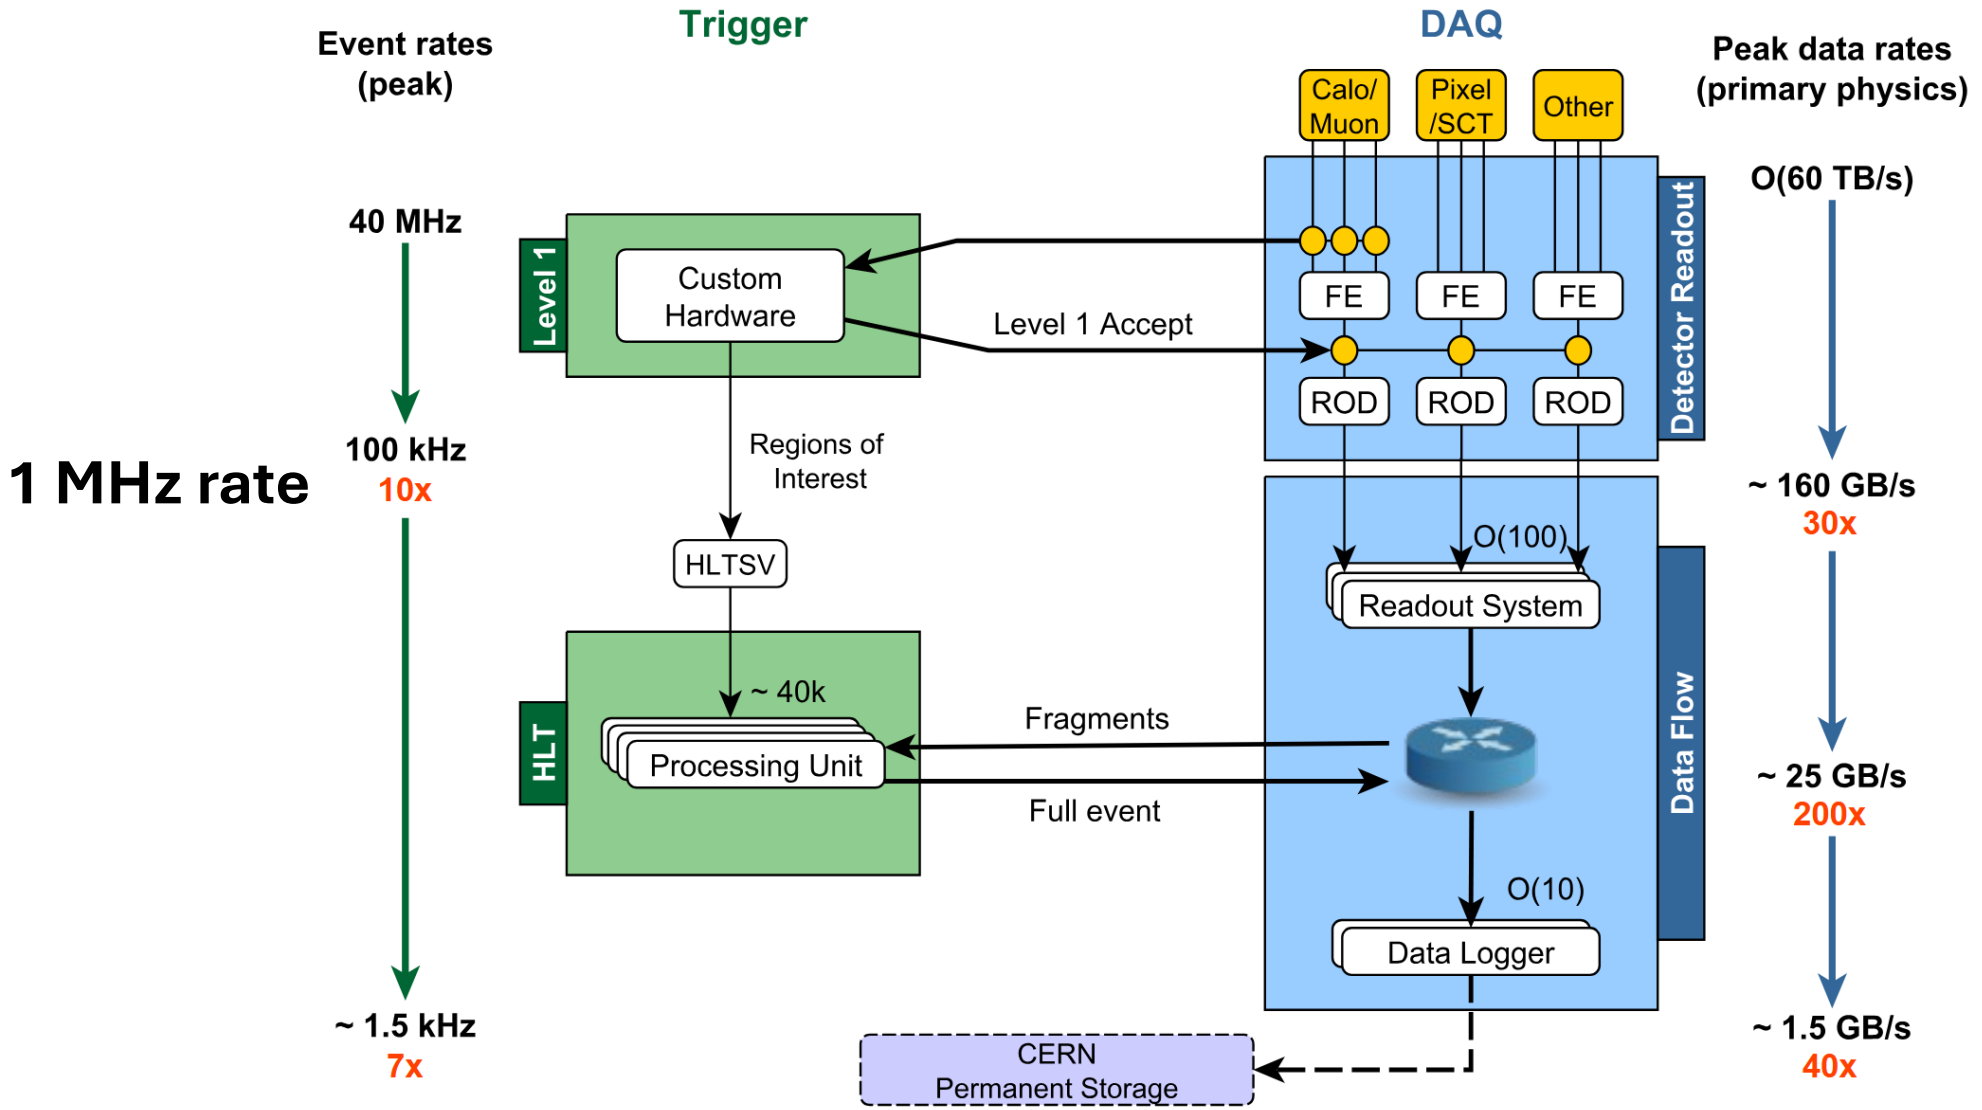

ATLAS Phase-II Upgrade Requirements

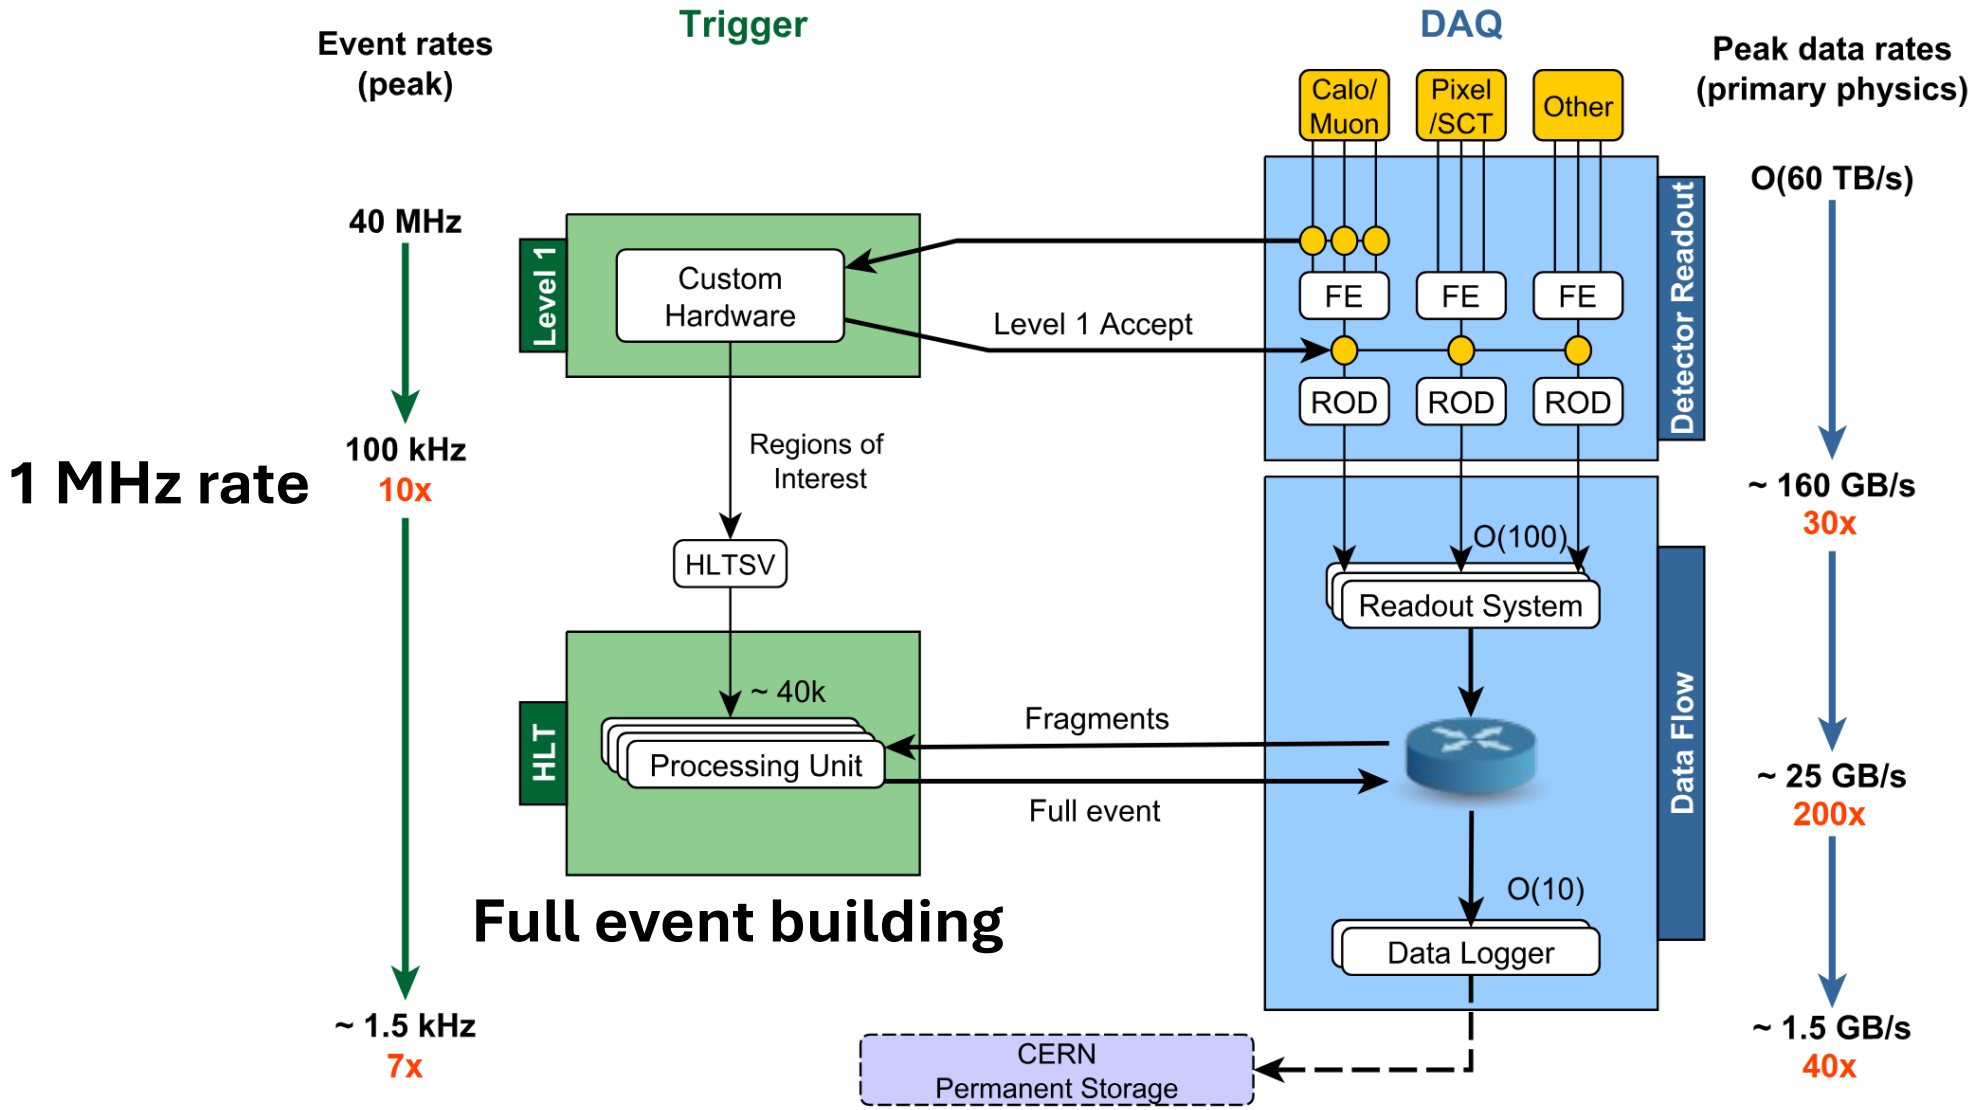

ATLAS Phase-II Upgrade Requirements

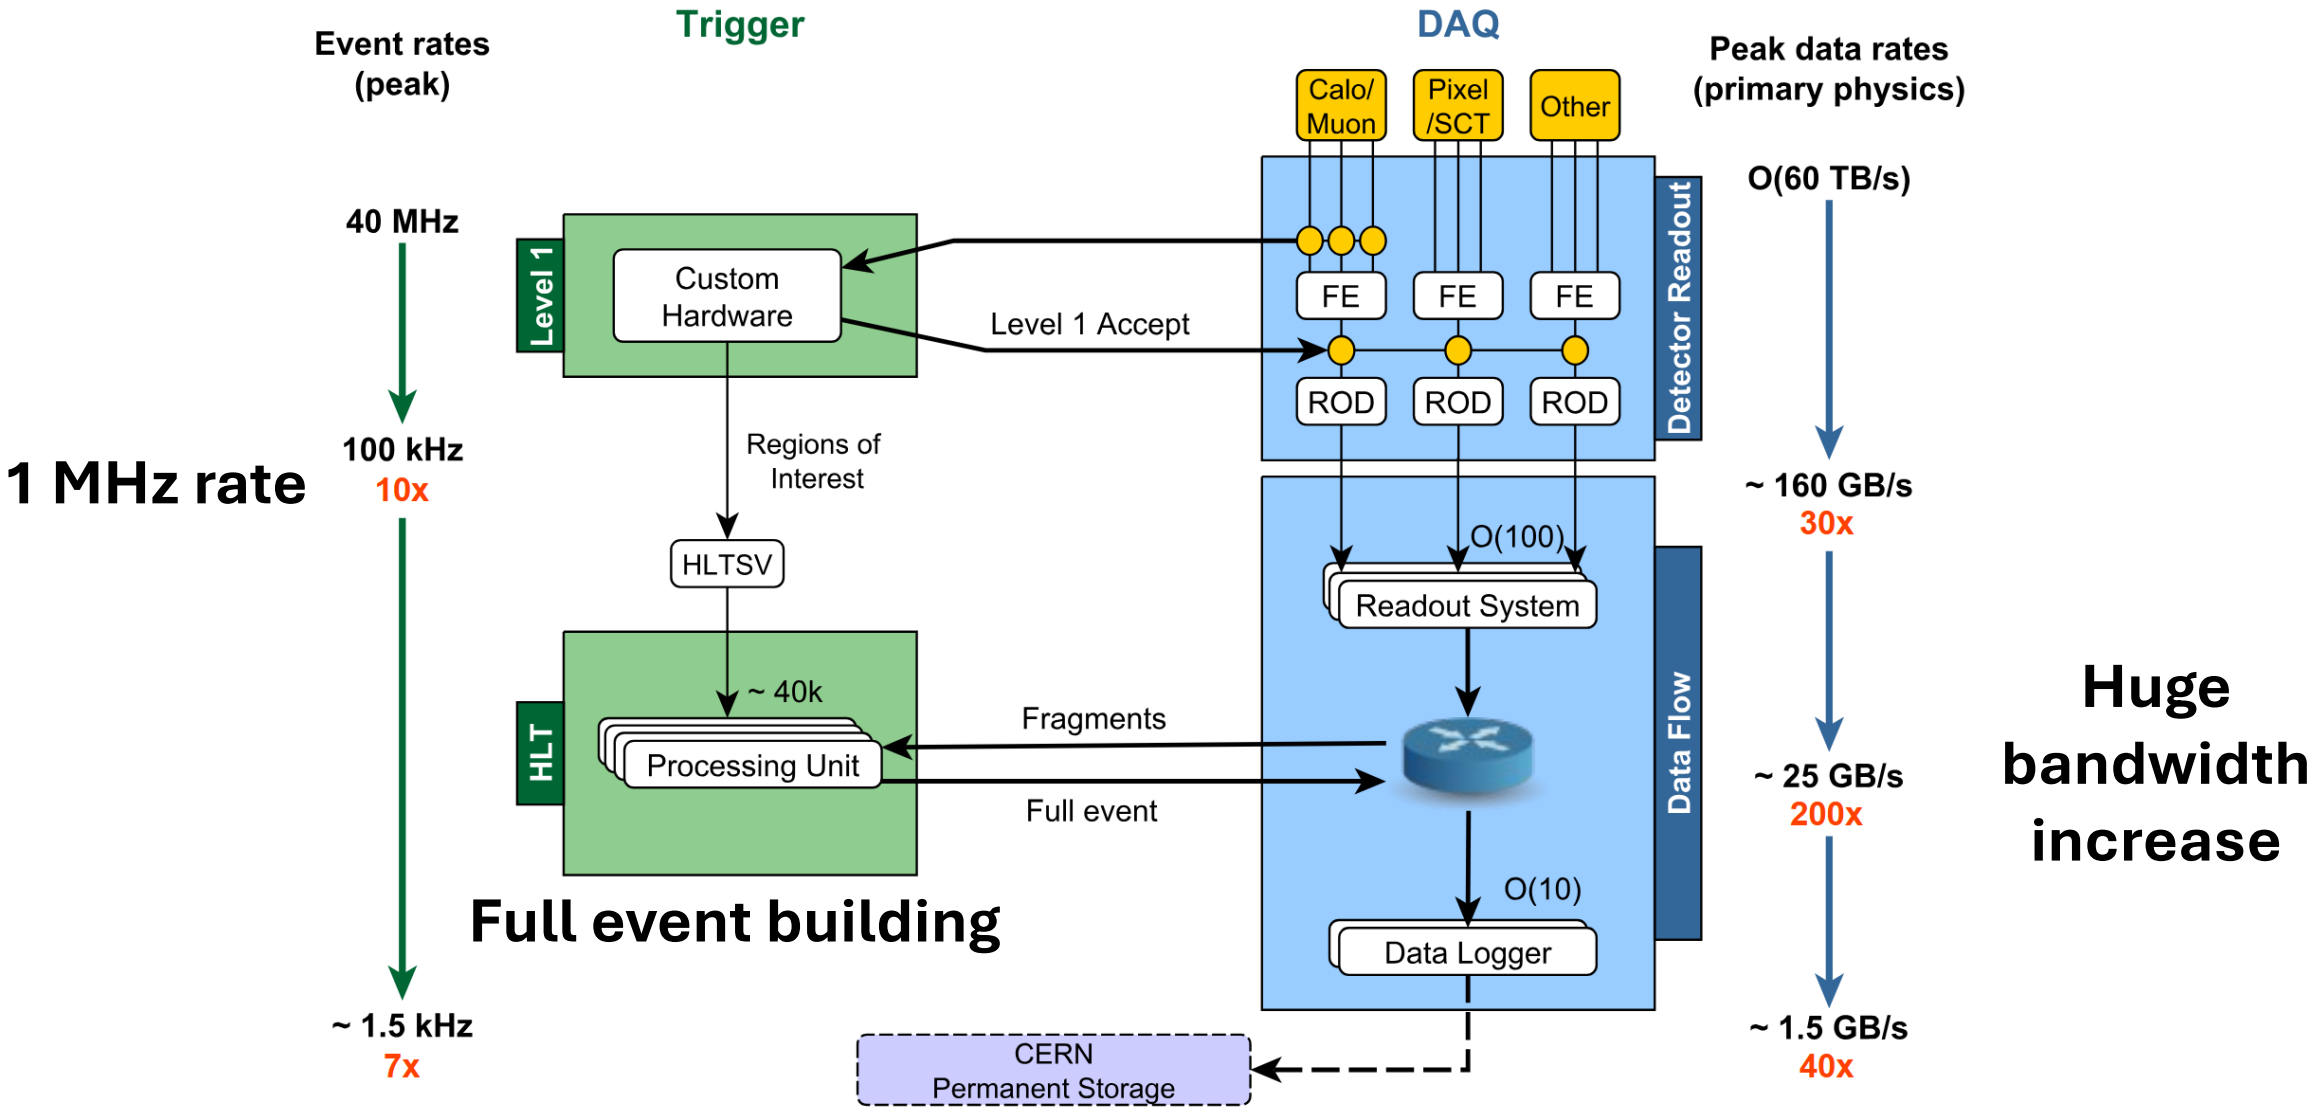

The PhD project

