DATA ACQUISITION
AT THE HIGH-LUMINOSITY LHC

Tommaso Colombo
CERN

CERN openlab Technical Workshop
4 March 2025

FINDING NEEDLES IN **HAYSTACKS**

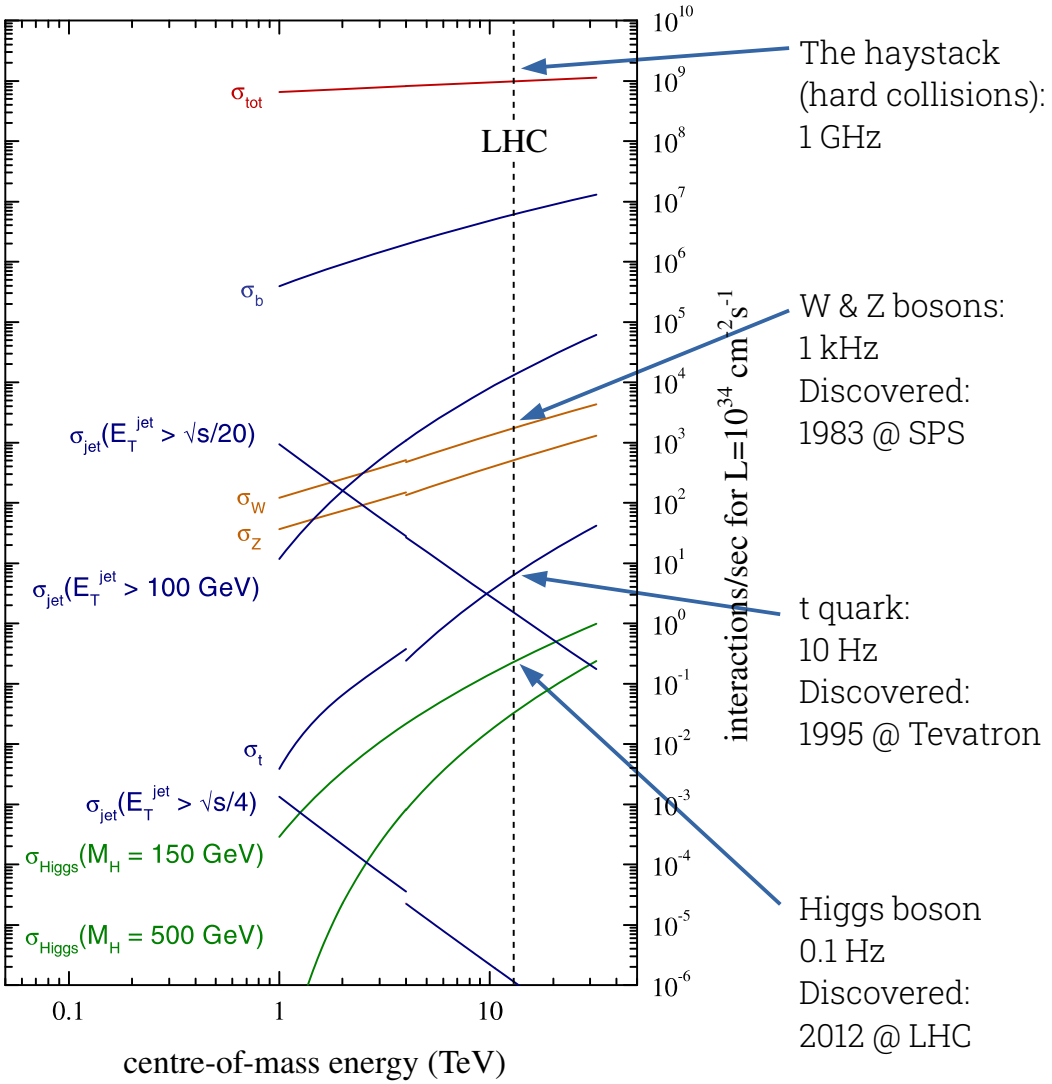

The haystack (hard collisions): 1 GHz

W & Z bosons: 1 kHz Discovered: 1983 @ SPS

t quark: 10 Hz Discovered: 1995 @ Tevatron

Higgs boson 0.1 Hz Discovered: 2012 @ LHC

- Contemporary particle physics: **rare** processes
- The vast majority of the collisions is “boring”:
  - They only result in particles and processes that were studied to death decades ago
  - Interesting physics is ≥ 9 orders of magnitude rarer: ≥ one in a billion

HIGH-LUMINOSITY LHC UPGRADE

- High-luminosity LHC,
  aka HL-LHC, aka HiLumi
- Luminosity:
  physicist-speak for
  number of collisions
- More collisions =
  → more “rare” phenomena
  → happy physicists!

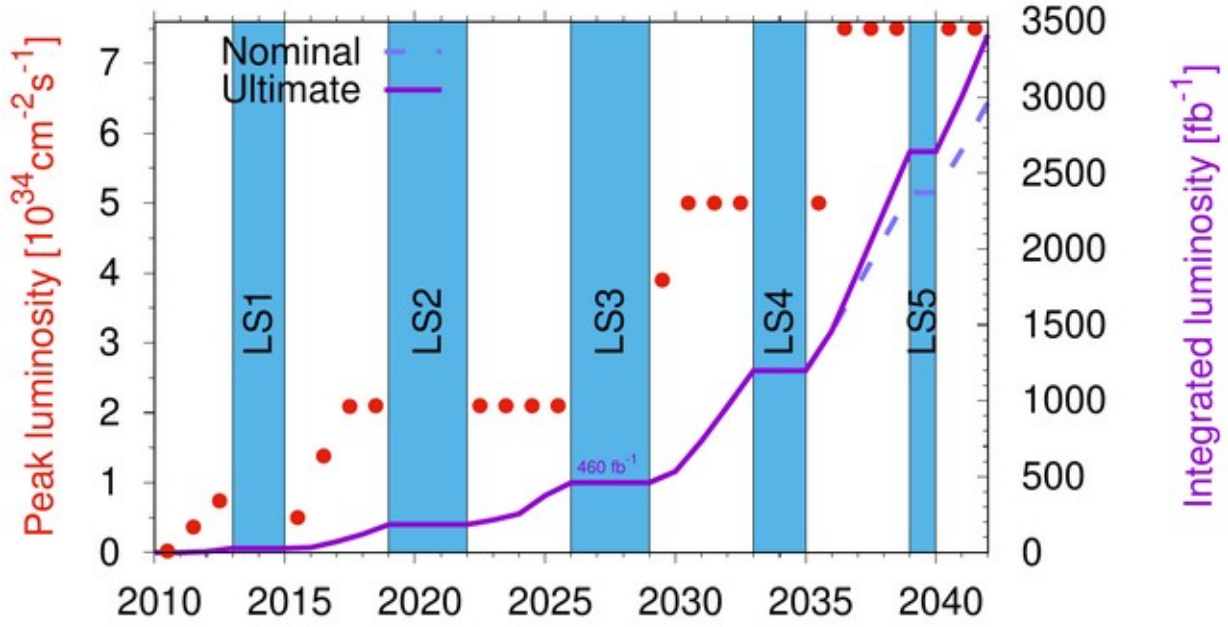

- Goal: deliver 6-8 times more collisions in 11 years (2030-2041)
  than “original” LHC in 16 years (2010-2026)

