ATLAS Data Acquisition Phase 2 Outlook

Matías Bonaventura, CERN
On behalf of ATLAS TDAQ Collaboration

CEPC International Workshop, September 2022

LHC baseline plan for the next decade

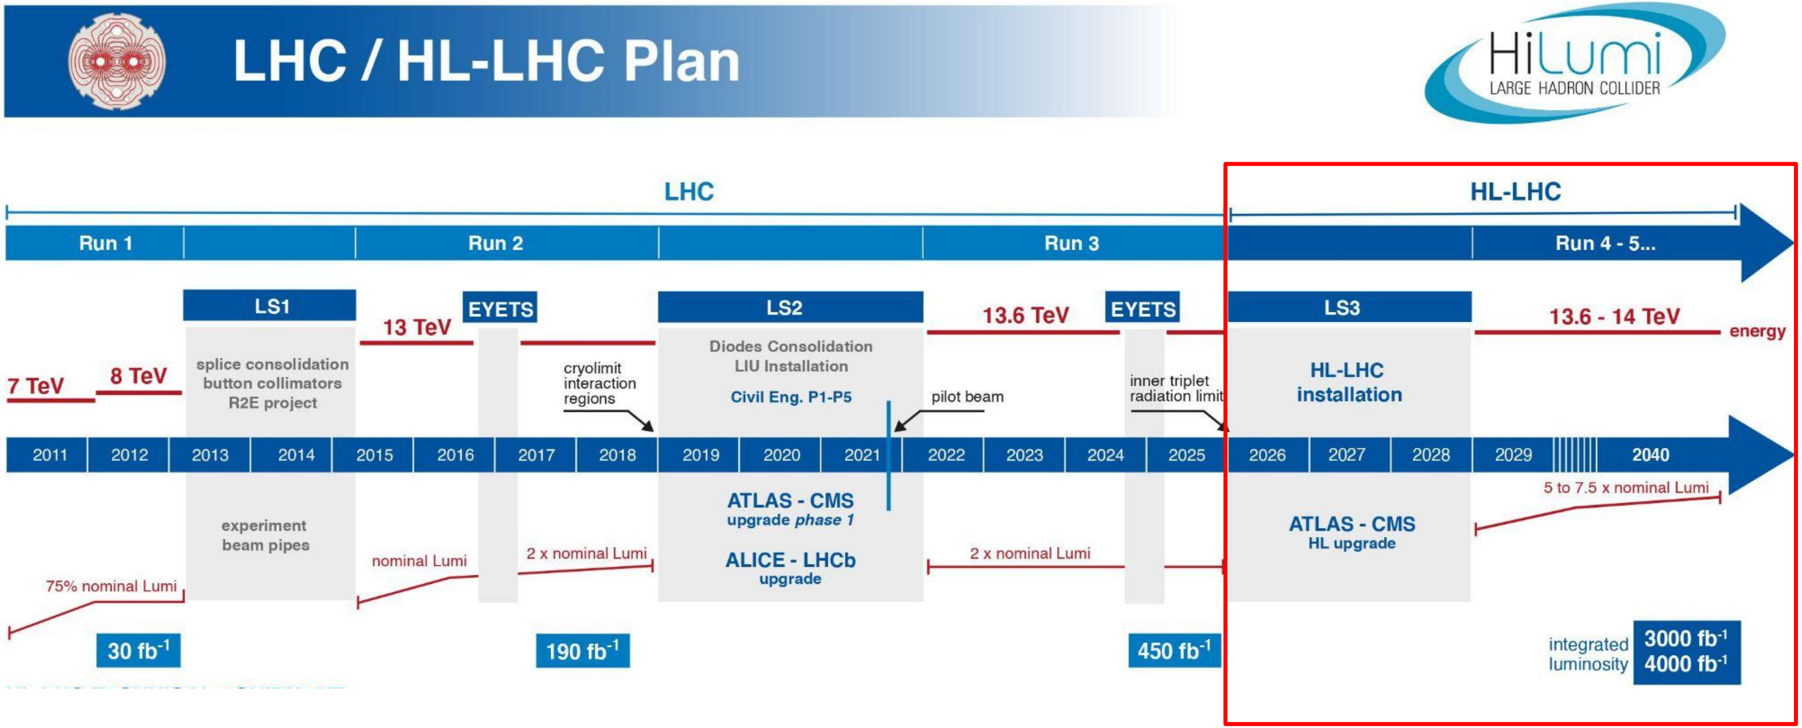

Data Acquisition Challenges for the HL-LHC

Read-out at **1 MHz**
10 x more than today's

Throughput: **4.6 TB/s**
~20 x more than today's

Selected events: **10kHz, 10 PB**
10 x more than today's

Major upgrade of all TDAQ sub-systems

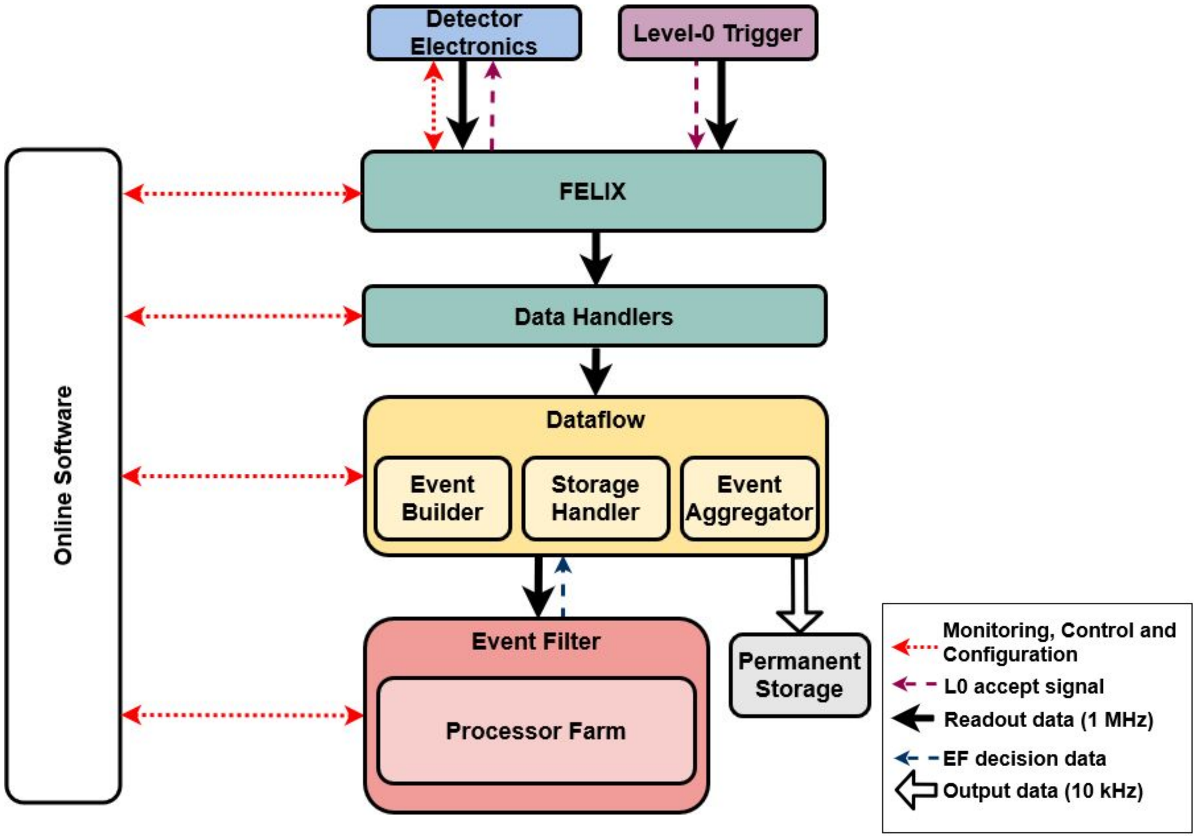

Outline

Dataflow

Network

Online software

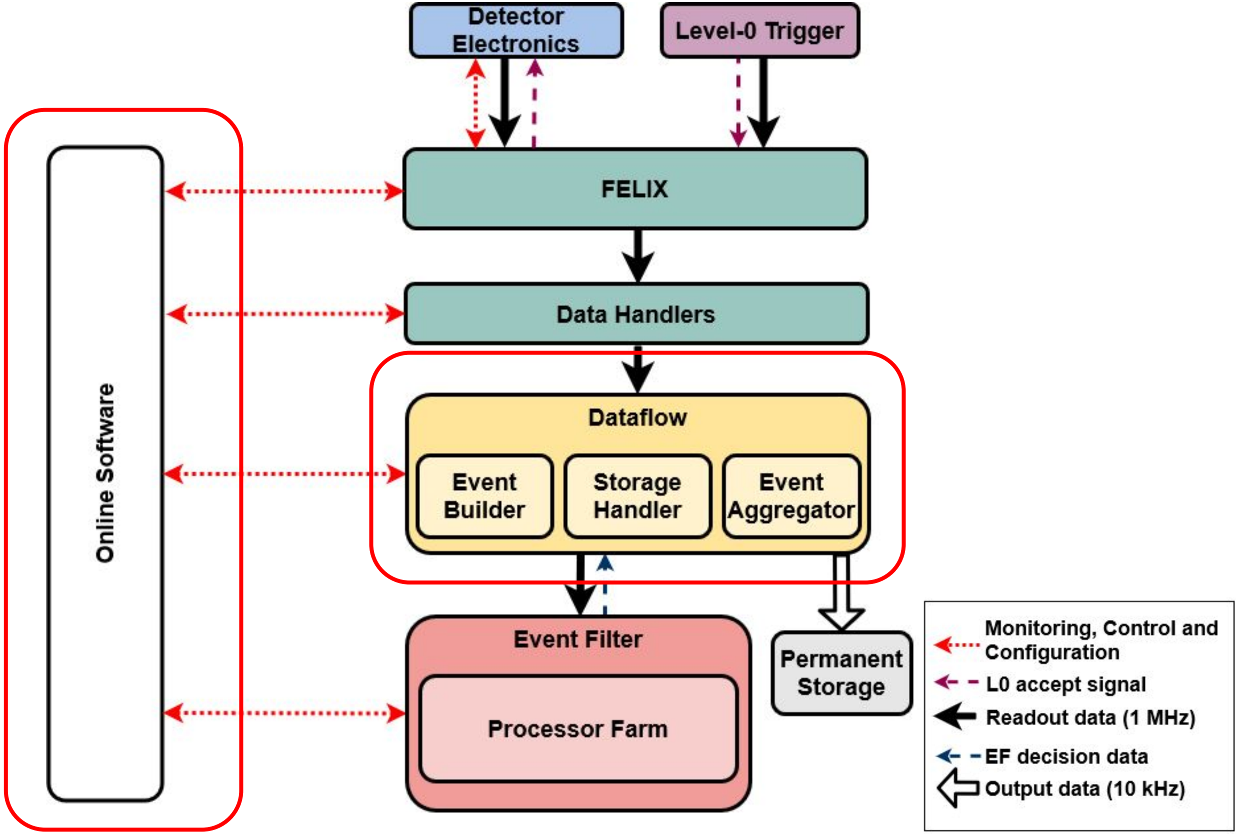

Dataflow