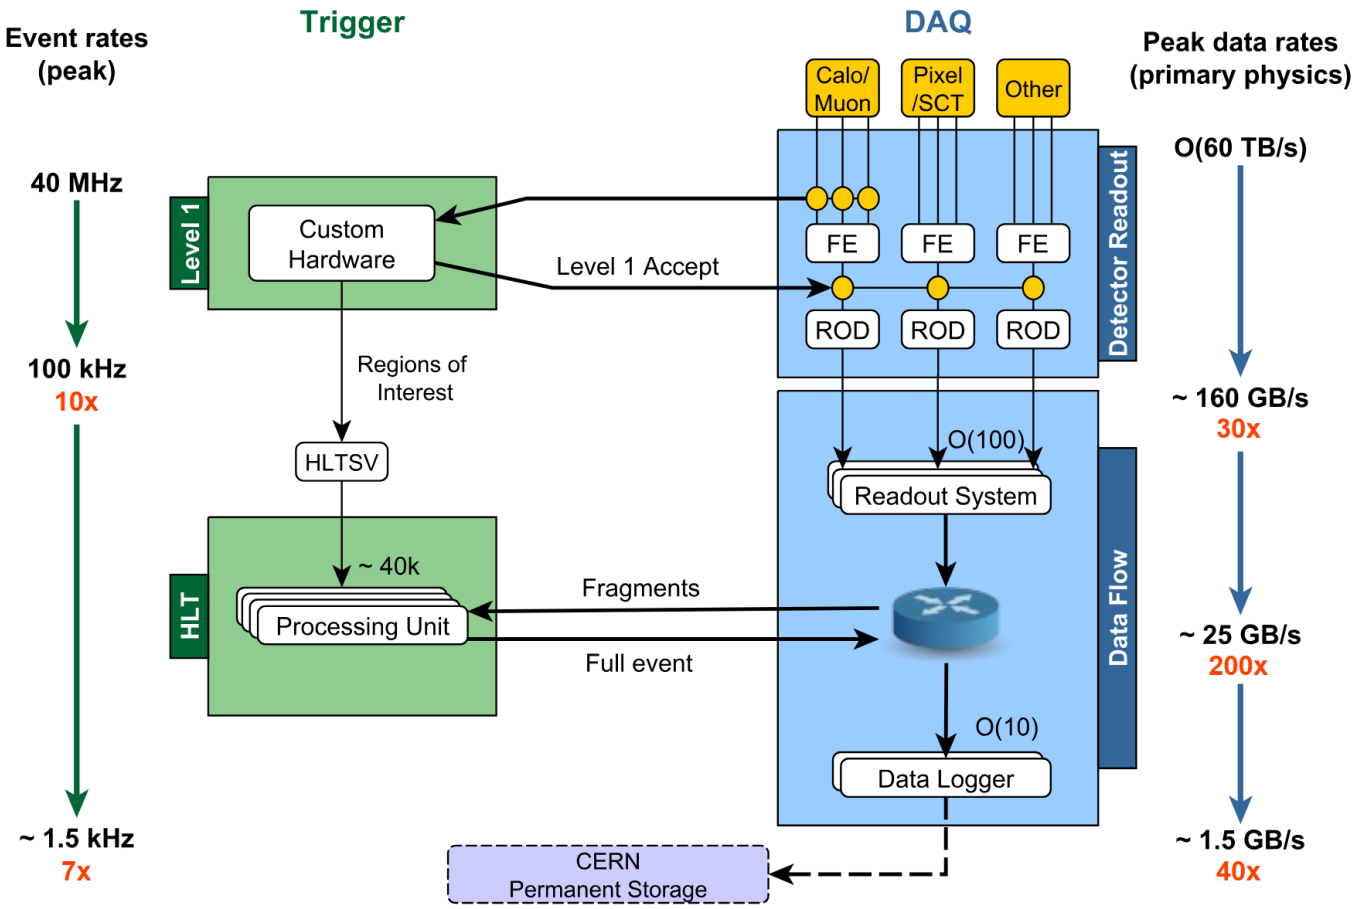

The PhD project

The PhD project mainly focuses on the ATLAS **Dataflow** Phase II upgrade.

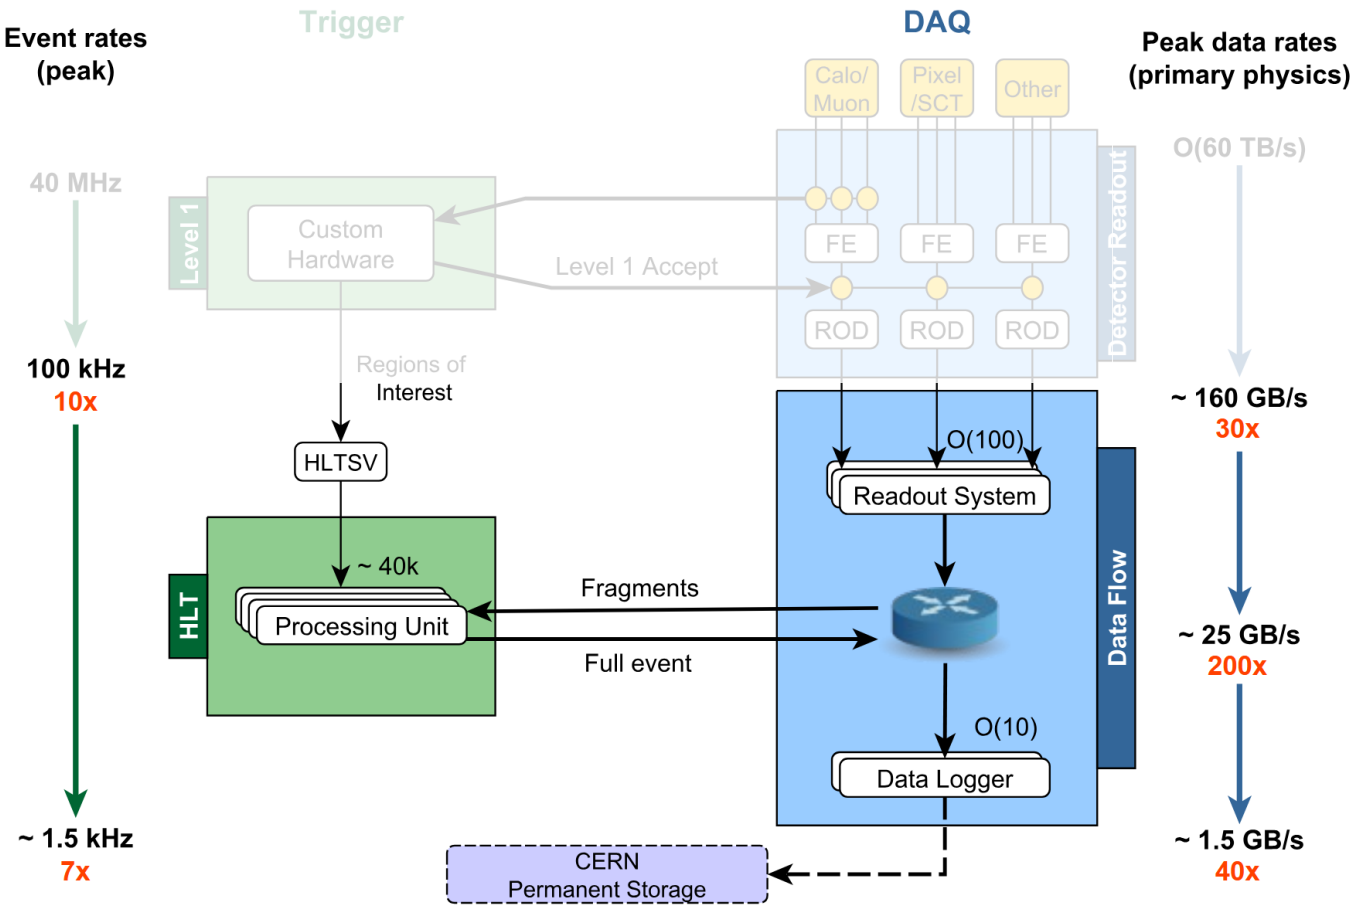

The PhD project

The PhD project mainly focuses on the ATLAS **Dataflow** Phase II upgrade.

The aim of the project is:

- **Assessment** of the current system capabilities
- **Design and development** of the Dataflow application for Phase II
- **Testing and integration** with the rest of TDAQ