HIGH-LUMINOSITY LHC UPGRADE

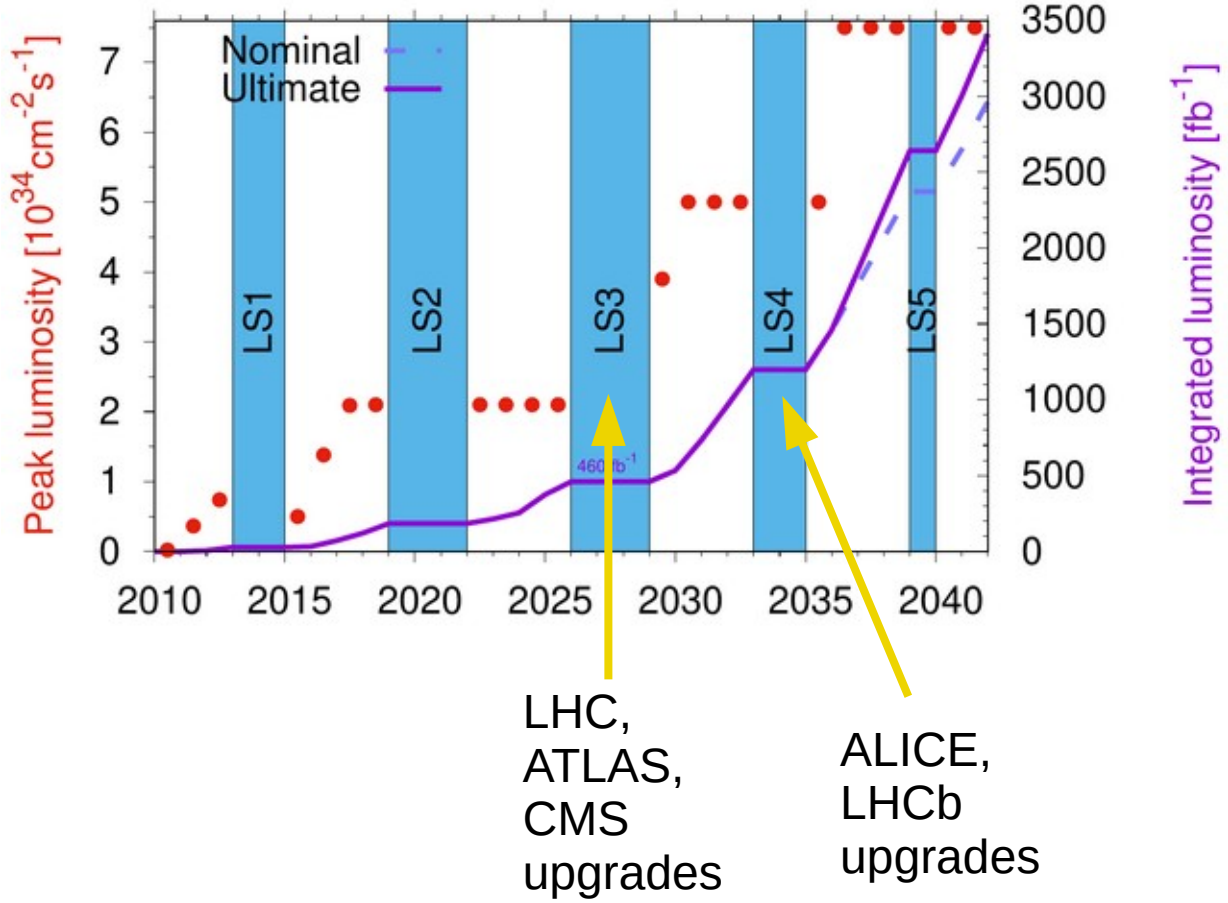

THE NEEDLE

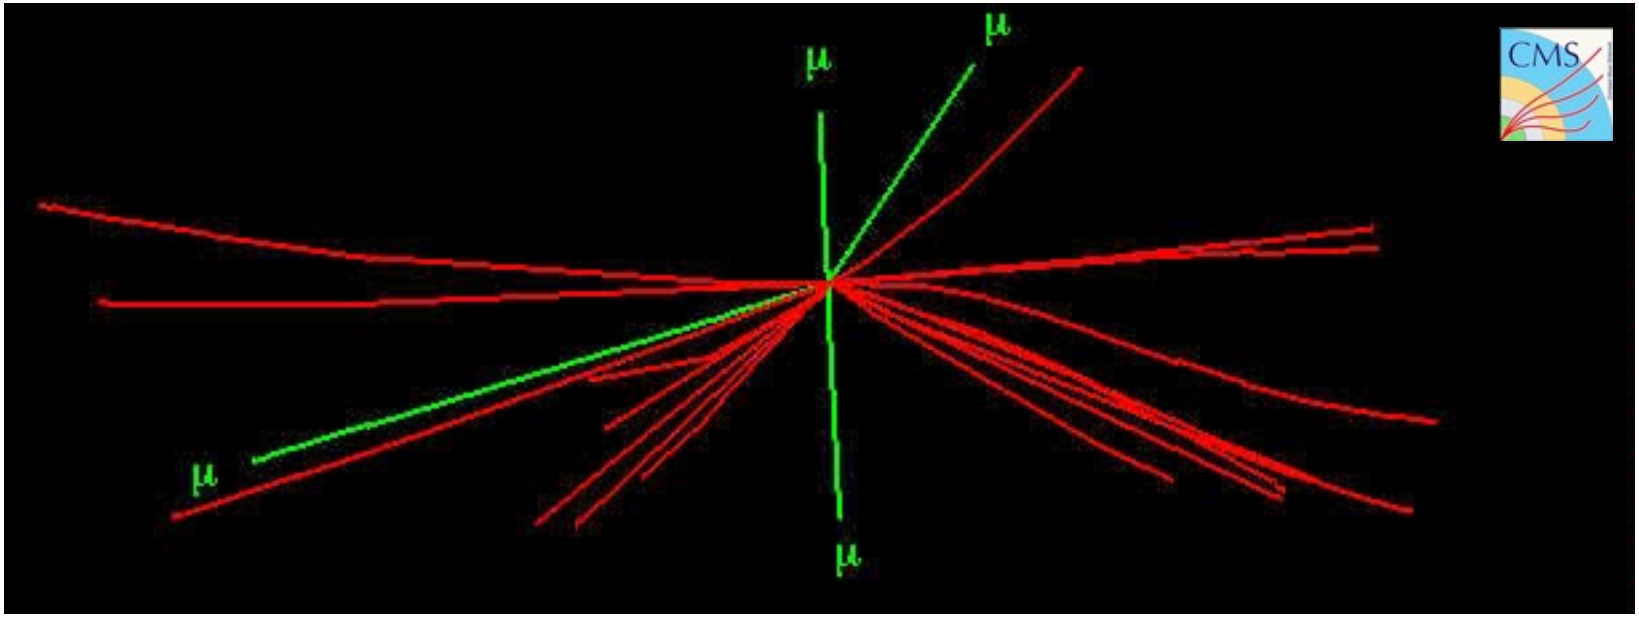

- This is what we're looking for:
  a Higgs boson decaying in four easily identifiable muons
- LHC produces a few of these **per day**, HL-LHC will increase 6-8x

THE HAYSTACK

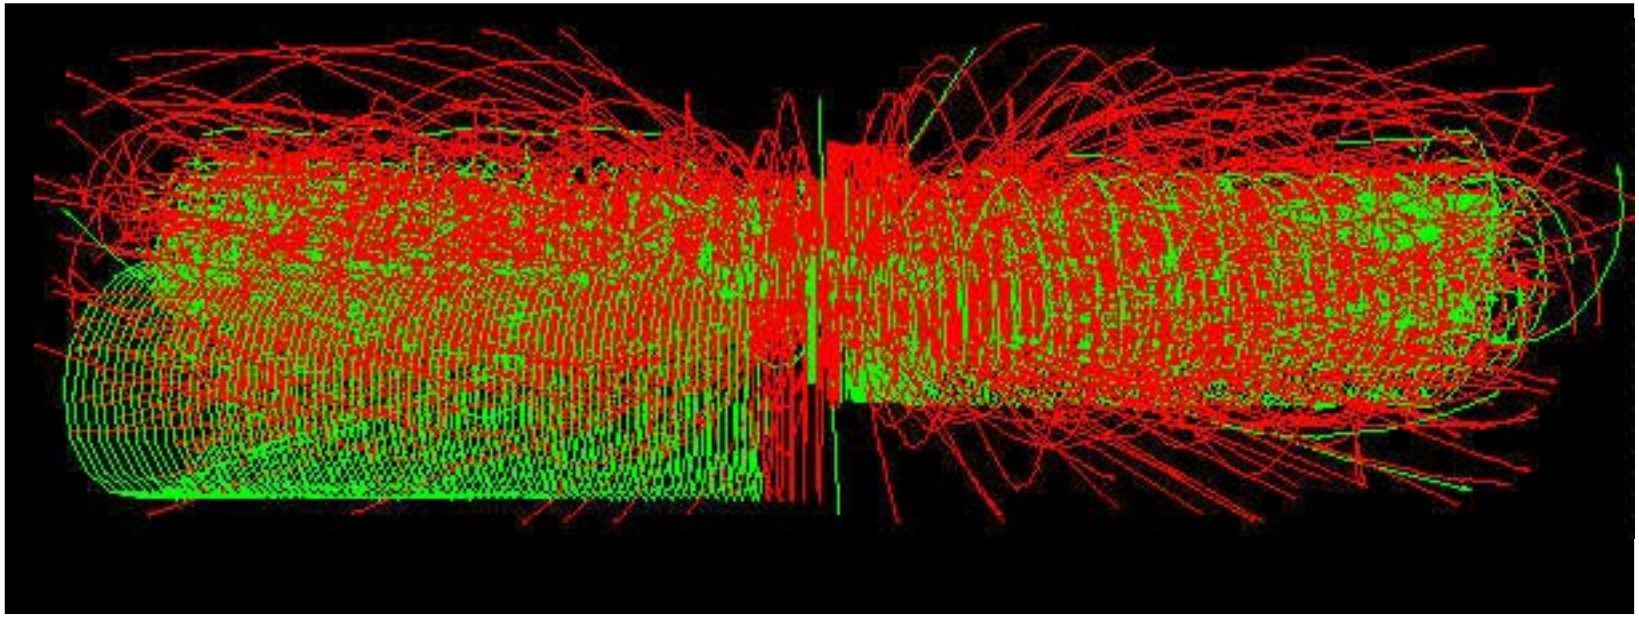

- This is where it hides:
  ~ 60 (LHC) – 200 (HL-LHC) other collisions producing 1000s of particles
- The LHC makes 40 million of these per second!