Phase 2 Dataflow

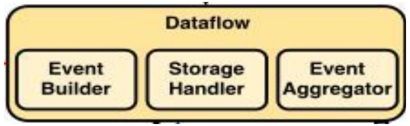

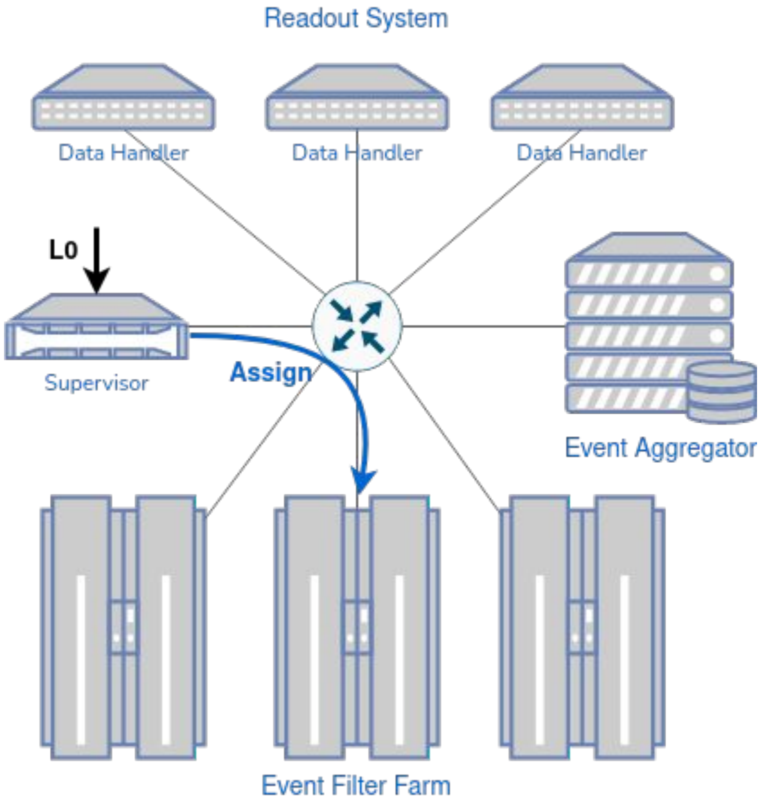

Phase 2 Dataflow

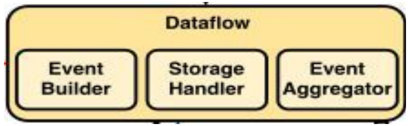

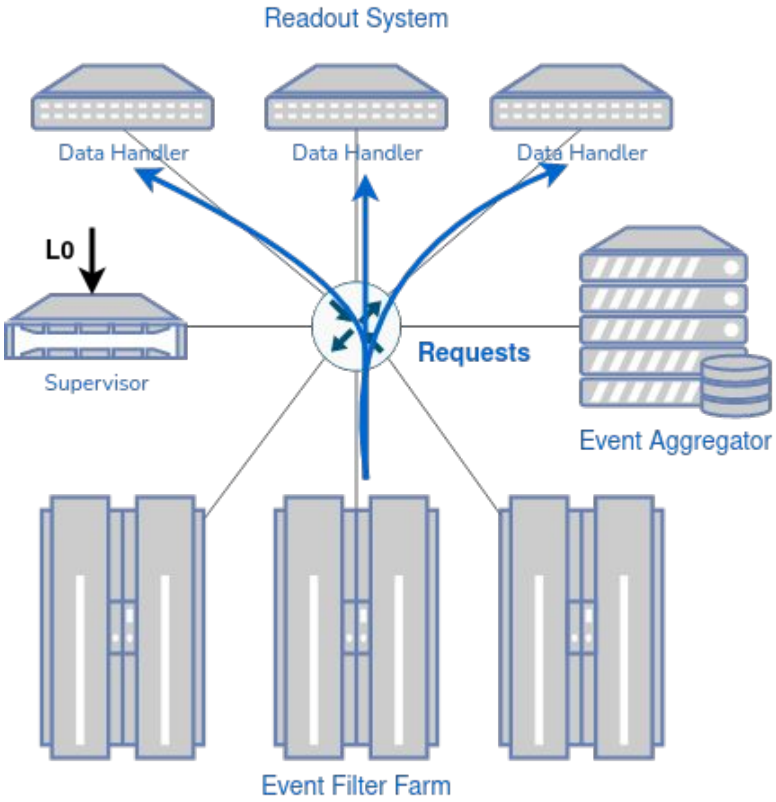

Phase 2 Dataflow

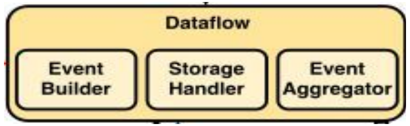

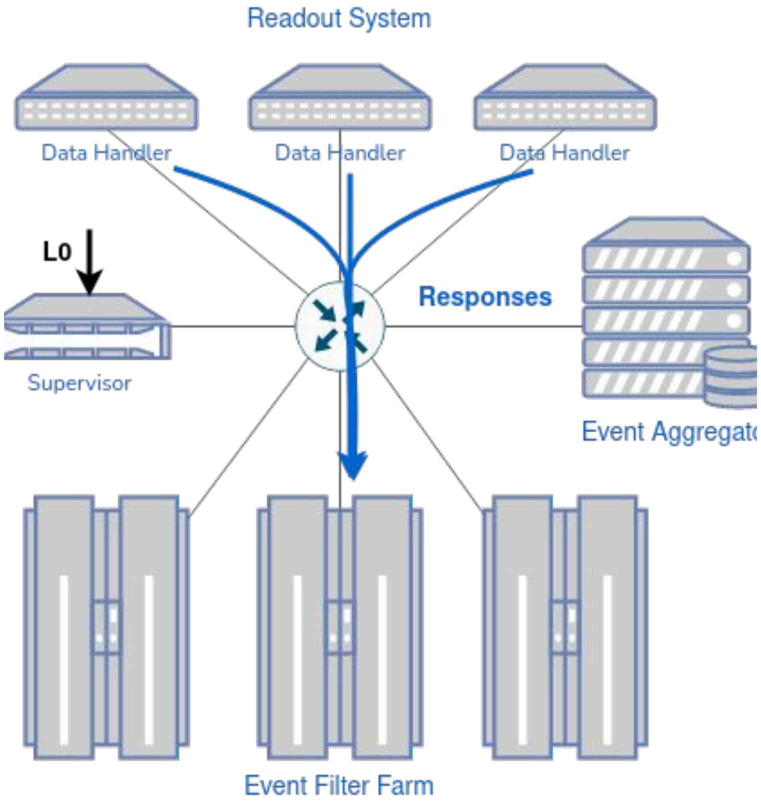

Phase 2 Dataflow

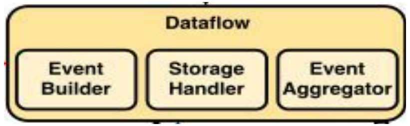

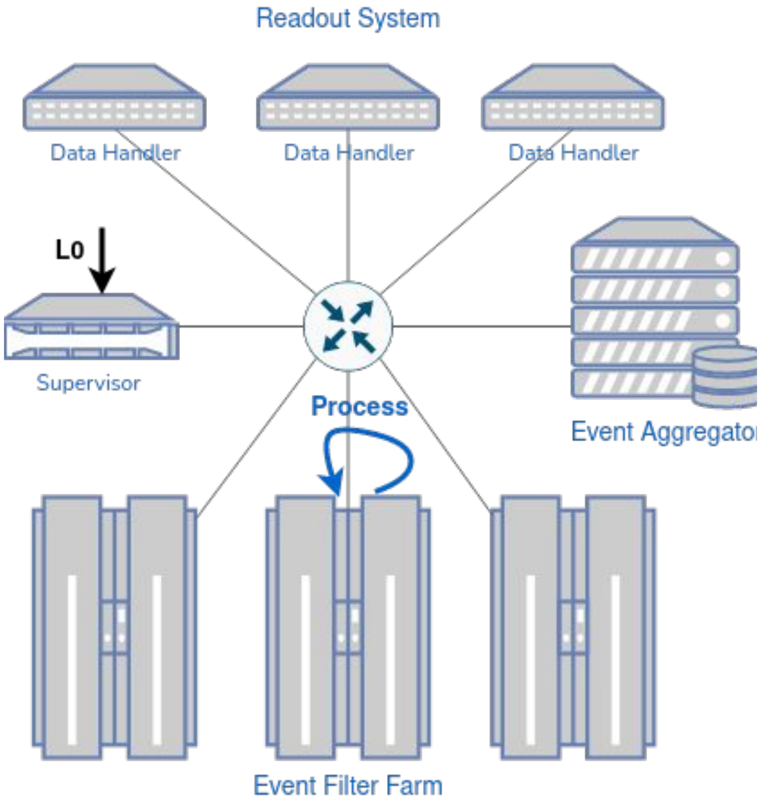

Phase 2 Dataflow

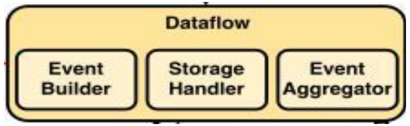

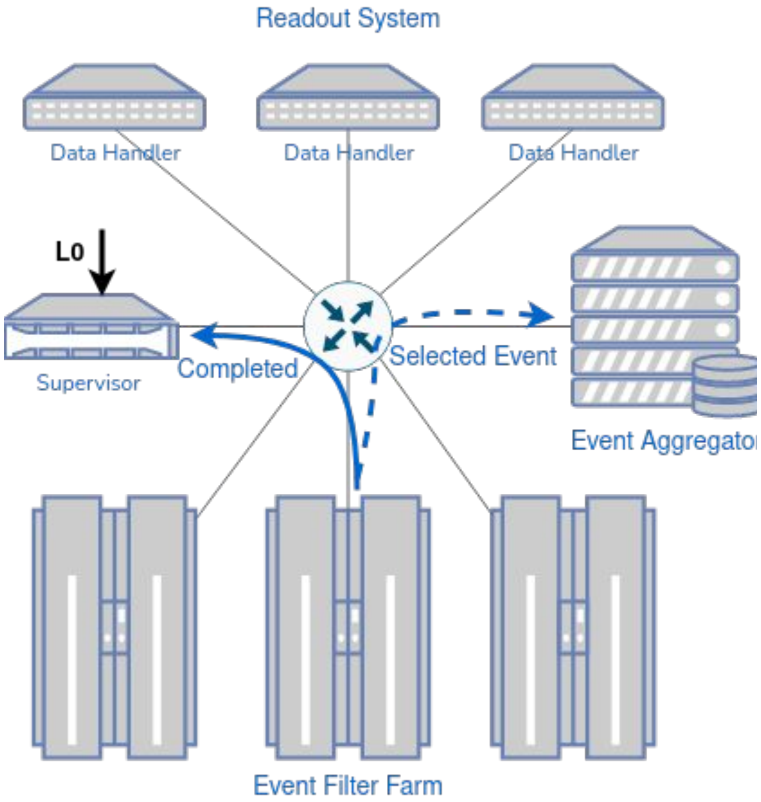

Dataflow design evolution

Full event building at 1 MHz
w.r.t Incremental partial build in previous runs