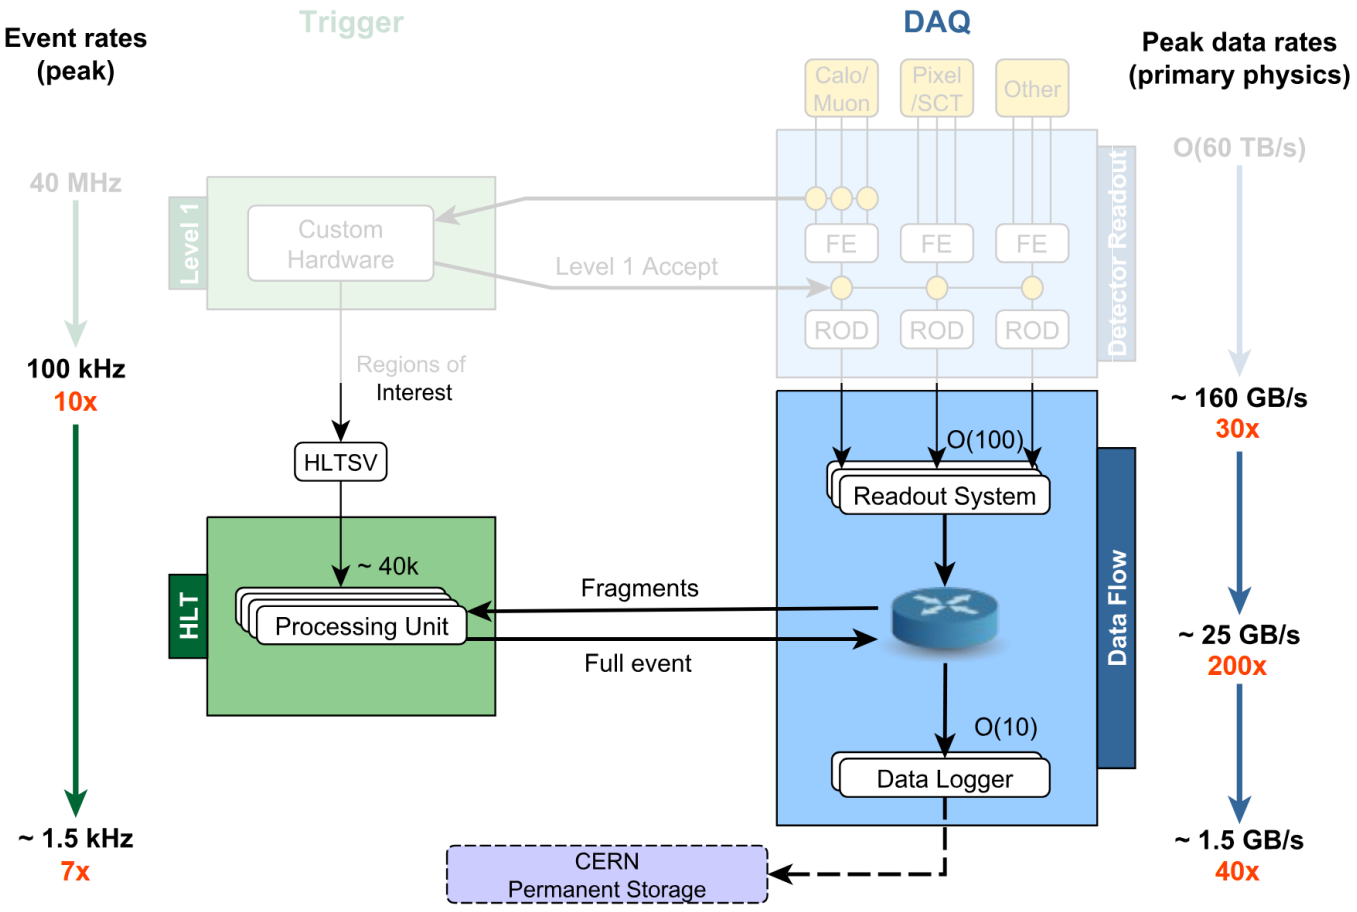

The PhD project: performance assessment

The current Dataflow system uses asynchronous programming* to manage data transfer.

The PhD project: performance assessment

The current Dataflow system uses asynchronous programming* to manage data transfer.

*Do you know what asynchronous programming is?

The PhD project: performance assessment

The current Dataflow system uses asynchronous programming* to manage data transfer.

- To **optimize** resources usage
- Chosen solution by big tech and social network companies in **datacenters**
- Can **maximize** event **processing rate** while using the least amount of processing power
- **Baseline** for the Phase II upgrade

*Do you know what asynchronous programming is?

The PhD project: performance assessment

Assessment started during last year.

- Current system not optimized enough
- **Ex.** The new Event Dispatcher:
  1 MHz assignment rate over 3000 nodes

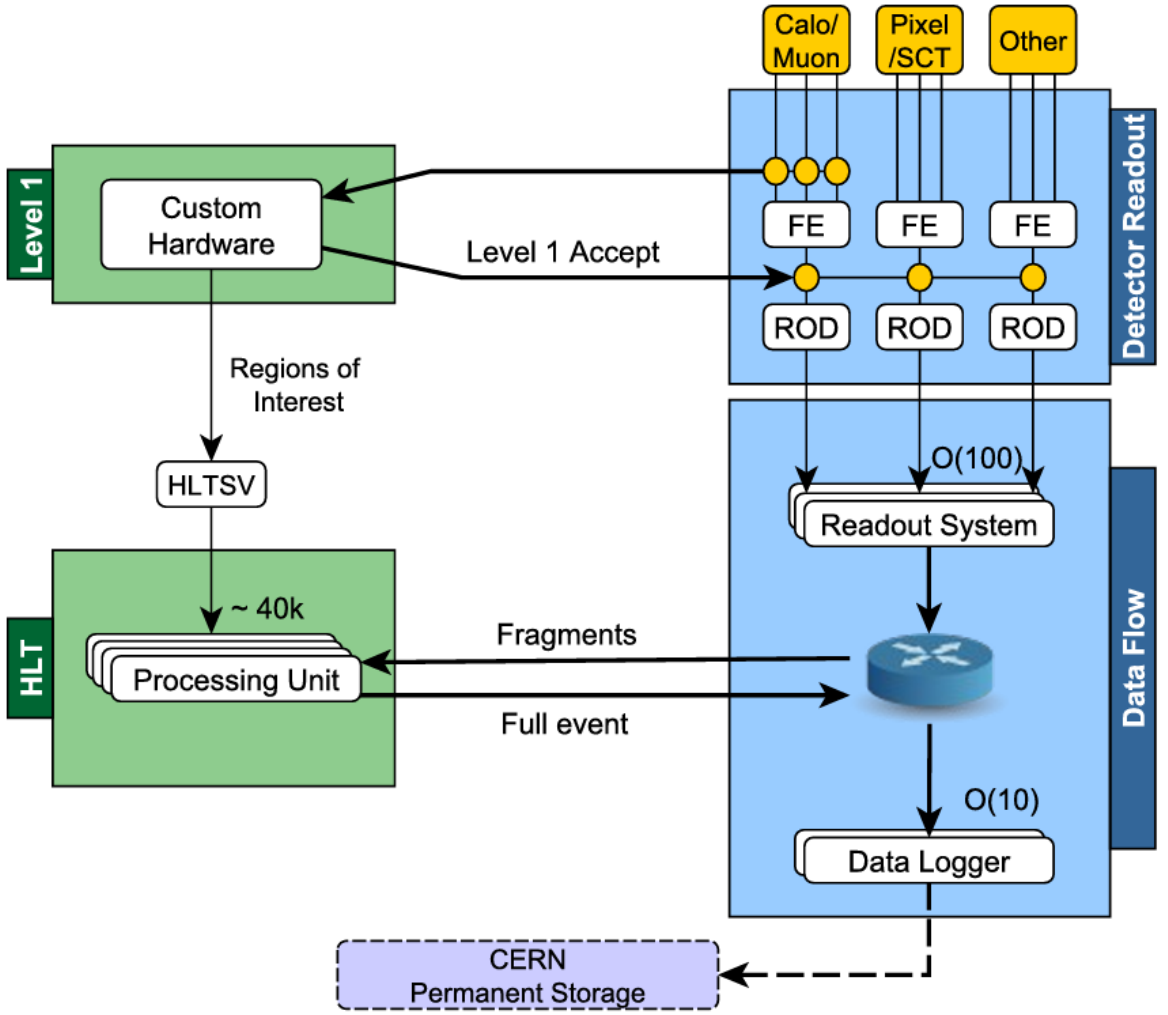

The PhD project: performance assessment

Assessment started during last year.

- Current system not optimized enough
- **Ex.** The new Event Dispatcher:
  1 MHz assignment rate over 3000 nodes
- Proposed changes effect:
  - Up to **2x** data rate increase
  - Up to **80%** CPU usage reduction

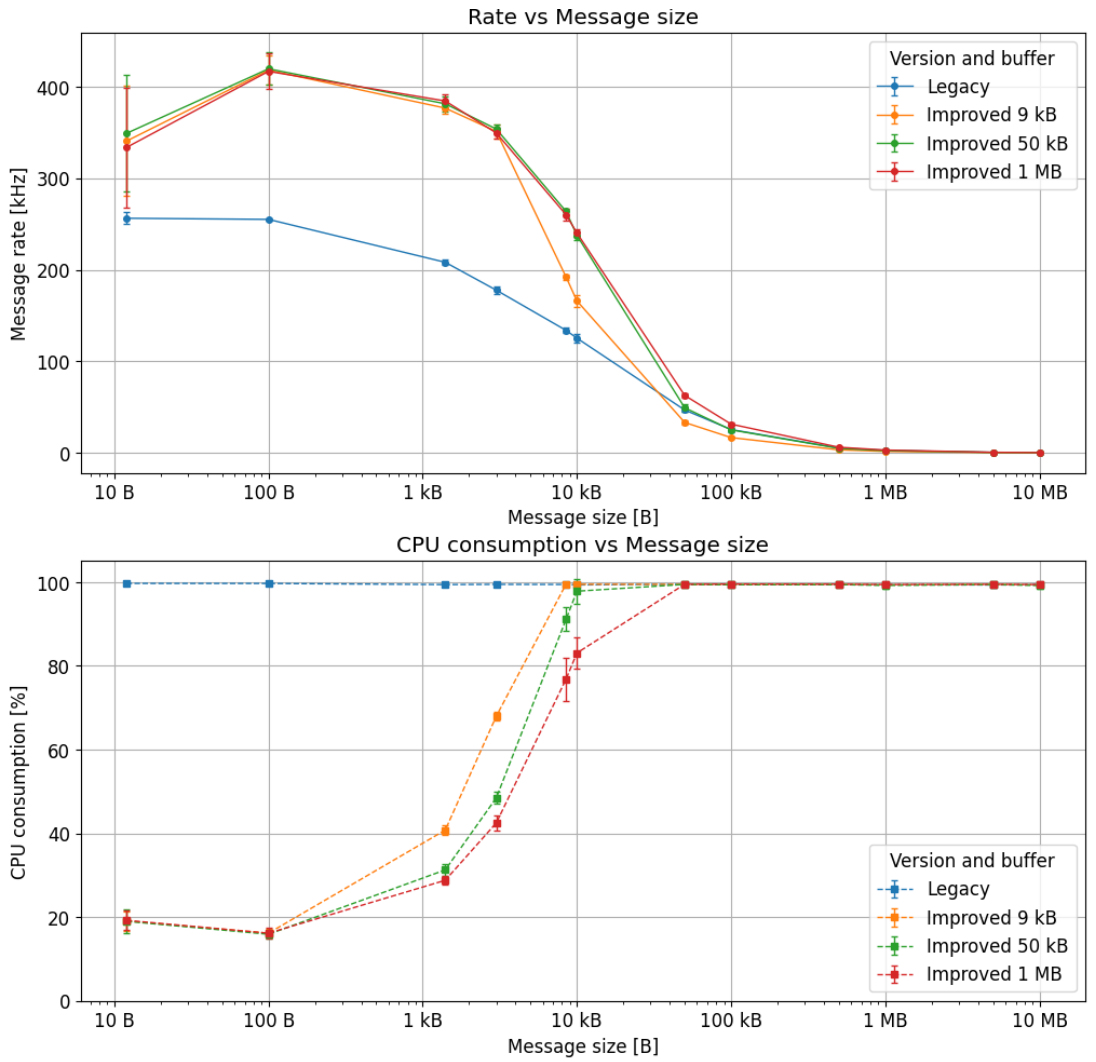

The PhD project: performance assessment

Assessment started during last year.