DEALING WITH THE HAYSTACK

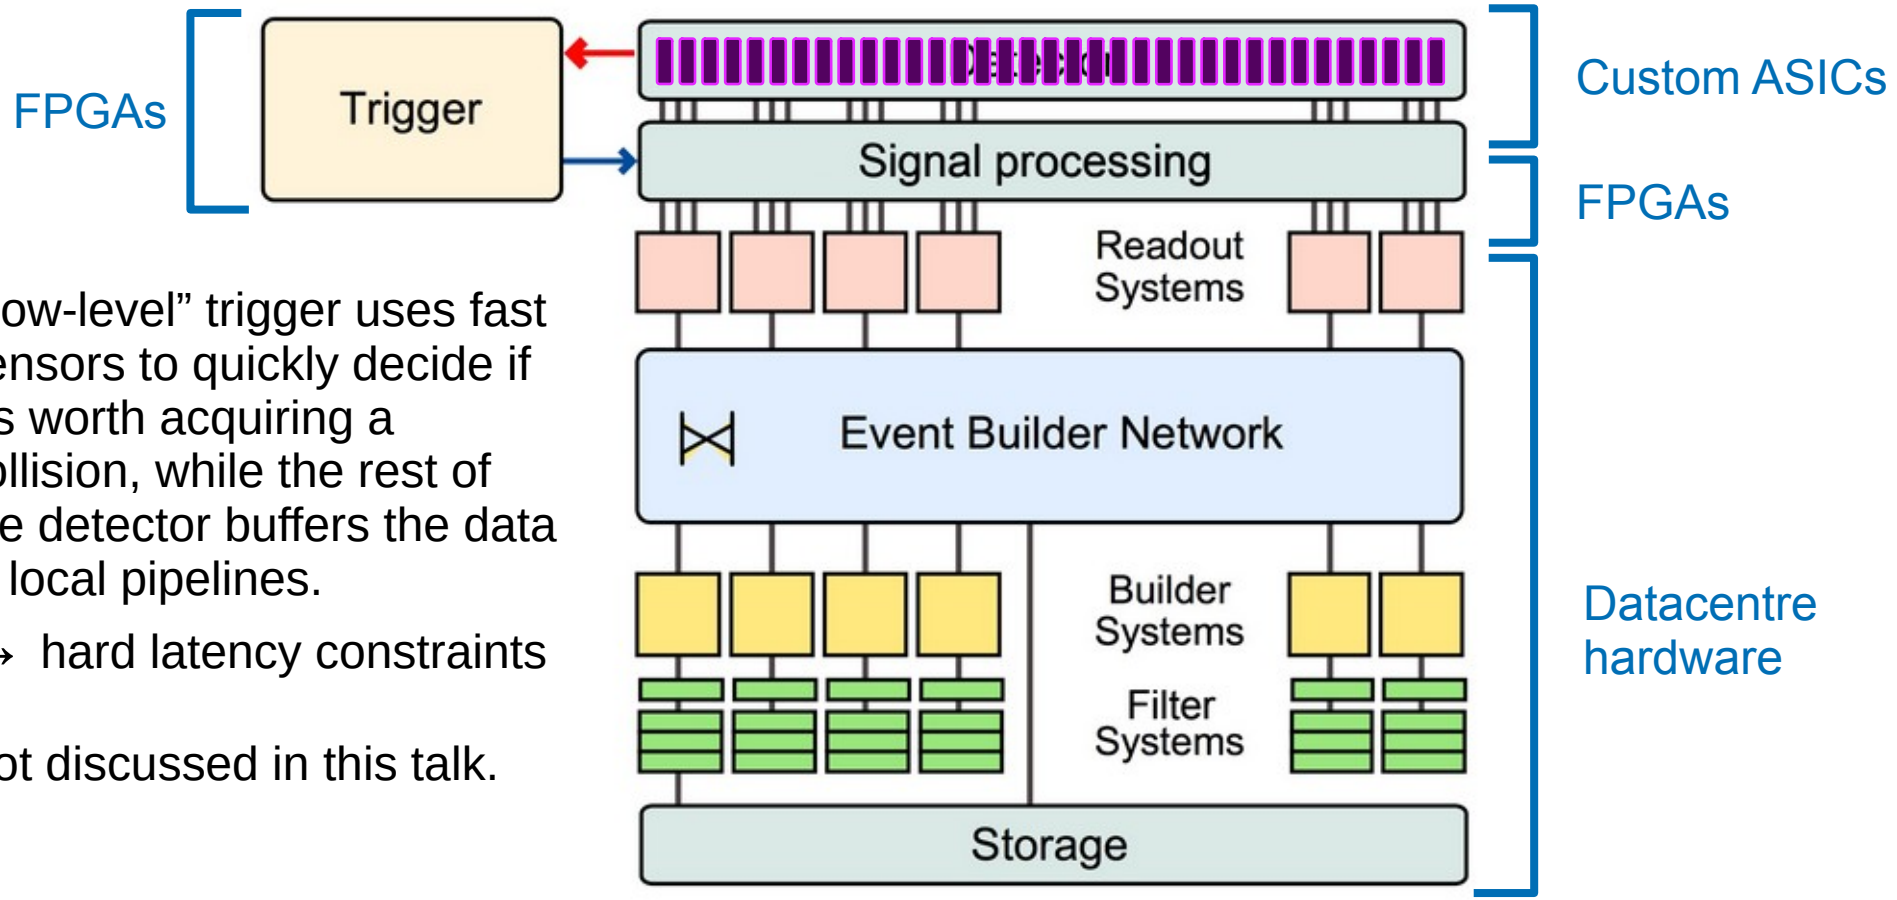

"Low-level" trigger uses fast sensors to quickly decide if it's worth acquiring a collision, while the rest of the detector buffers the data in local pipelines.

→ hard latency constraints

Not discussed in this talk.