Events are fully built in RAM in the processing nodes

Decouple data collection from data analysis
Readout transfer all data to the processing farm (contrary to previous runs)

Dedicated system for selected data
Based on HDDs (~10PB)

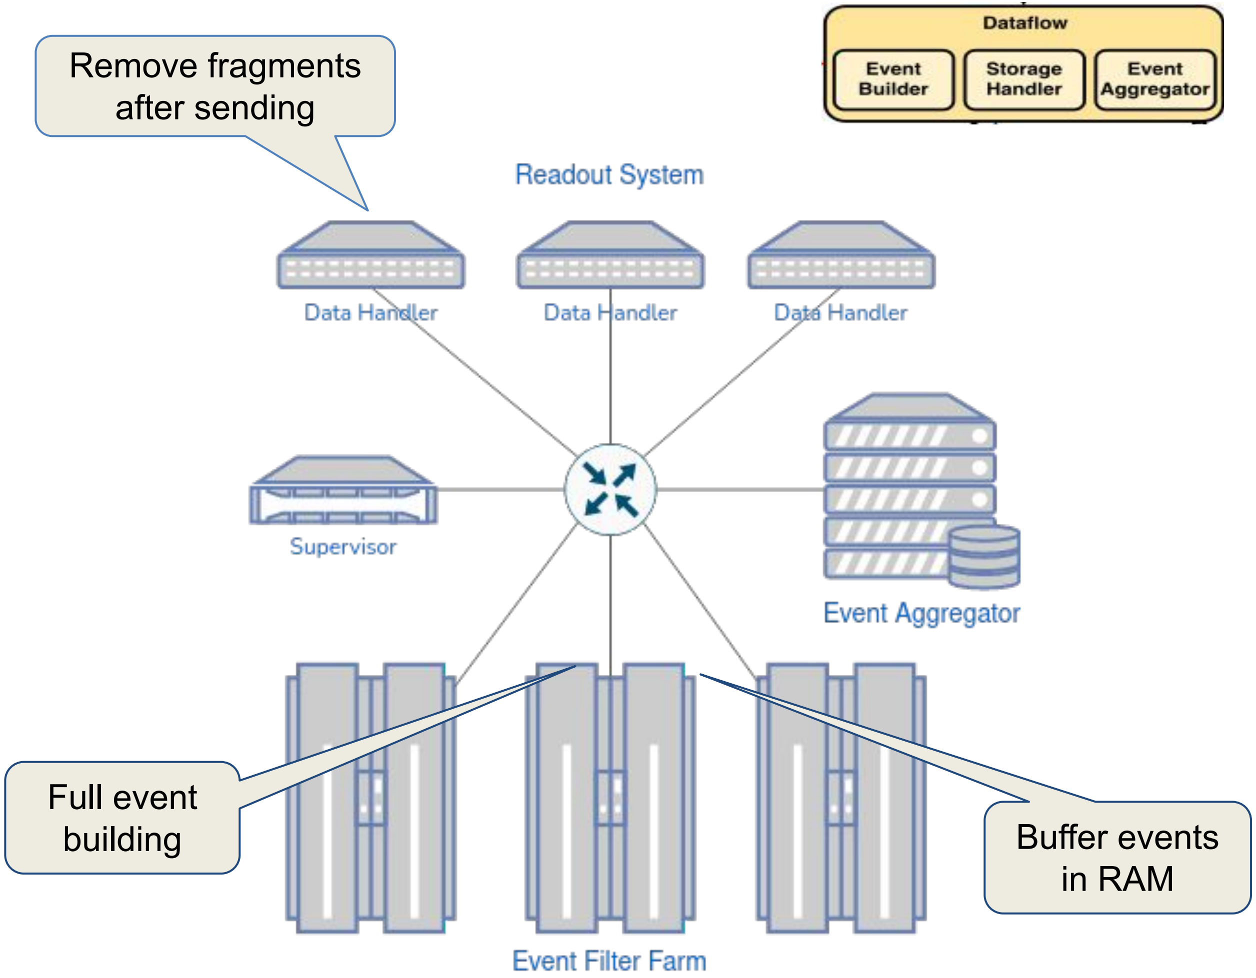

Ongoing R&D - *two complementary strategies*

Main concerns and focus of studies
Increased data rate at 1 MHz - x10 with respect to. run 3
Efficiently use network resources in a in 100 GbE environment

Design prototypes

Goal: benchmark applications in realistic environment

Based on run 2 applications
Strategy: maintain message exchange rates under control

Test lab - O(100) servers
Acquired equipment that resembles Phase 2 hardware (including switches with 400 Gbps uplinks)
Network efficiency studies

Dataflow simulations

Goal: evaluate different designs and scale up system

PowerDEVS, open-source general purpose discrete event simulation used for run 2

Simulation scenarios for Phase 2
Readout buffer evaluation
Network impact of TCP incast
Supervisor load balancing policies

Dataflow Simulations

Allows **evaluating different designs** simulating unavailable hardware (e.g. **system scale up**)