- Current system not optimized enough
- **Ex.** The new Event Dispatcher:
  1 MHz assignment rate over 3000 nodes
- Proposed changes effect:
  - Up to **2x** data rate increase
  - Up to **80%** CPU usage reduction
- **New applications needed** to get the performance boost

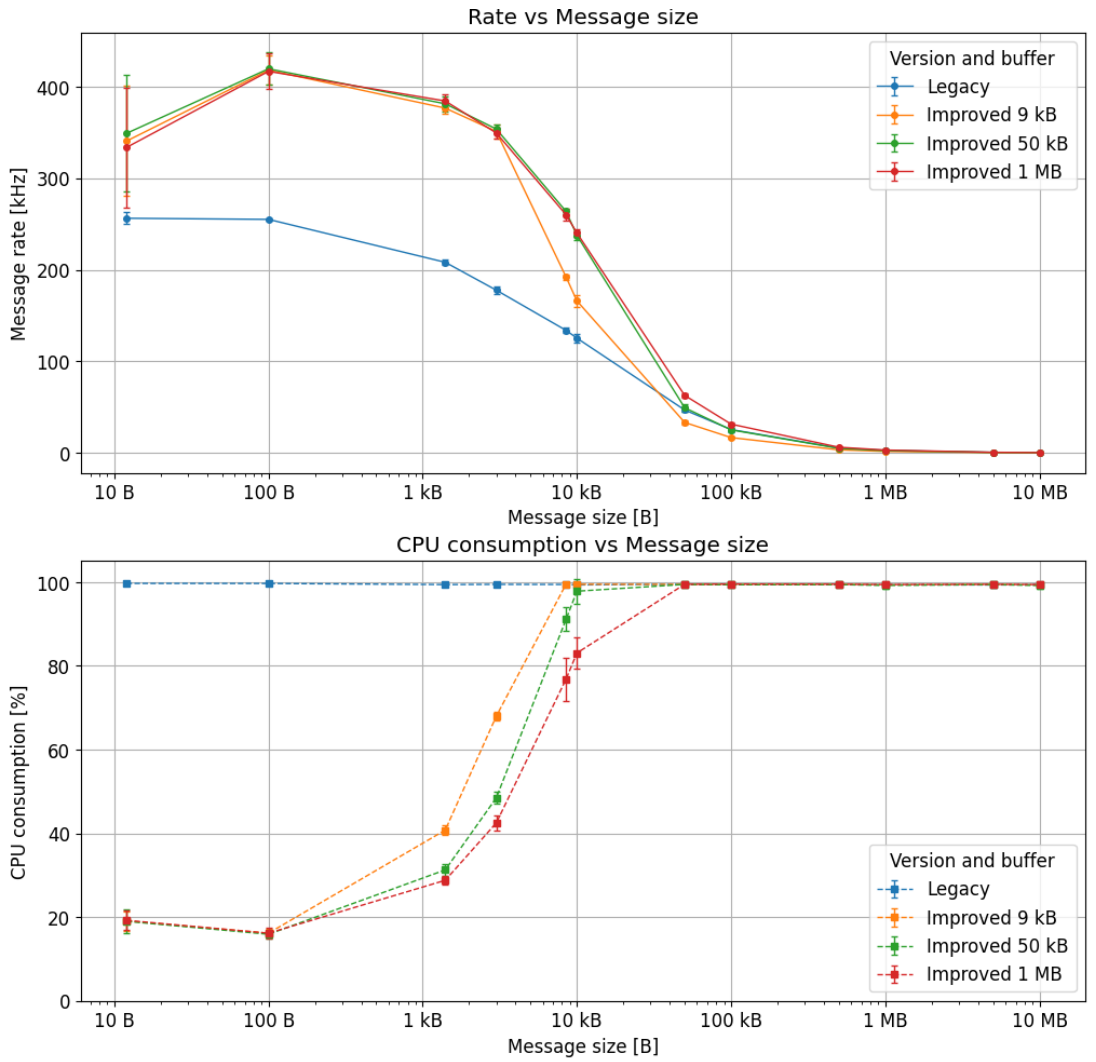

The PhD project: system development

The next year is focused on applications development.

The PhD project: system development

The next year is focused on applications development.

- New application are needed to reach the target performances within the limited budget

The PhD project: system development

The next year is focused on applications development.

- New application are needed to reach the target performances within the limited budget
- Research on cutting edge technologies used in industry: **R**emote **D**irect **M**emory **A**ccess, distributed file system

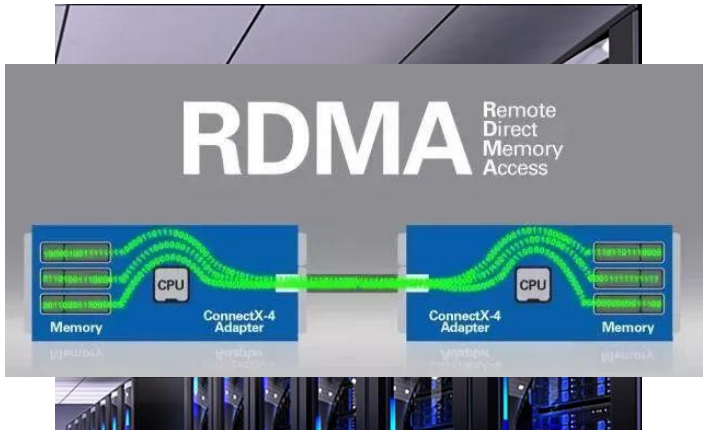

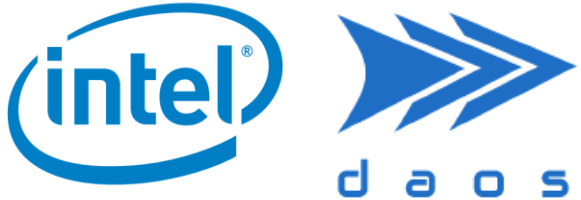

The PhD project: system development

The next year is focused on applications development.

- New application are needed to reach the target performances within the limited budget
- Research on cutting edge technologies used in industry: **R**emote **D**irect **M**emory **A**ccess, distributed file system
- Development of **ad-hoc solutions** to tailor functionalities on the available hardware and network

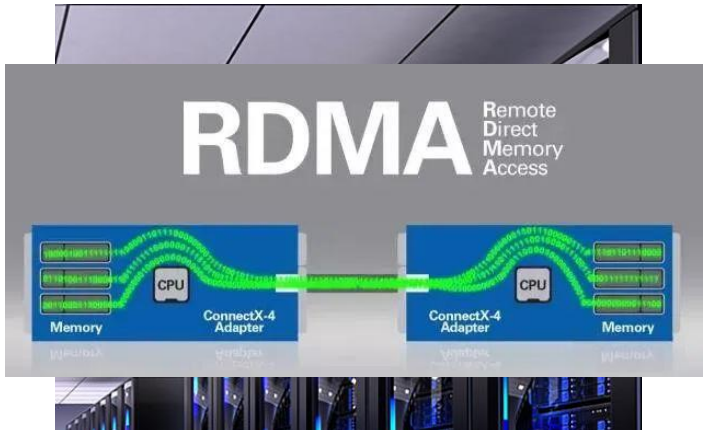

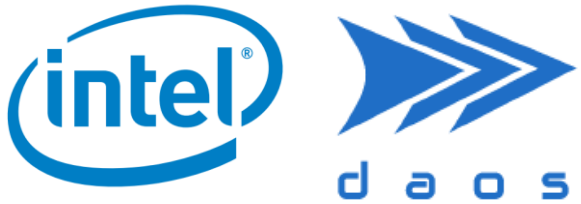

The PhD project: system development

The next year is focused on applications development.

- New application are needed to reach the target performances within the limited budget
- Research on cutting edge technologies used in industry: **R**emote **D**irect **M**emory **A**ccess, distributed file system
- Development of **ad-hoc solutions** to tailor functionalities on the available hardware and network
- Testing during the **Technical Stop** at the end of 2025 on the entire farm

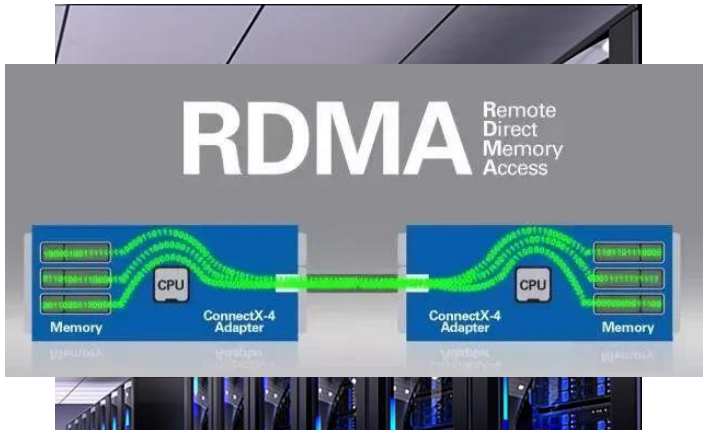

intel

daos

The PhD project: testing and integration

The last step is the system integration.

Dataflow

The PhD project: testing and integration

The last step is the system integration.

- Applications prototype needs to be **tested** in the upgraded TDAQ system, starting from **2027**

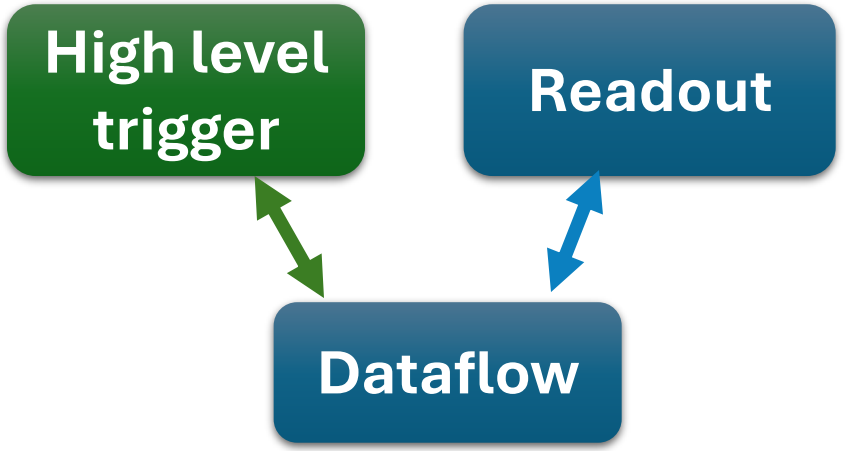

The PhD project: testing and integration

The last step is the system integration.