DEALING WITH THE HAYSTACK

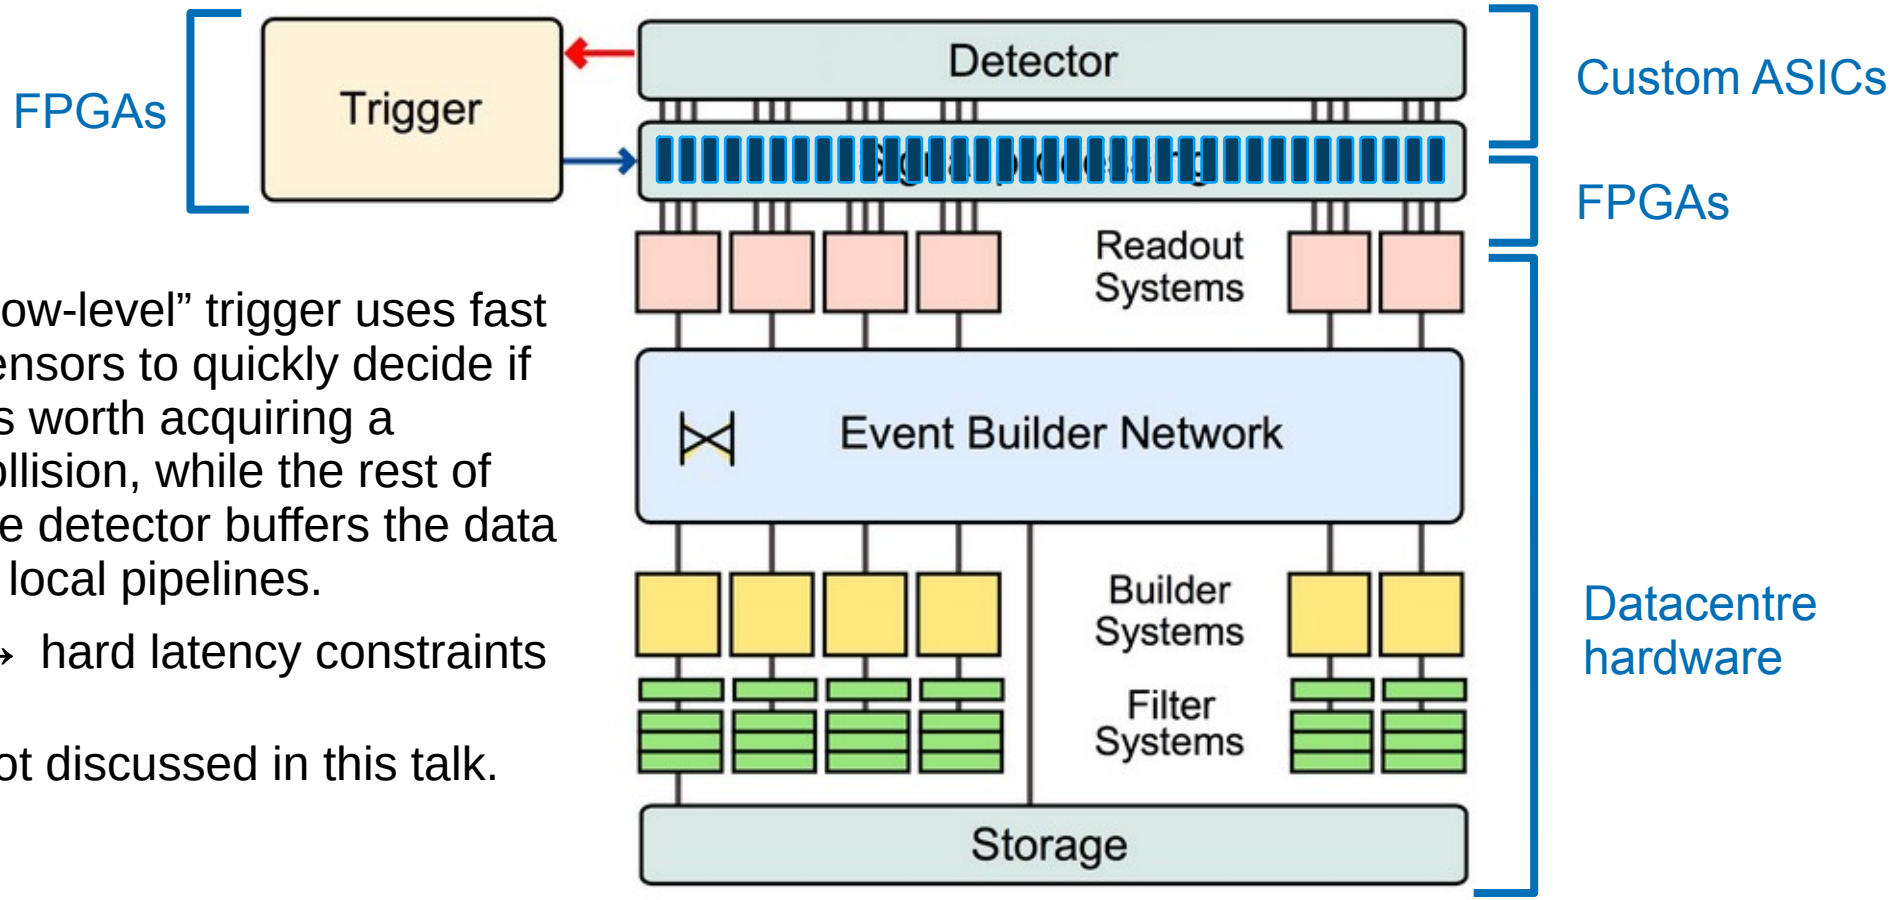

"Low-level" trigger uses fast sensors to quickly decide if it's worth acquiring a collision, while the rest of the detector buffers the data in local pipelines.

→ hard latency constraints

Not discussed in this talk.