Model includes
Network devices and protocols
Dataflow DAQ applications
Adaptive network topology

Model does not include
OSI layers 3 and below
Trigger algorithms
Hardware usage (e.g. CPU usage)

Dataflow Phase 2 simulation model

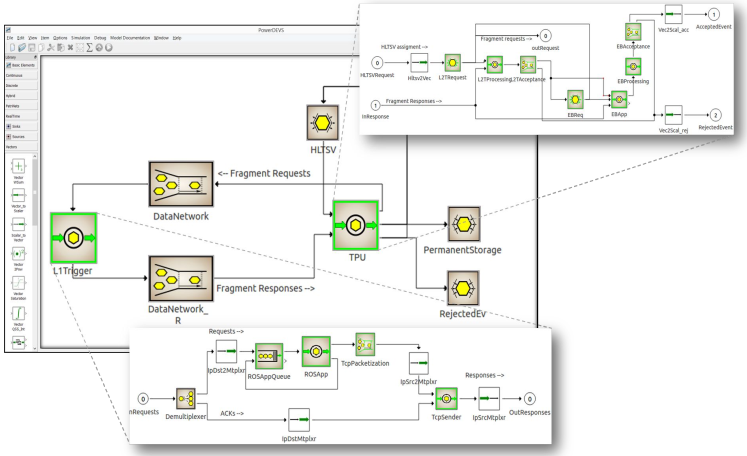

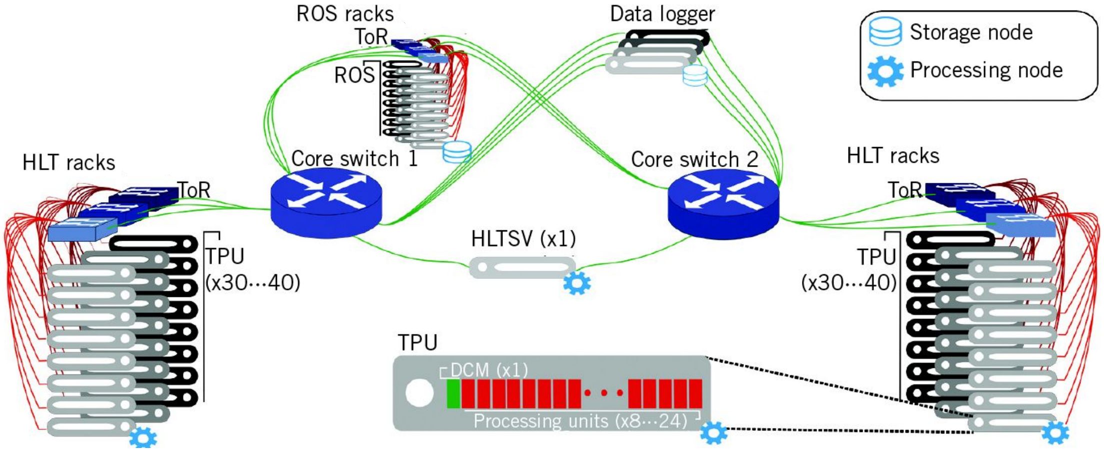

Dataflow Simulations

Allows **evaluating different designs** and simulating unavailable hardware (e.g. **system scale up**)

Model includes
Network devices and protocols
Dataflow DAQ applications
Adaptive network topology

Model does not include
OSI layers 3 and below
Trigger algorithms
Hardware capacity

Comparison of DAQ architectures by simulation

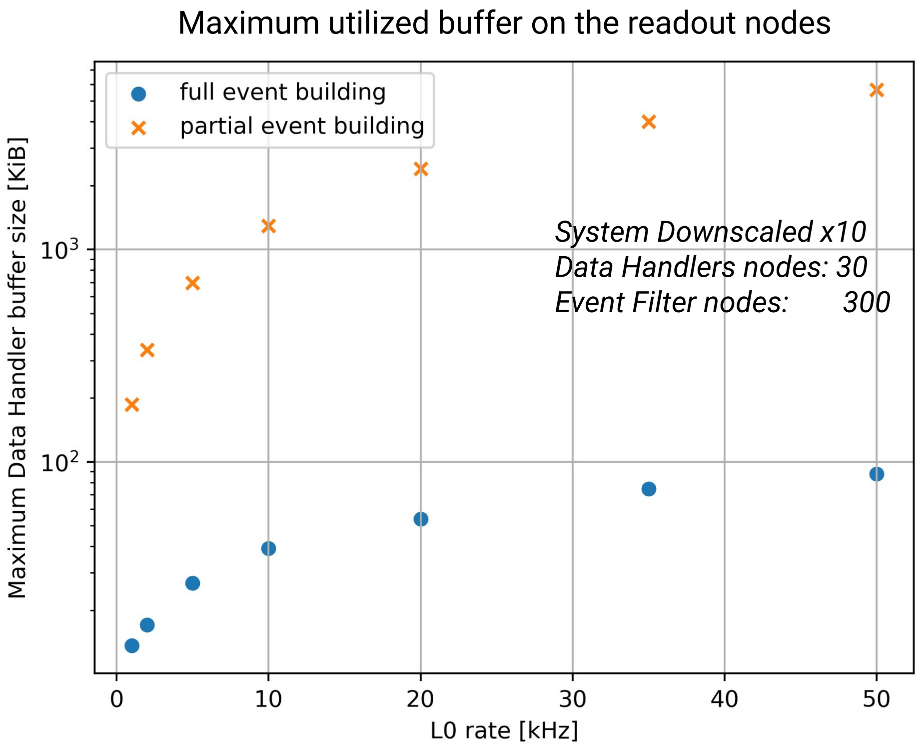

Adam Abed Abud, PhD Thesis 2022

Persistent Storage Design

Original Storage Handler
Persistent cache decoupling the detector readout from the event filter

Based on Solid State Drives (SSD)
Total throughput of ~7.8 TB/s
~1800 SSDs provide ~36 PB of storage