- Applications prototype needs to be **tested** in the upgraded TDAQ system, starting from **2027**
- Software **optimization** in the new farm, to reduce **resource usage**

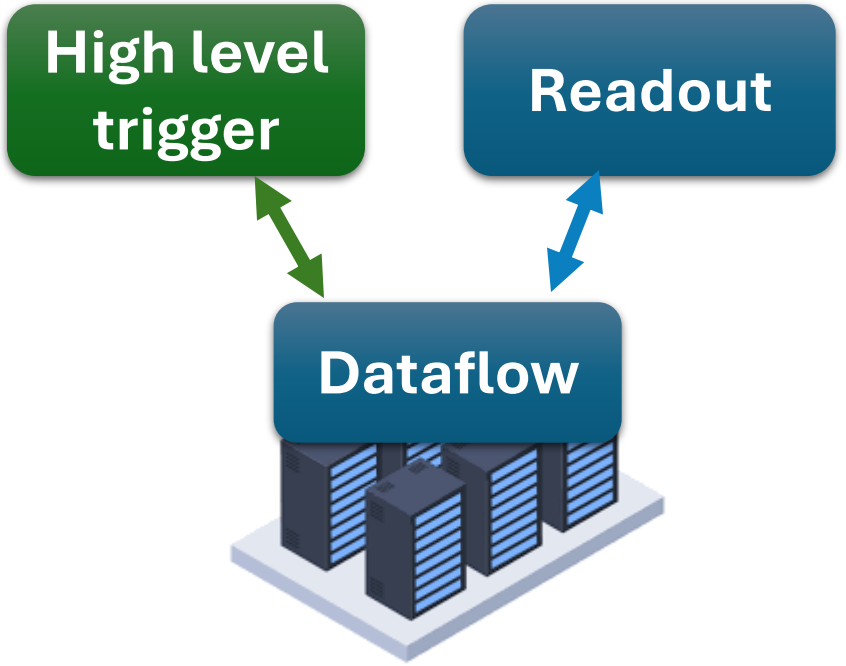

The PhD project: testing and integration

The last step is the system integration.

- Applications prototype needs to be **tested** in the upgraded TDAQ system, starting from **2027**
- Software **optimization** in the new farm, to reduce **resource usage**
- Extensive test for **reliability** and **robustness** to ensure maximum **efficiency**

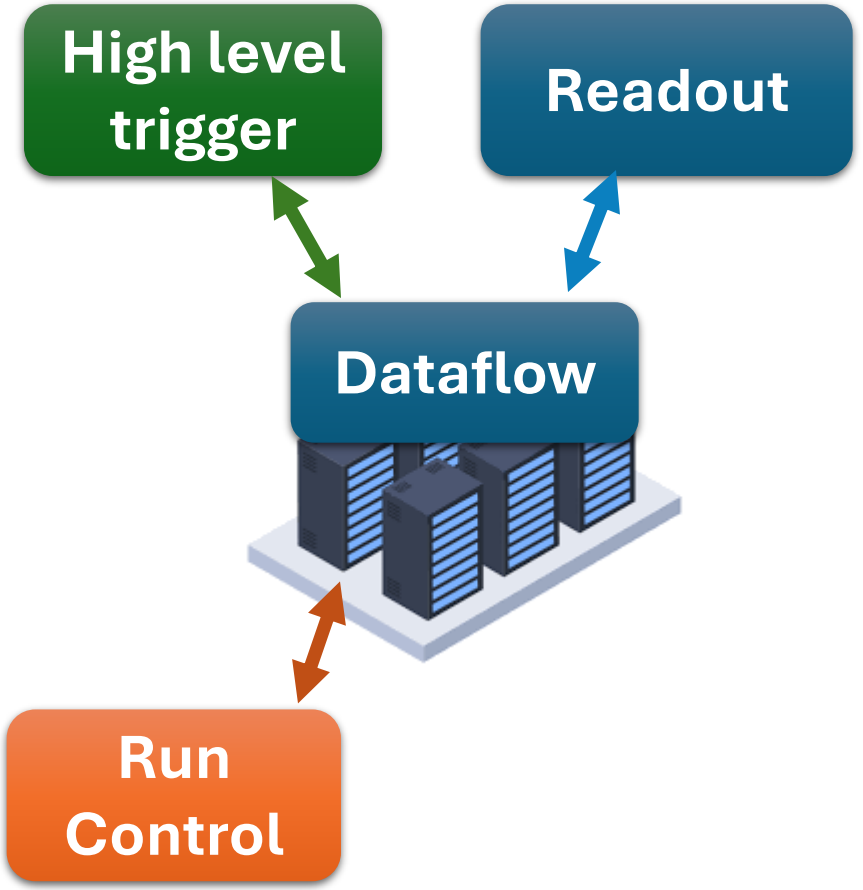

The PhD project: testing and integration

The last step is the system integration.

- Applications prototype needs to be **tested** in the upgraded TDAQ system, starting from **2027**
- Software **optimization** in the new farm, to reduce **resource usage**
- Extensive test for **reliability** and **robustness** to ensure maximum **efficiency**
- Integration with **monitoring** tools to provide early **problems detection**

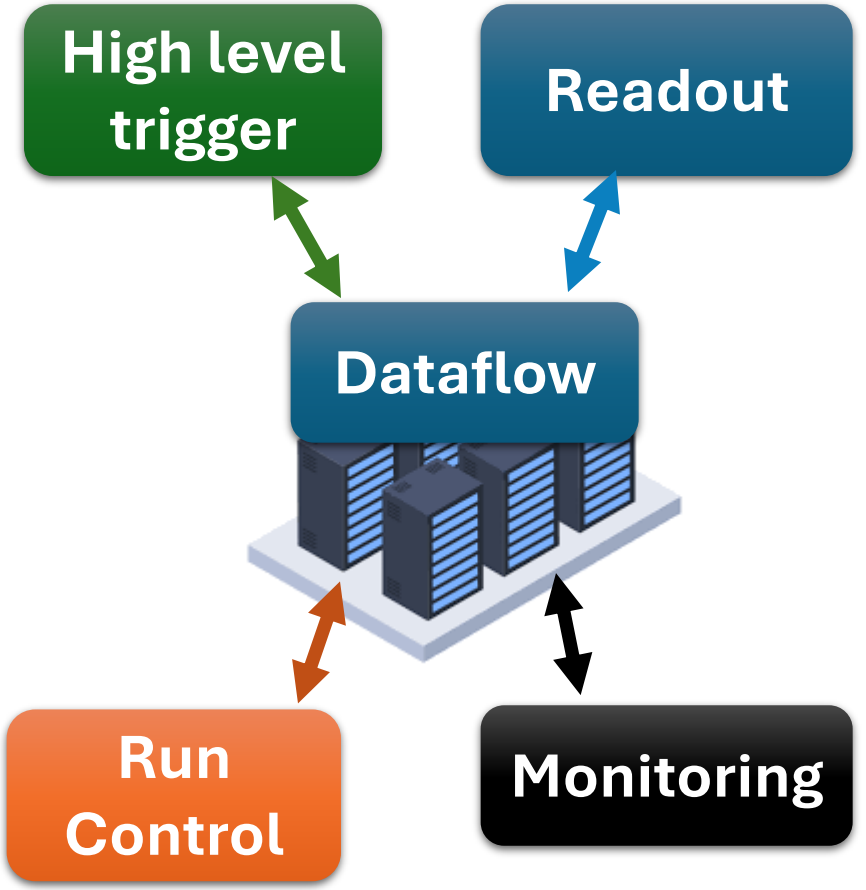

Conclusions

LHC will start its new **High Luminosity Phase**, for high precision physics measurement.

- **4x** peak luminosity, **3x** pile-up
- ATLAS experiment TDAQ needs a significant **upgrade**

The PhD project main focus is the new **Dataflow system**:

- Research **new solution** to meet the performance requirements
- Design the new system to ensure **performance**, **efficiency** and **reliability**

The other PhD focus will be in the **muon data quality** to support the online **operations** for this Run and beyond