DEALING WITH THE HAYSTACK

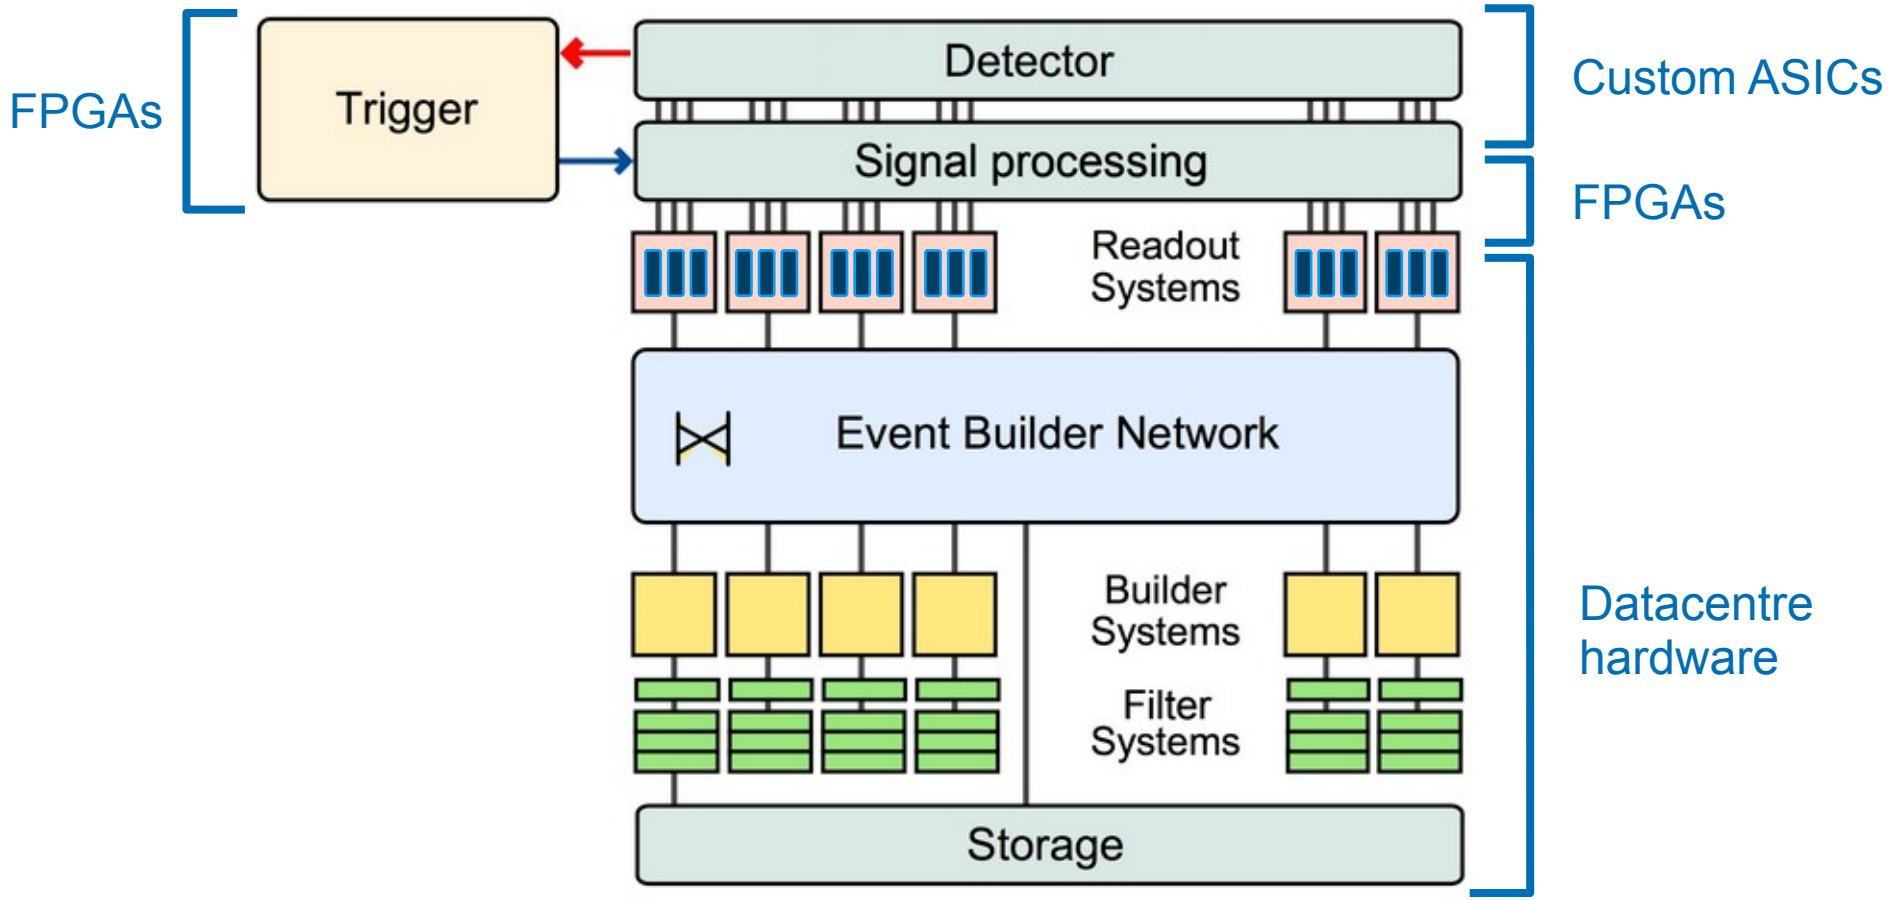

DEALING WITH THE HAYSTACK

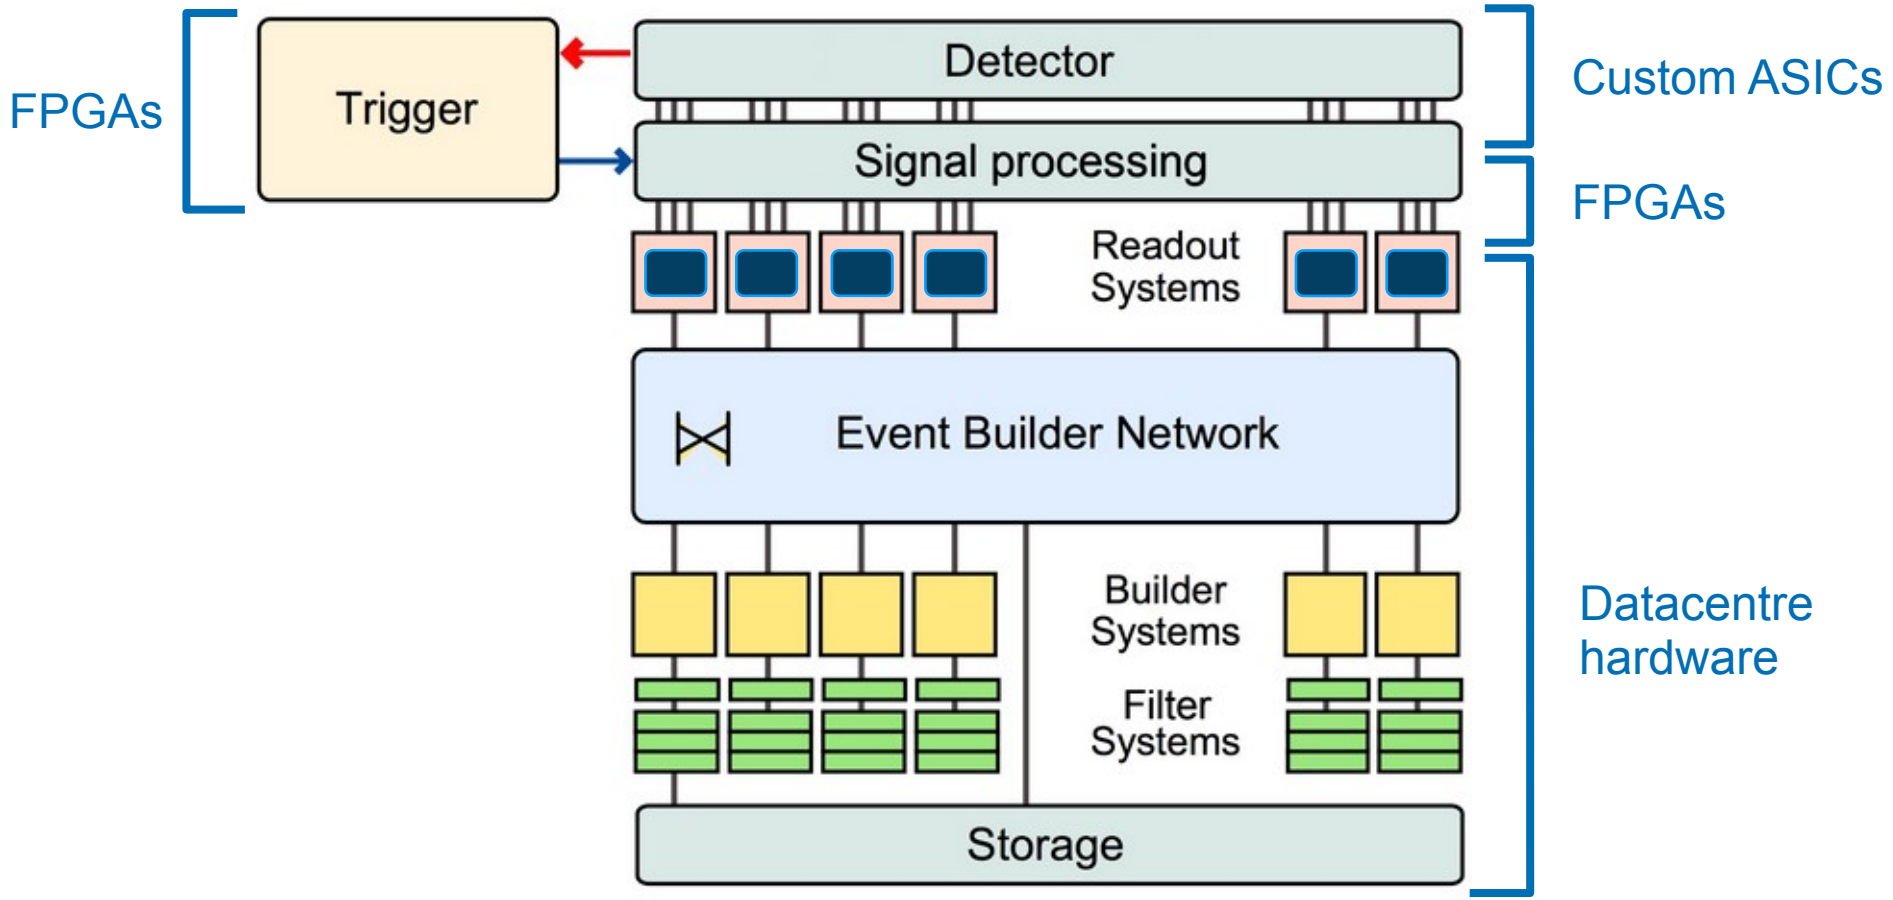

DEALING WITH THE HAYSTACK

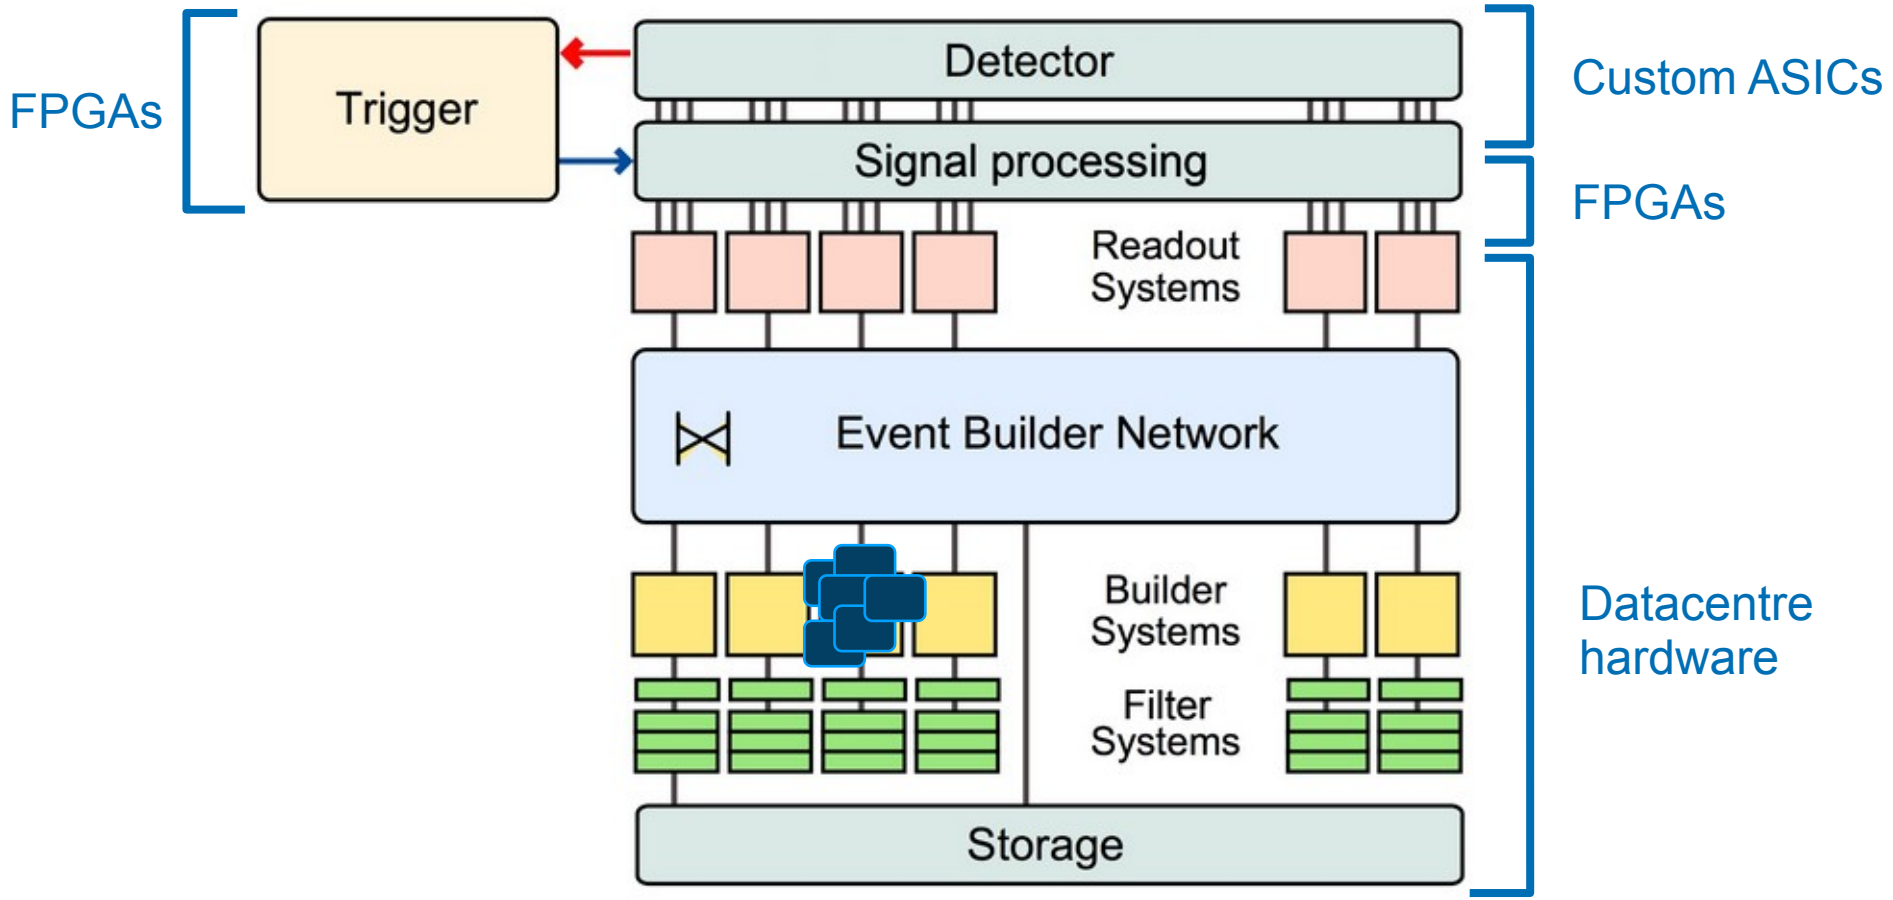

DEALING WITH THE HAYSTACK

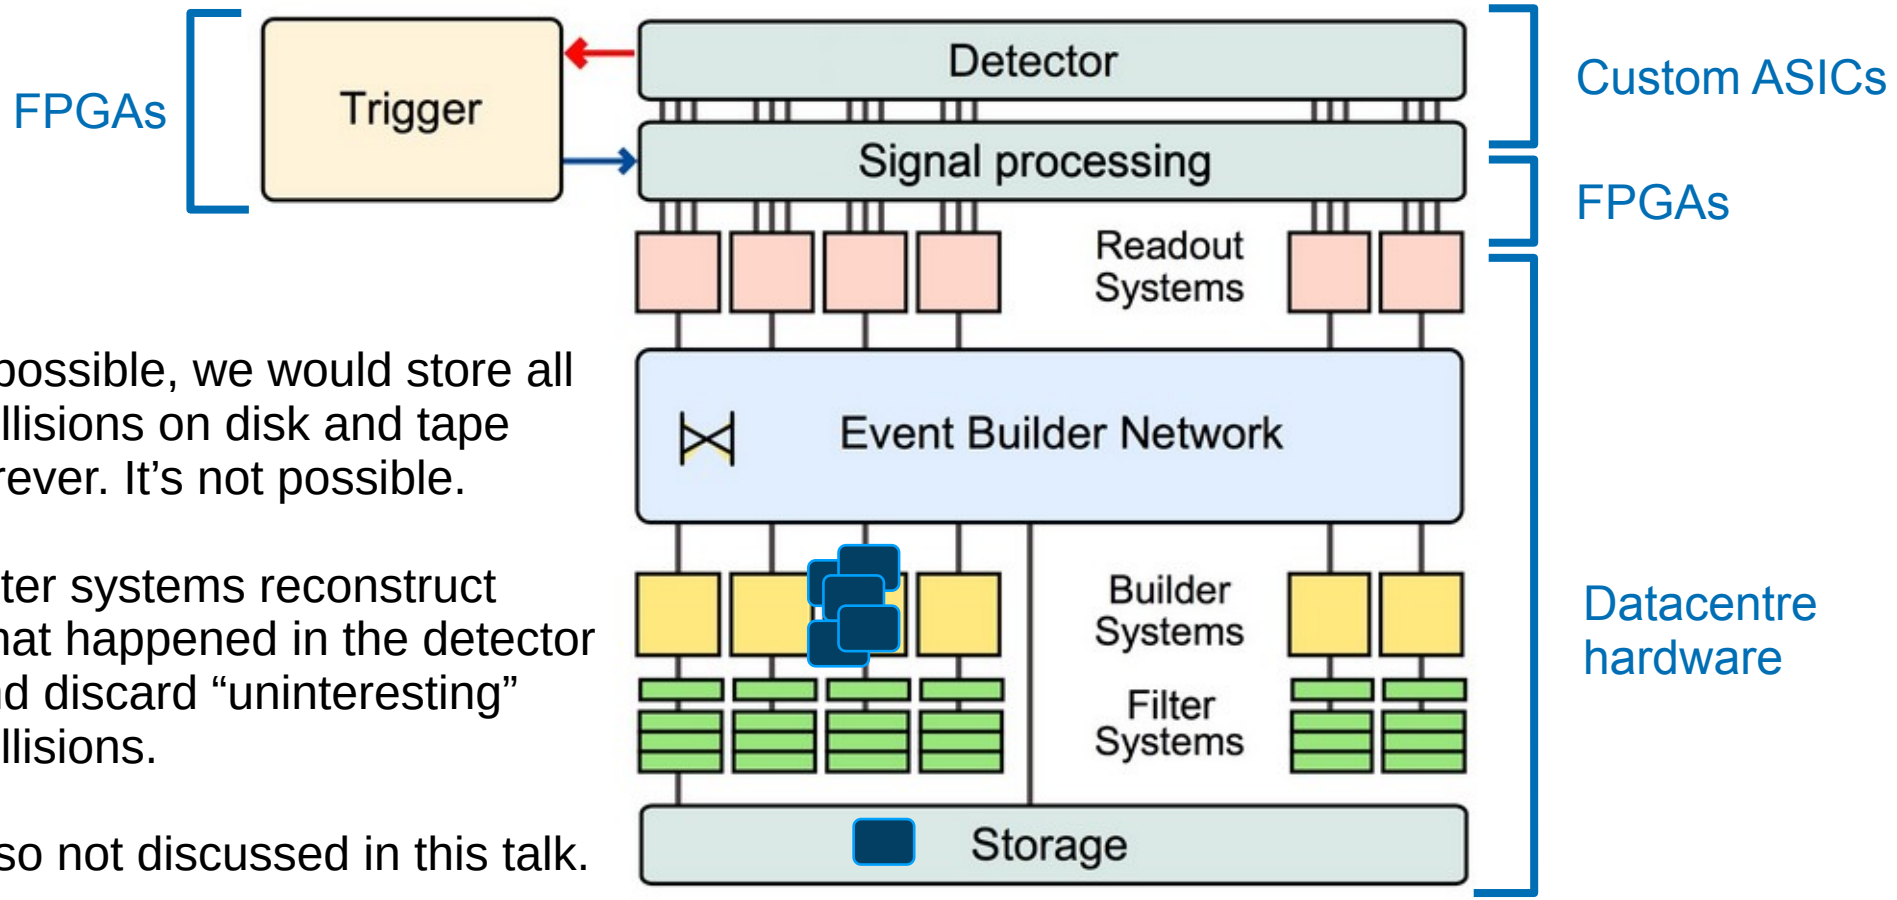

If possible, we would store all collisions on disk and tape forever. It's not possible.

Filter systems reconstruct what happened in the detector and discard "uninteresting" collisions.

Also not discussed in this talk.

IN NUMBERS

Experiment	Data size	Rate before trigger	Rate after trigger	DAQ throughput	
ALICE now	12000 kB	50 kHz	50 kHz	700 GB/s	
ATLAS now	1500 kB	27000 kHz	90 kHz	55 GB/s	90x
ATLAS 2030	5000 kB	27000 kHz	1 MHz	5000 GB/s	
CMS now	2000 kB	27000 kHz	100 kHz	200 GB/s	32x
CMS 2030	8400 kB	27000 kHz	750 kHz	6300 GB/s	
LHCb now	80 kB	24000 kHz	24000 kHz	2400 GB/s	5x
LHCb 2036	200 kB	24000 kHz	24000 kHz	12000 GB/s	

RADIATION

- Collisions → ionising radiation → bit flips, dead electronics
- Solution: radiation-hard custom ASICs inside the experiment
- Get data to easier datacentre environment over “slow” (10 Gb/s) optical links