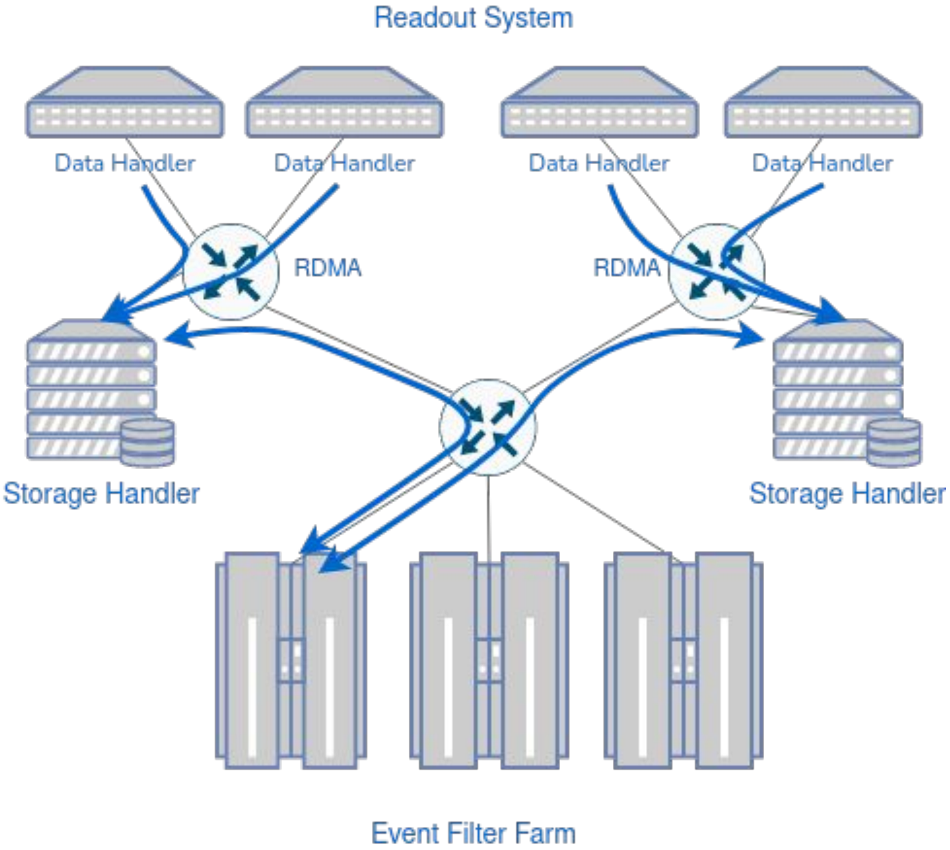

Persistent Storage Design

Market trends updated since project original design in 2017
Cost and endurance

Risk assessed too high
Baseline design without persistent storage
Significant architecture change

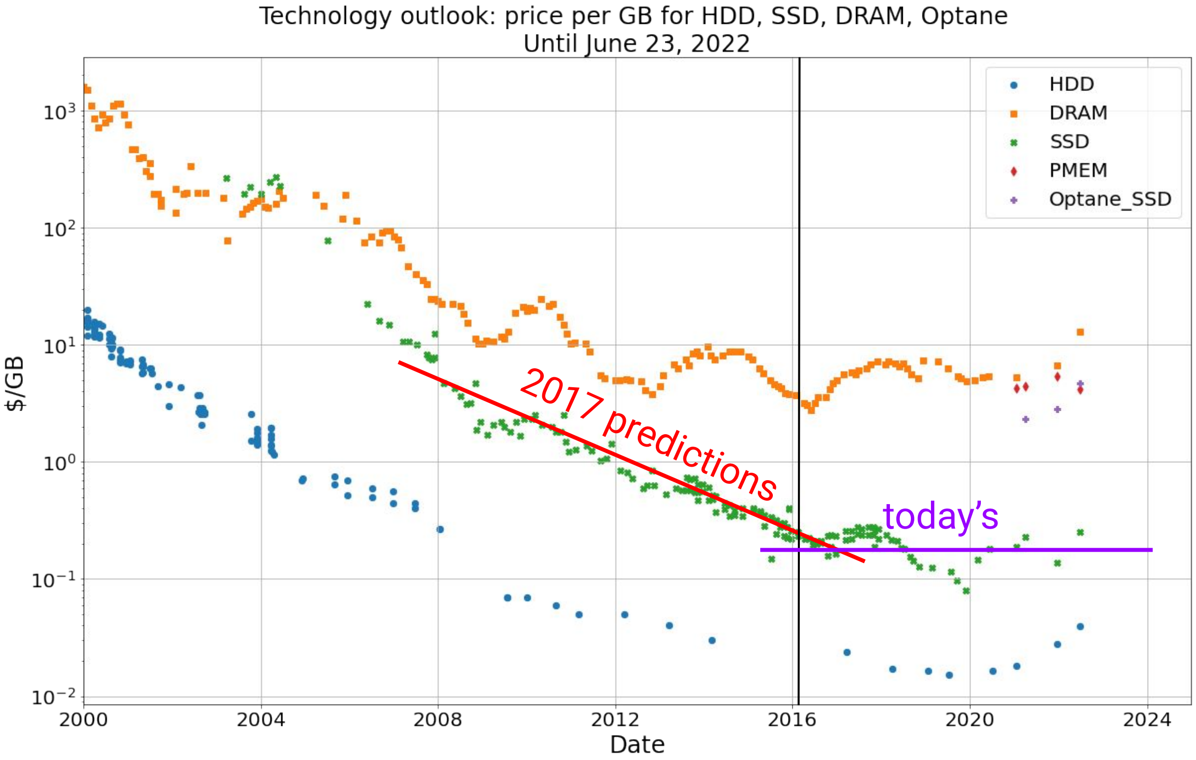

Data collected by John C. McCallum.
Data collected by Adam Abed Abud since 2018

Persistent Storage Design

As alternative secondary design

Viability depends on technology improvement and/or cost reduction

Closely interacting with industry
Current Openlab collaboration with Intel

DAOS: A high-performance distributed key-value store, open-source and developed by Intel, taking advantage of persistent memory

DAOS as a drop-in replacement for a persistent Storage Handler

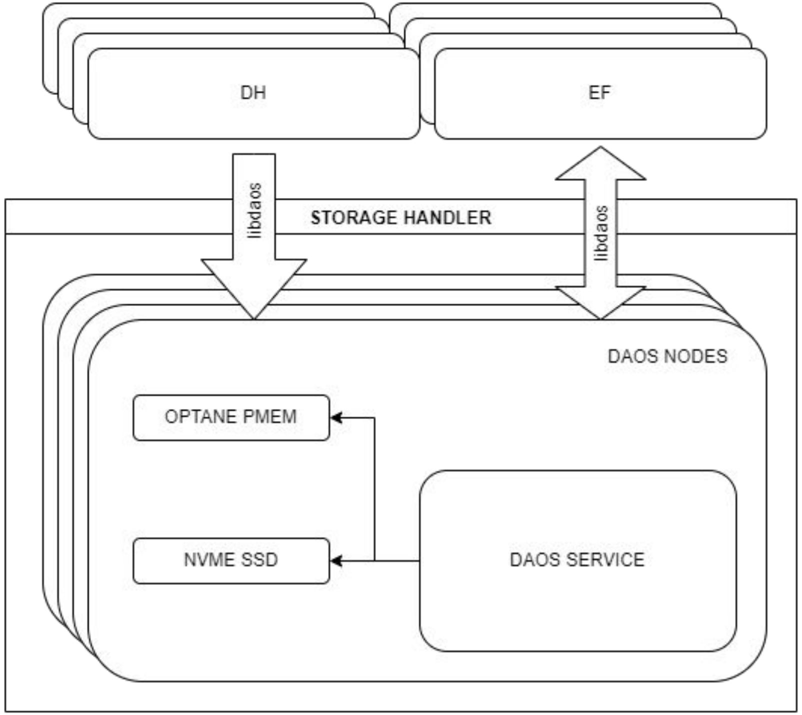

Network

RDMA evaluation and 100 GbE environment

Phase 2 Data Acquisition Network

2 independent network domains

Readout network
High-throughput (5 TB/s)
Sliced
Redundant
RDMA (RoCEv2, Infiniband, iWARP)