Backup slides

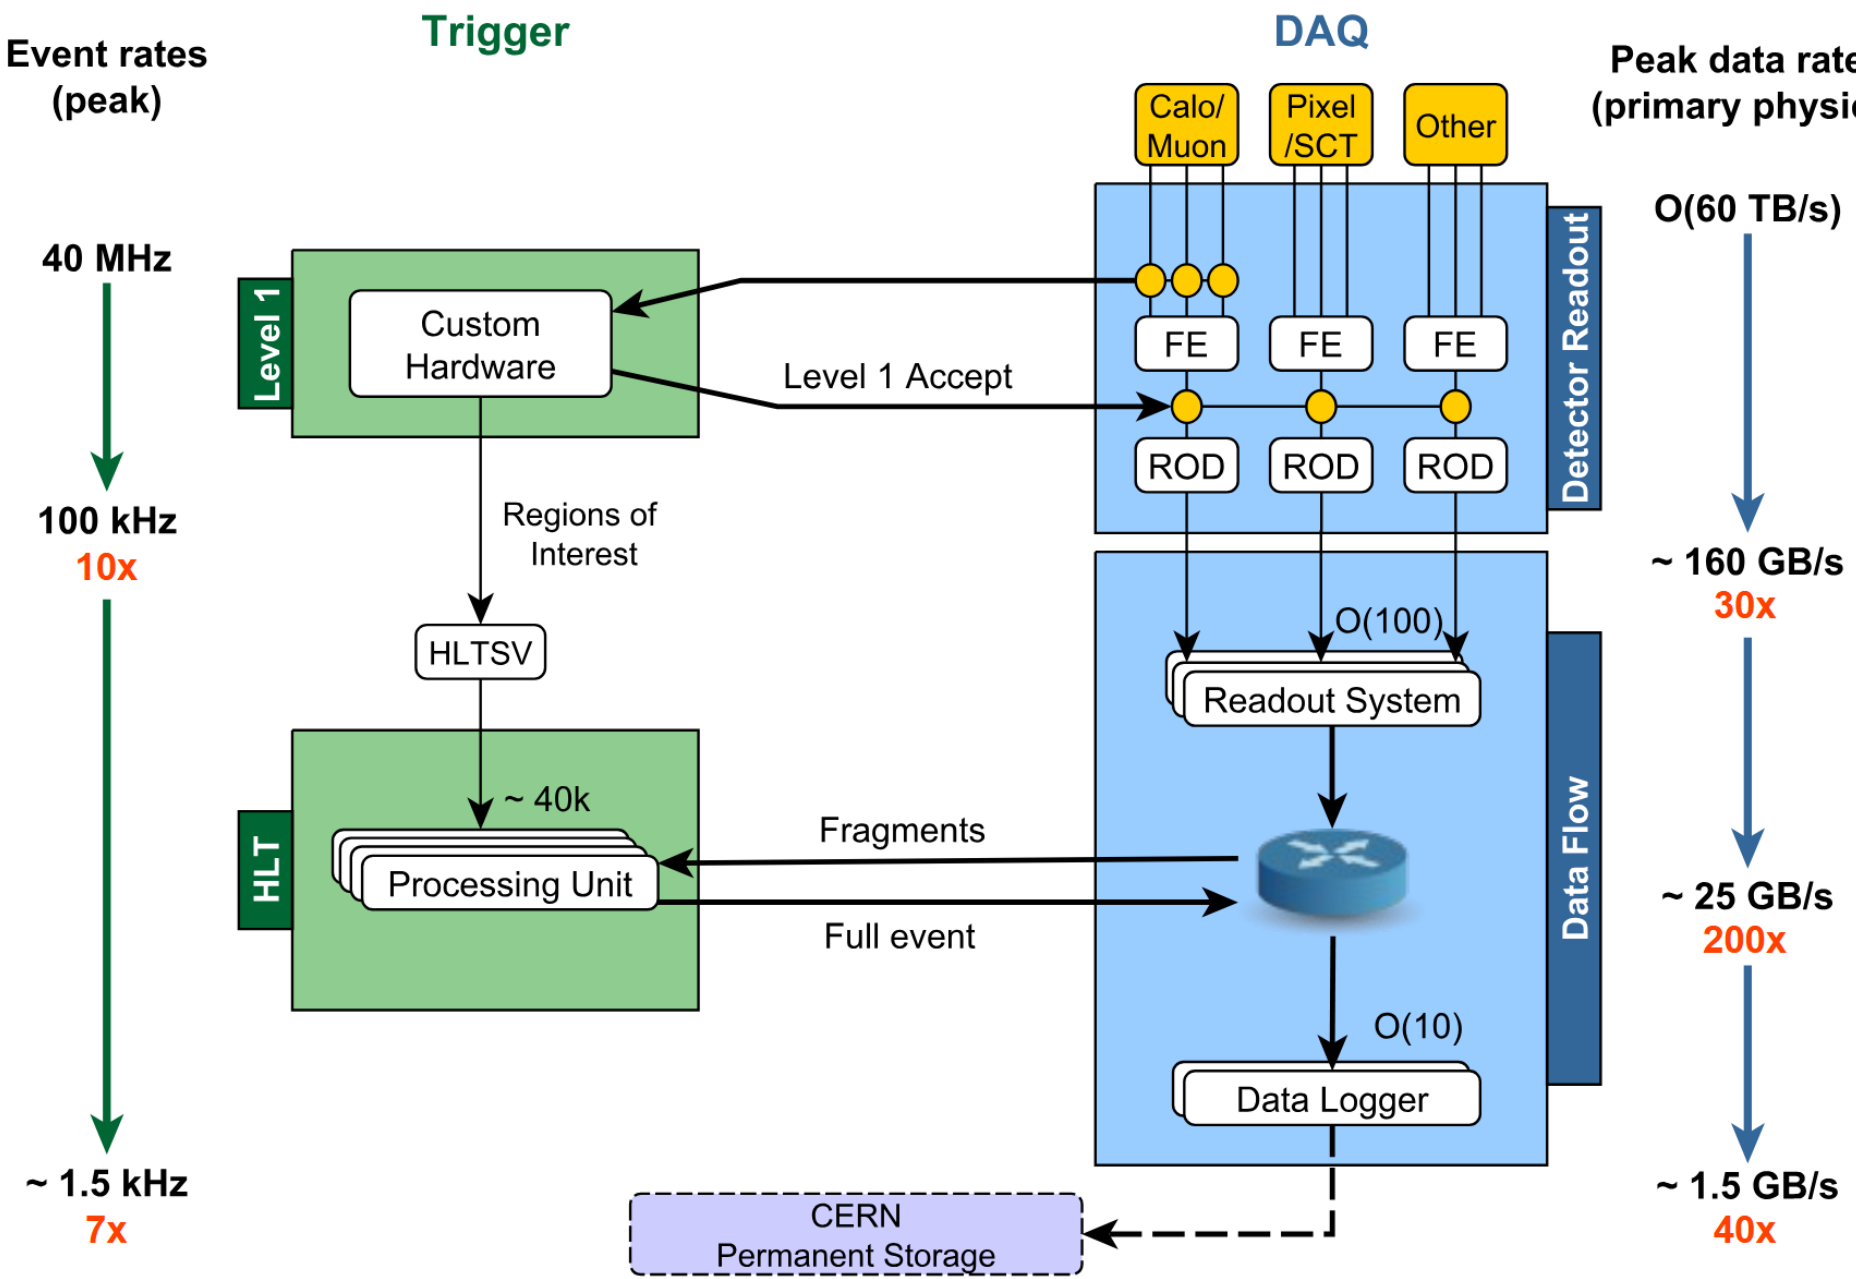
Event rates
(peak)
Trigger
DAQ
Peak data rates
(primary physics)
Calo/
Muon
Pixel
/SCT
Other
40 MHz
O(60 TB/s)
Level 1
Custom
Hardware
Level 1 Accept
FE
FE
FE
Detector Readout
ROD
ROD
ROD
Regions of
Interest
100 kHz
10x
~ 160 GB/s
30x
HLTSV
O(100)
Readout System
~ 40k
HLT
Processing Unit
Fragments
Full event
Data Flow
~ 25 GB/s
200x
O(10)
Data Logger
~ 1.5 kHz
7x
CERN
Permanent Storage
~ 1.5 GB/s
40x