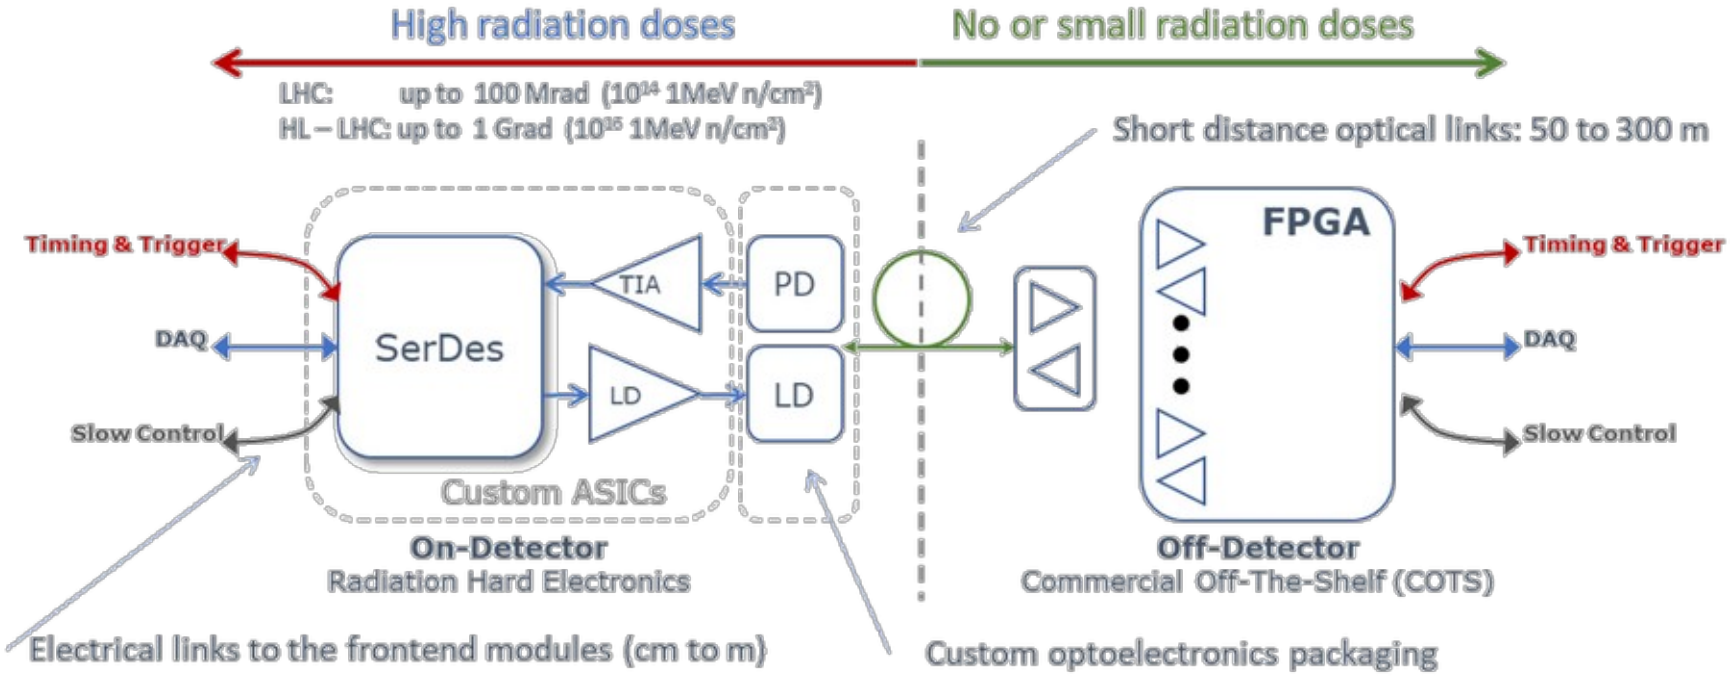

PRECISE TIMING

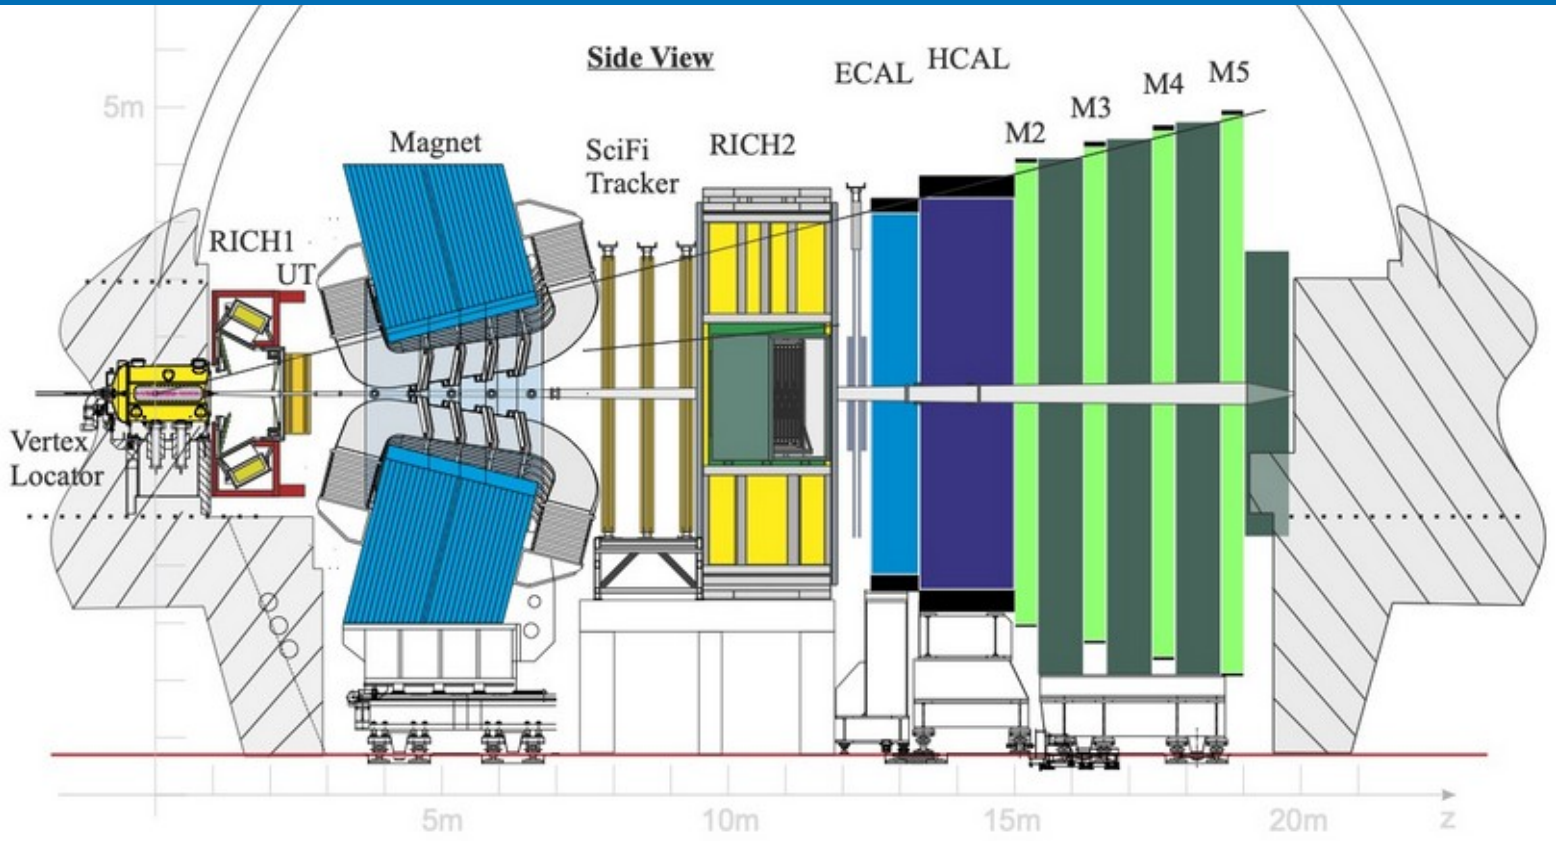

LHC collisions: Every 25 ns

If a local clock is off by 1 ns (4%), the particle you wanted to measure is already gone!

EVEN MORE PRECISE TIMING

- Even though hundreds of collisions are delivered by HL-LHC every 25 ns the particles coming out of them won't produce signals in detectors at precisely at the same time
- With a sufficiently precise clock (10 ps), time could be used as a "fourth dimension" to separate these superimposed signals
- All LHC experiments are working on this, in anticipation of HL-LHC, tuning circuit design and FPGAs