Dataflow network
High-throughput (5 TB/s)
All-to-all communication pattern
Redundant core layer
TCP/IP

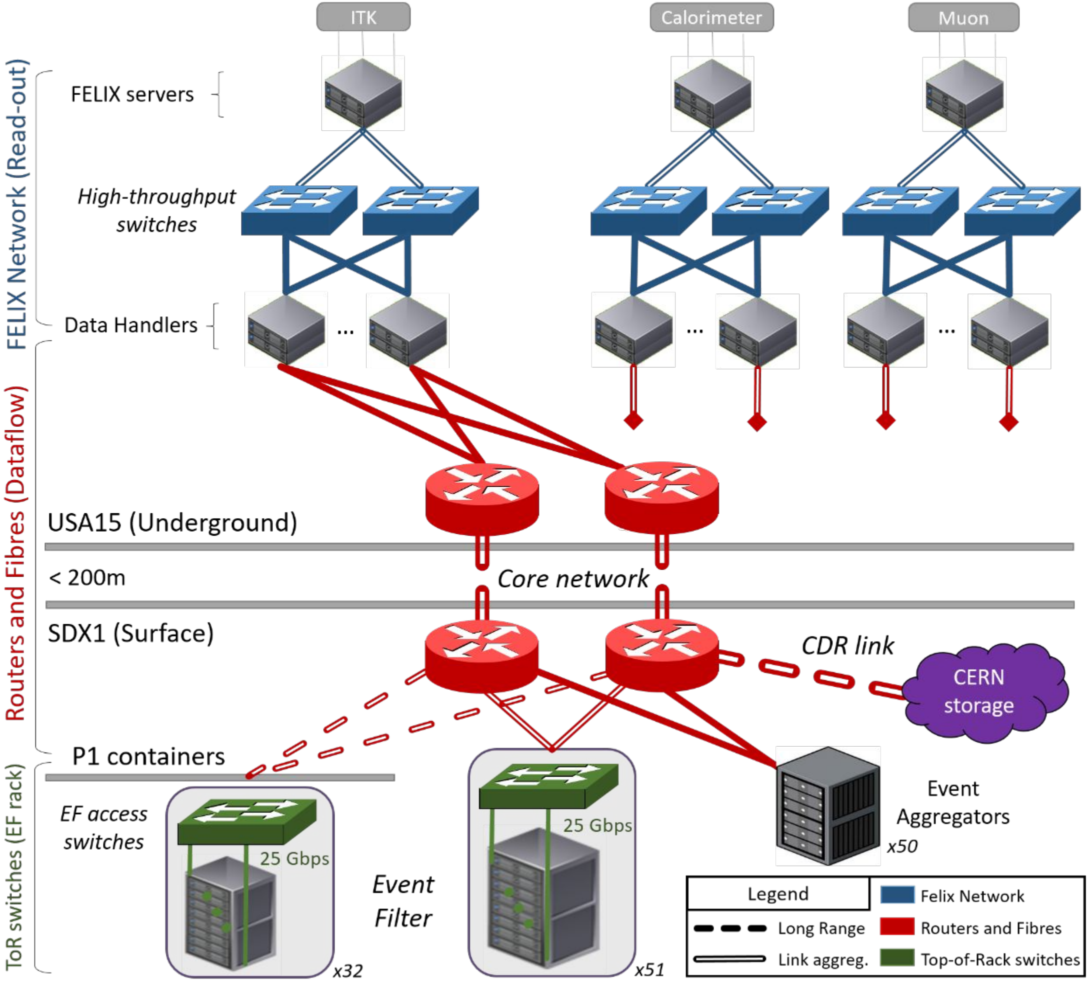

Readout network

Remote Direct Memory Access (RDMA) protocol evaluation

InfiniBand (IB)
High-throughput and low-latency protocol stack used in HPC
Requires dedicated network hardware

RDMA over Converged Ethernet (RoCEv2)
RDMA encapsulating IB packet over non-reliable UDP
Requires support/configuration from the underlying network

iWARP
RDMA over reliable TCP
Less requirements on the underlying Ethernet network

Simple sender-receiver scenario over 100 Gbps link

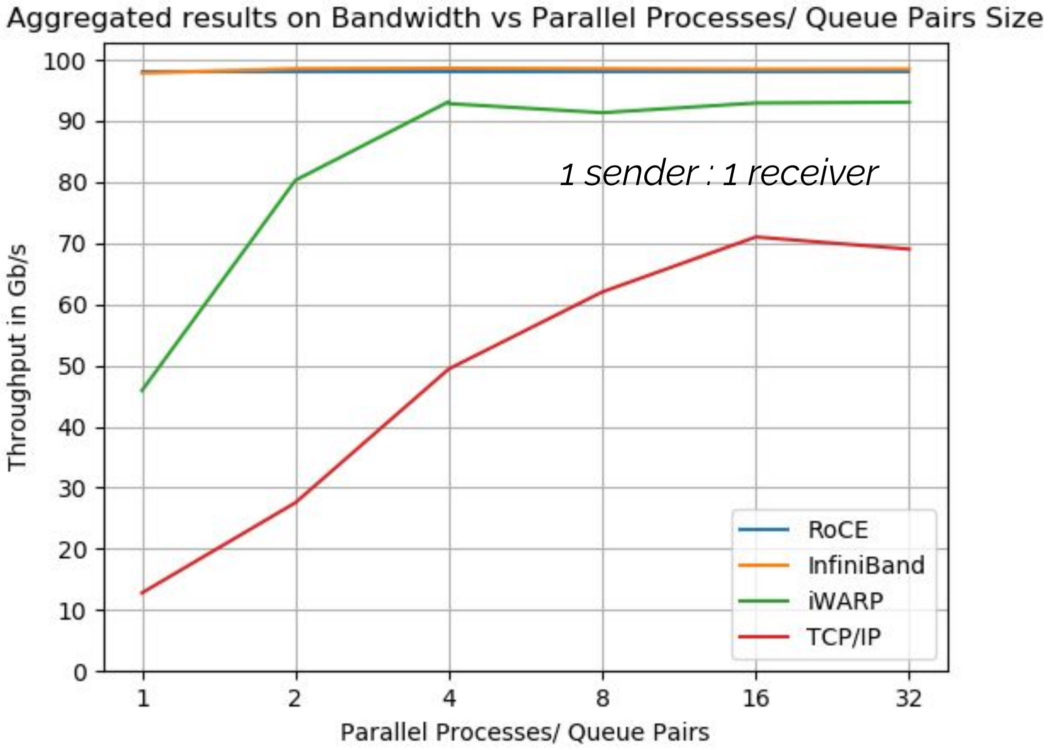

In a many to 1 scenario, all RDMA implementations fill up the link

Dataflow network

RDMA does not provide clear advantages
Network complexity makes iWARP the only option
Hardware increased cost