Buffer size evaluation

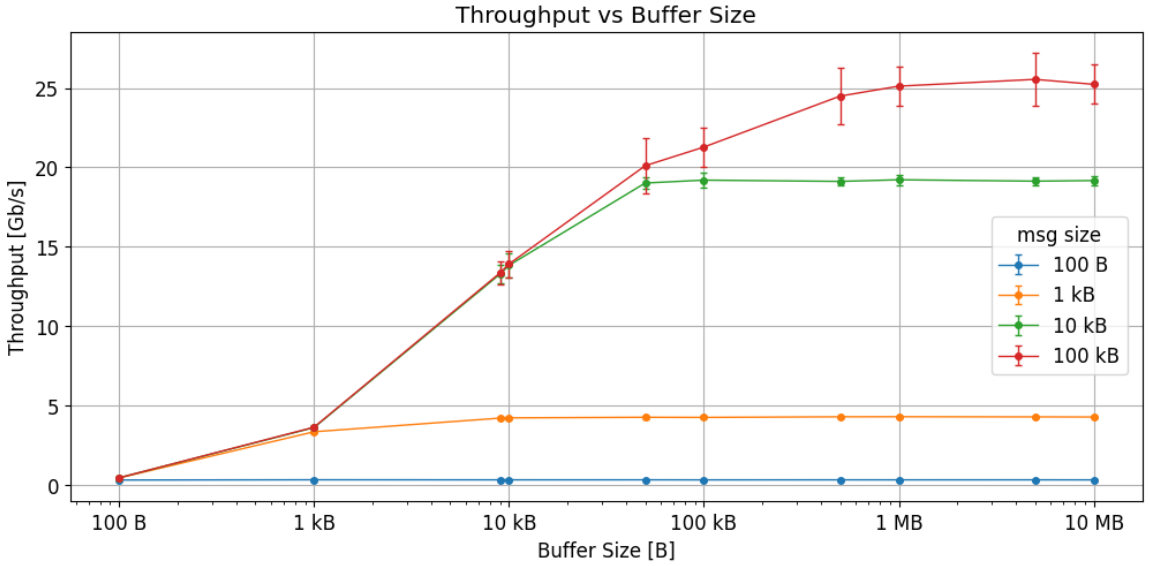

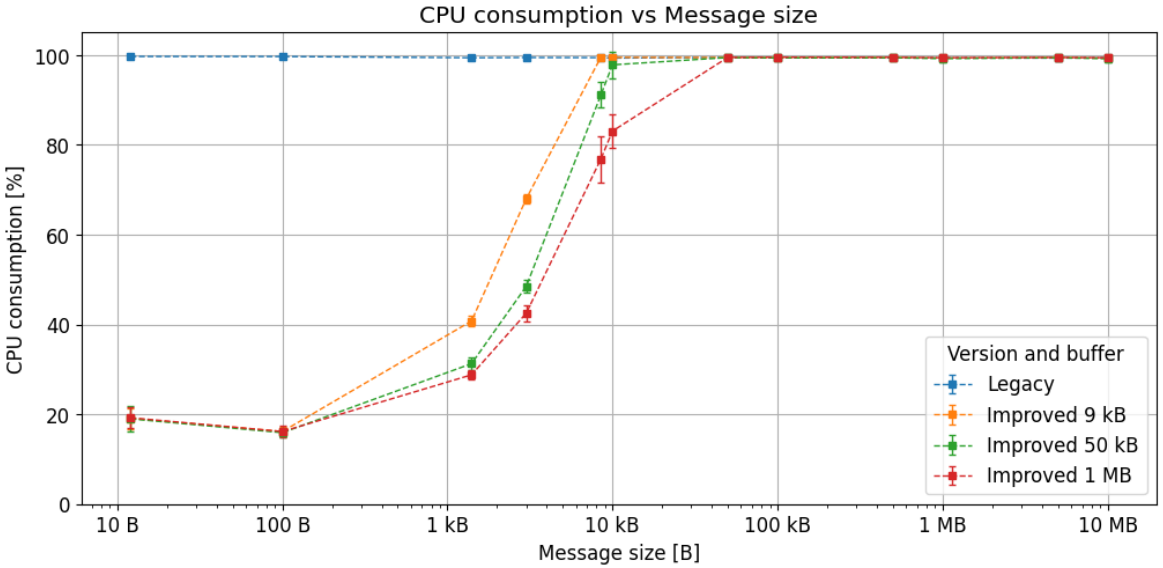

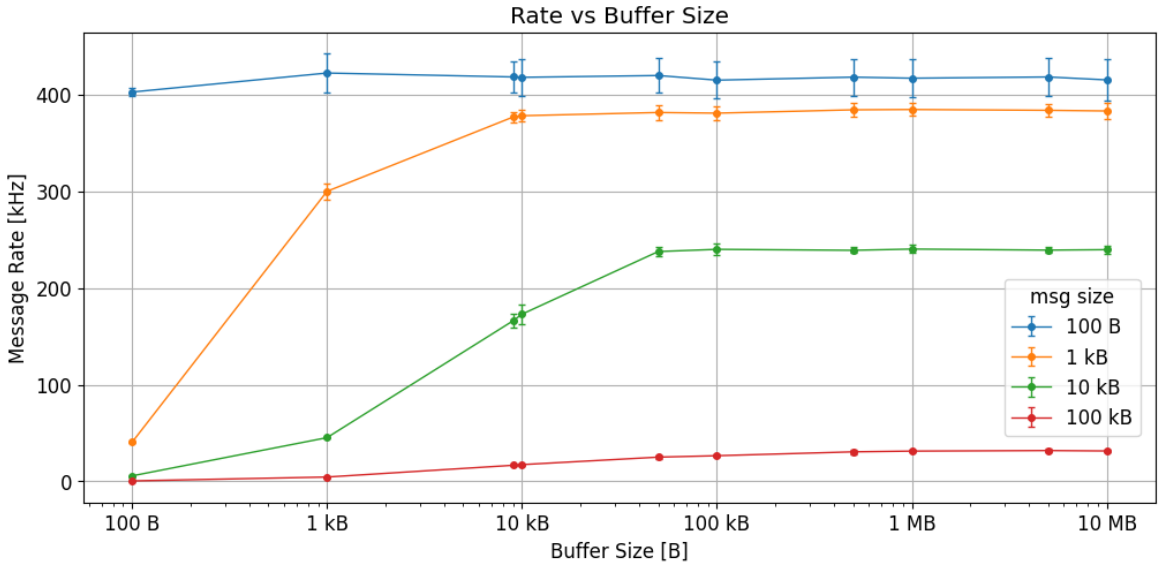

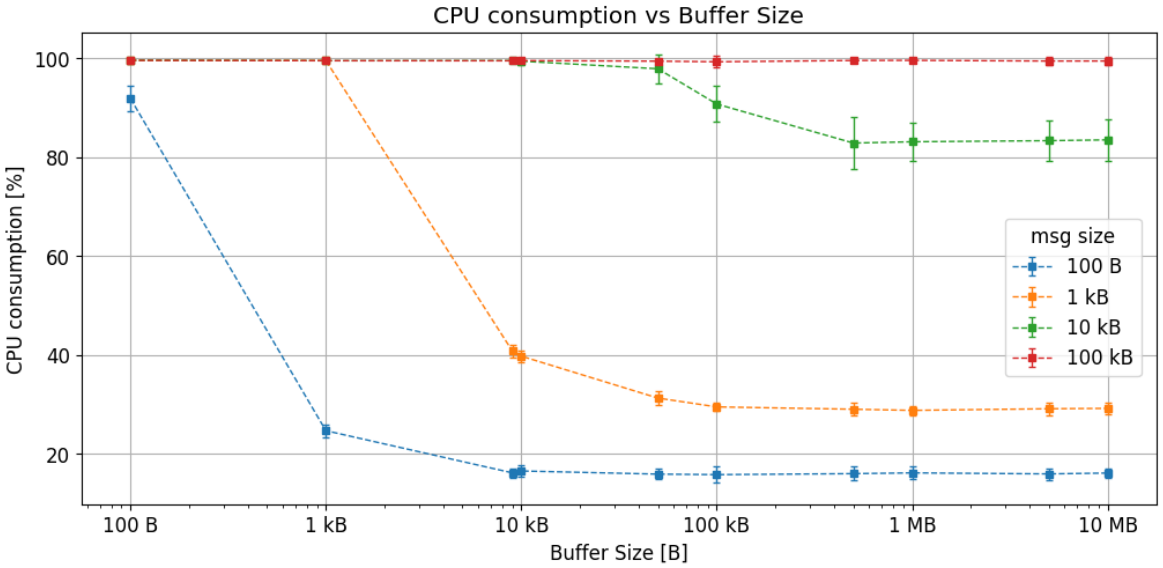

One-to-one performance tests spec

Test conditions:

- 1 connection, 1 thread
- Single test duration: 30 s
- ~ 30 tests per configuration
- Sender/receiver specification:
  AMD EPYC 7302P 16-Core Processor @ 3.1 GHz,
  32 GB RAM(send),
  128 GB RAM (recv),
  Single port Nvidia ConnectX-5 100 Gb/s NIC

Test results includes:

- Message rate
- Data bandwidth
- Total CPU usage (usr + sys)
- Computed average and deviation

Event Dispatcher / Event Builder tests

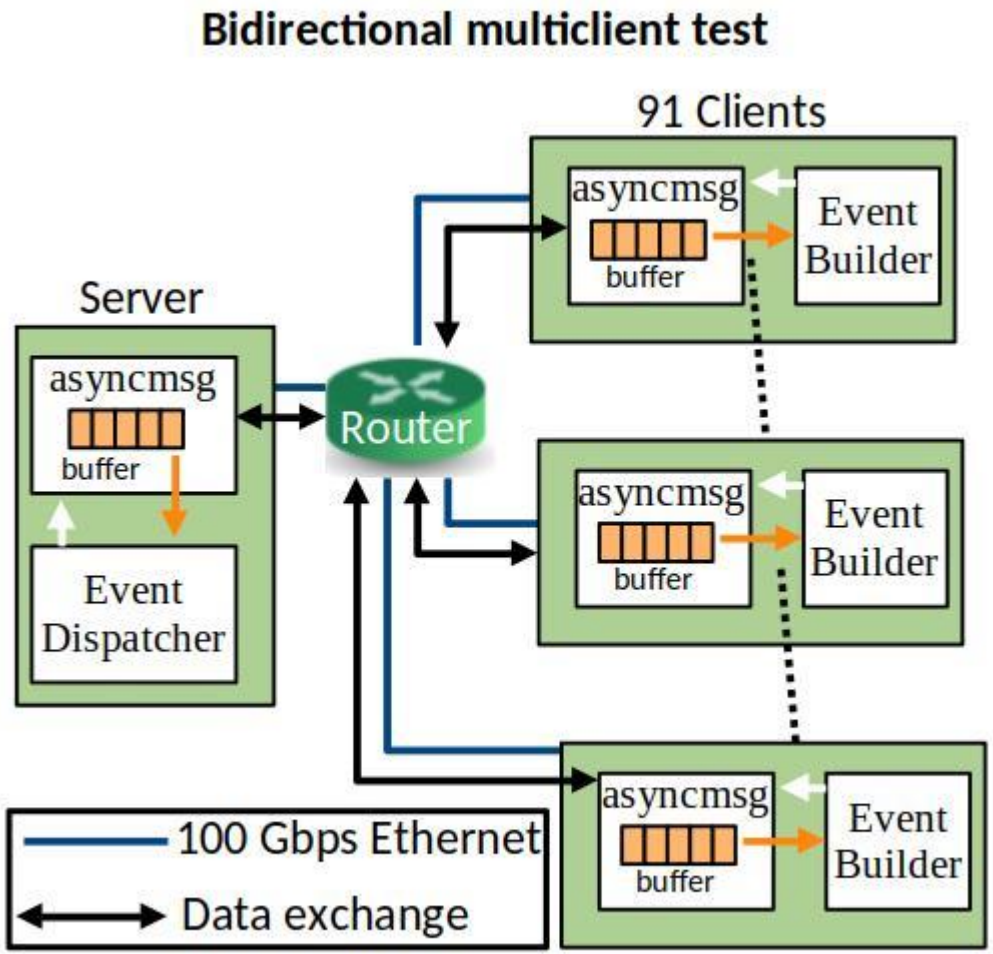

Event Dispatcher test

- Single Event Dispatcher sending
- 91 Event Builder receiving

Test metrics:

- Event Dispatcher Assignment rate
- Event Dispatcher CPU usage

Test parameters:

- Event Dispatcher threads
- Buffer size (new version)

Event Dispatcher / Event Builder tests

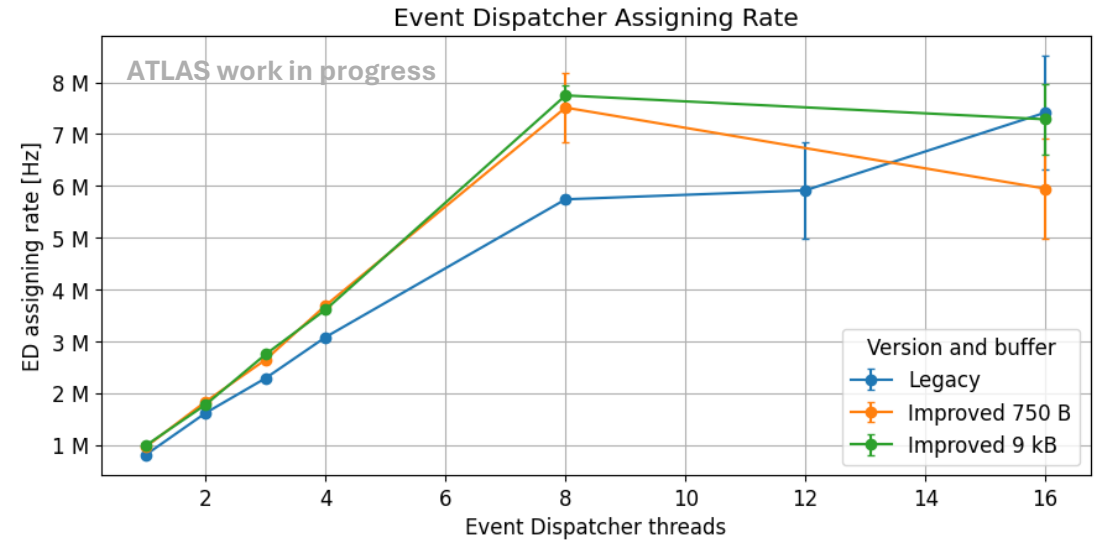

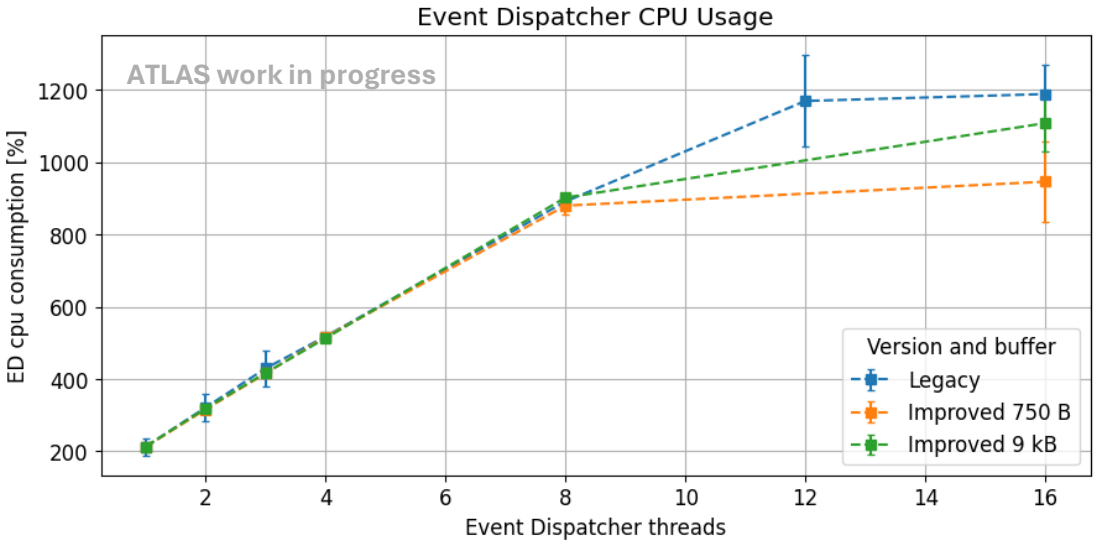

Assignment Rate

- New version always performed better
- Legacy version 810 kHz with 1 thread, 1.6 MHz with 2
- New version 970 kHz with 1, 1.83 MHz with 2

ED CPU Consumption

- Same until 8 threads, linear scaling
- With more threads, new version shows CPU reduction