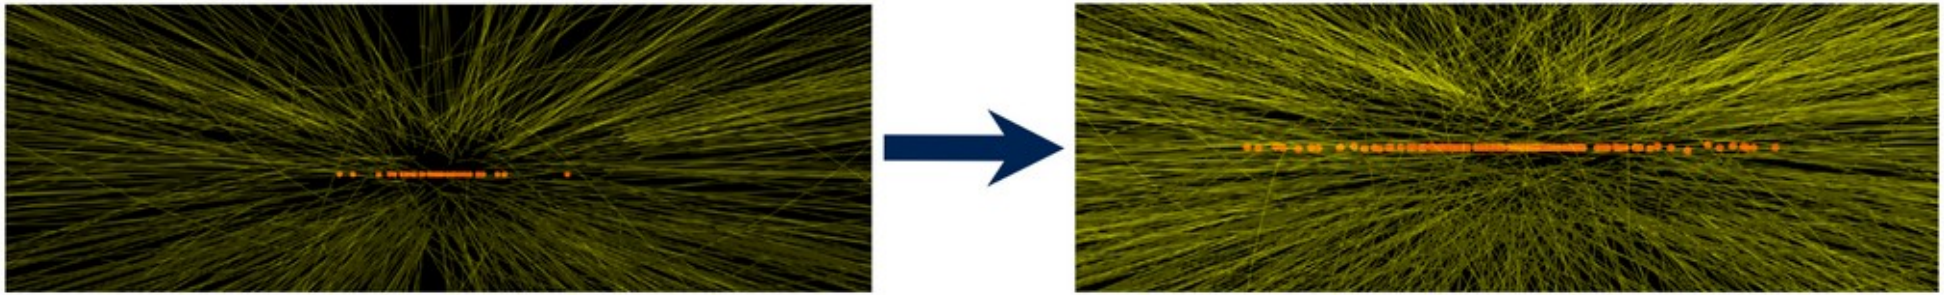

DATA COLLECTION IN THEORY

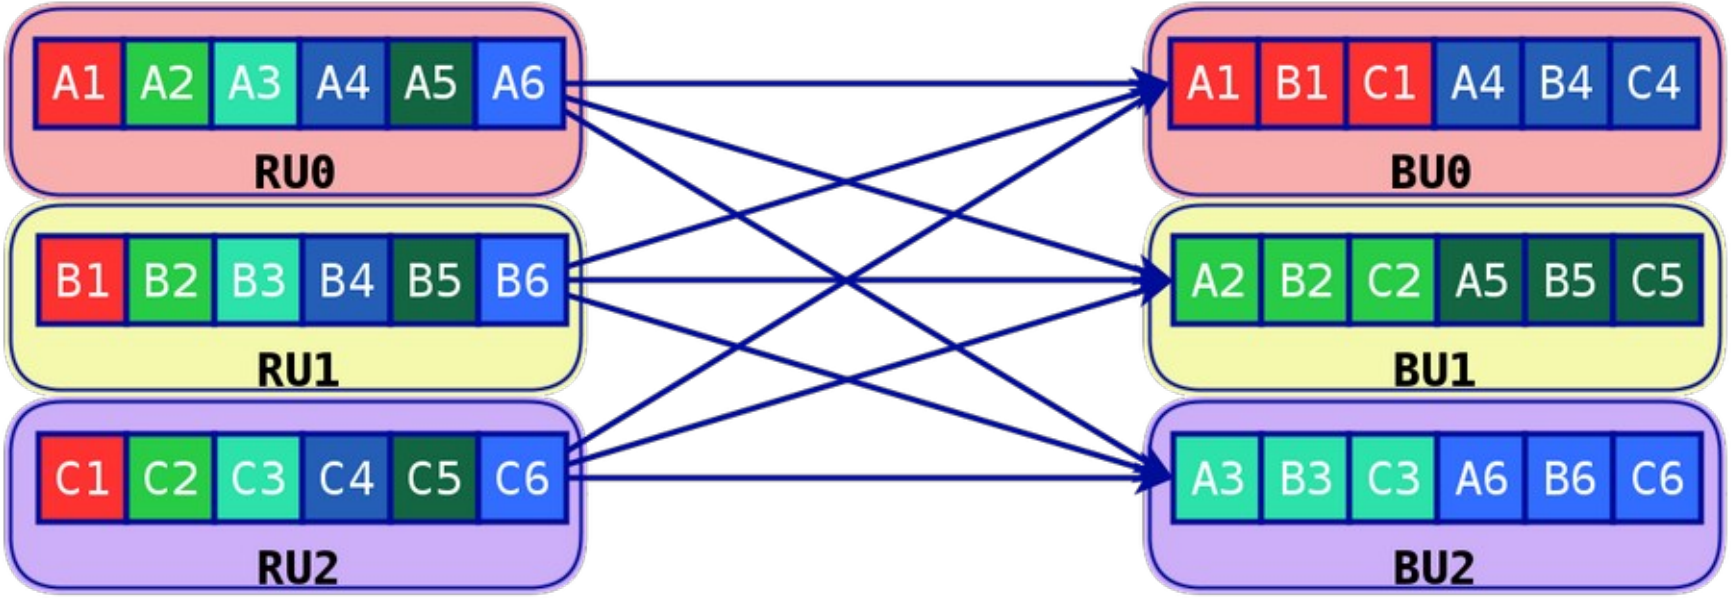

- Readout unit (RU): receives processed signals from some sensors
- Builder unit (BU): assembles all signals corresponding to the same observed phenomenon

DATA COLLECTION NETWORKS

- The BUs collect data from different RUs
  → Many-to-one communication
- Data transfers are driven by the availability of the data from the detector
  → Synchronous, bursty traffic
- When many sources send synchronous microbursts of data to a destination
  → Congestion
  → The network buffers are overflown
- Must be kept under control, otherwise: “Catastrophic throughput collapse”

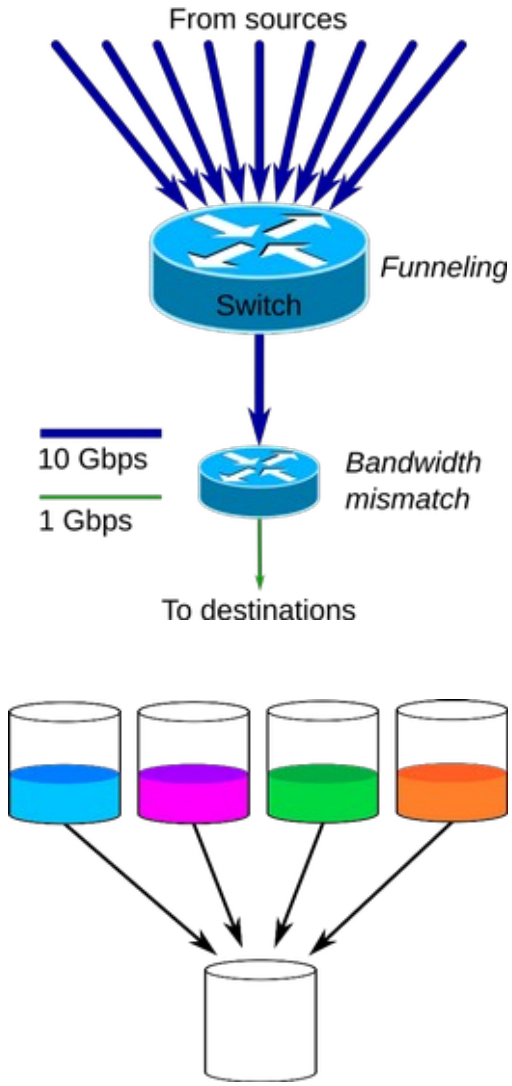

DATA COLLECTION NETWORKS

- The BUs collect data from different RUs
  → Many-to-one communication
- Data transfers are driven by the availability of the data from the detector
  → Synchronous, bursty traffic
- When many sources send synchronous microbursts of data to a destination
  → Congestion
  → The network buffers are overflown
- Must be kept under control, otherwise: “Catastrophic throughput collapse”

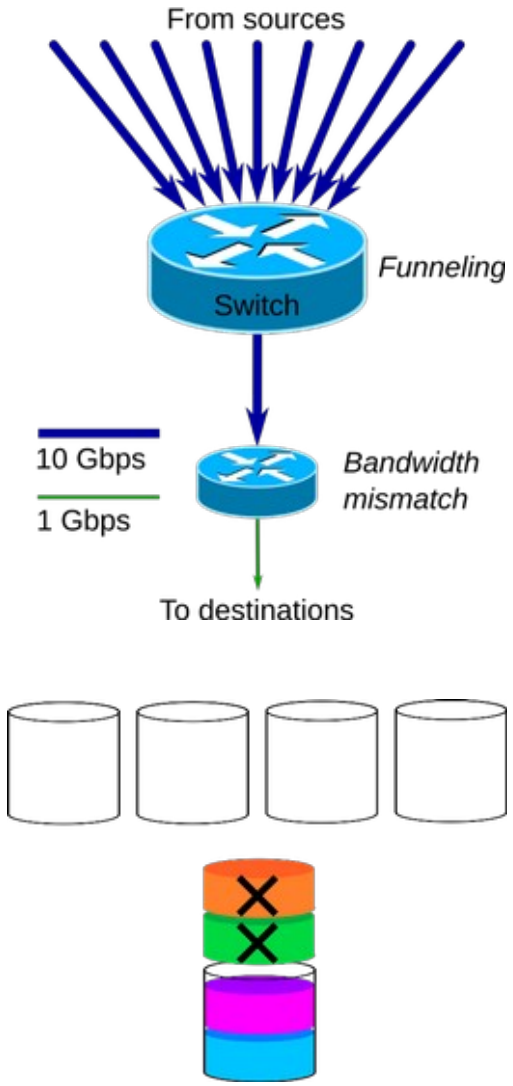

DATA COLLECTION SOLUTIONS

- Move the problem
  from shallow-buffered
  network devices
  to deep-buffered DAQ servers
  - “Lossless” networks
    (RoCE, InfiniBand, UltraEthernet)
  - Custom traffic shaping on servers
- Not a new problem, but no easy fix:
  painstaking testing
  and tuning needed

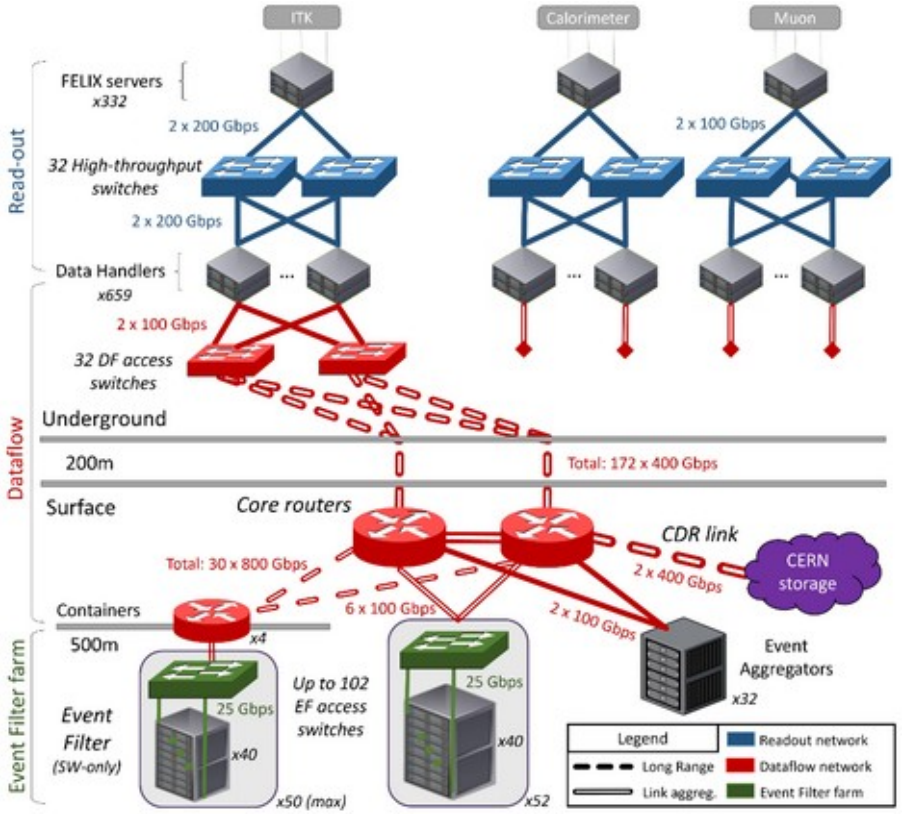

AND MORE!

Not discussed, but also constantly evolving:

- Initial data processing
  on detector electronics
- Data filtering systems
  on CPU/GPU/FPGA
  and their ever-growing datacentre needs
- Control and monitoring systems
- Long-term safe data storage