TCP connections over 100GbE
Optimized configuration improves performance: buffers, MSS, coalescing,
At least 3 TCP connections to reach a 90+ Gb/s total throughput

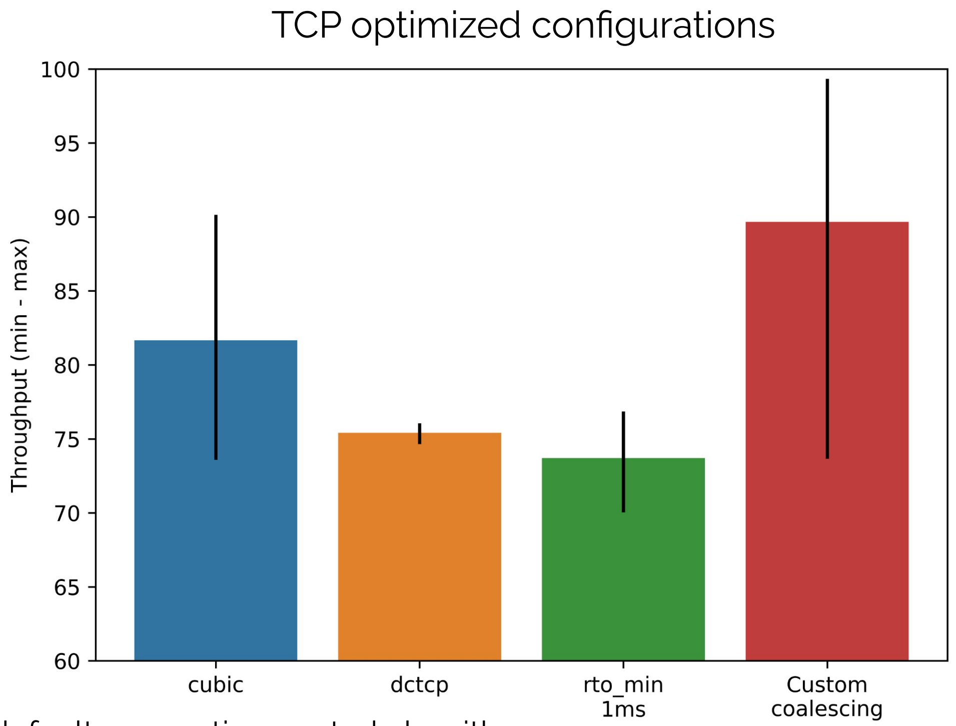

cubic: default congestion control algorithm
dctcp: uses Explicit Congestion Notification (ECN) to assess the extent of congestion and react more efficiently
rto_min: minimum retransmission timeout set to 1ms (default is 200 ms)
custom coalescing: interrupt receiver coalescing (512 μs or 1024 frames)

Dataflow network

TCP incast

Many to 1 traffic pattern
Highly synchronized collection traffic produces instantaneous buffer *oversubscription and drops*
Under utilization of network capacity and increased latency

Possible strategies
Deep buffer switches (expensive)
Credit-based congestion control (latency)
Supervisor controlled assignments

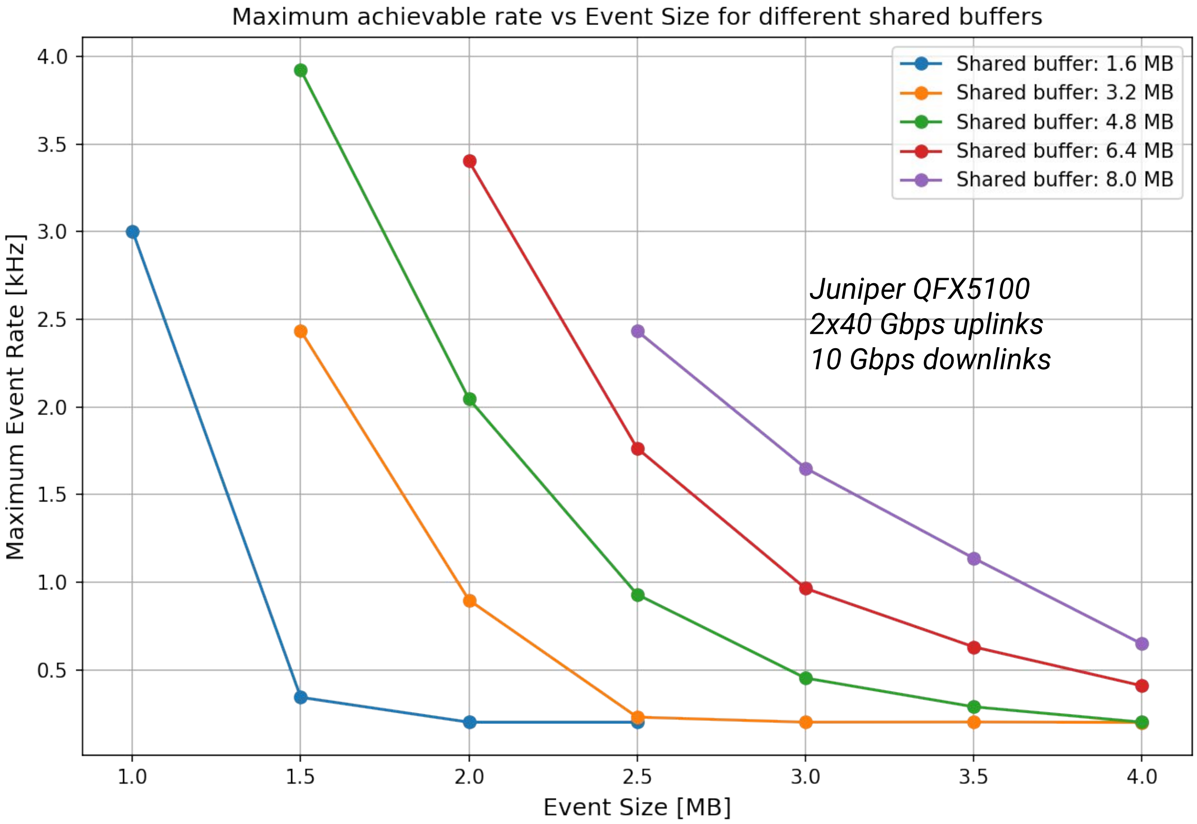

Online Software

Kubernetes Orchestration

Online Software

Manages and operates the components of the system integrating them with the wider ATLAS data taking environment

Highly distributed system
Tens of thousands of applications to supervise

Large IT infrastructure
The Event Filter farm only will consist of more than 3000 computing nodes

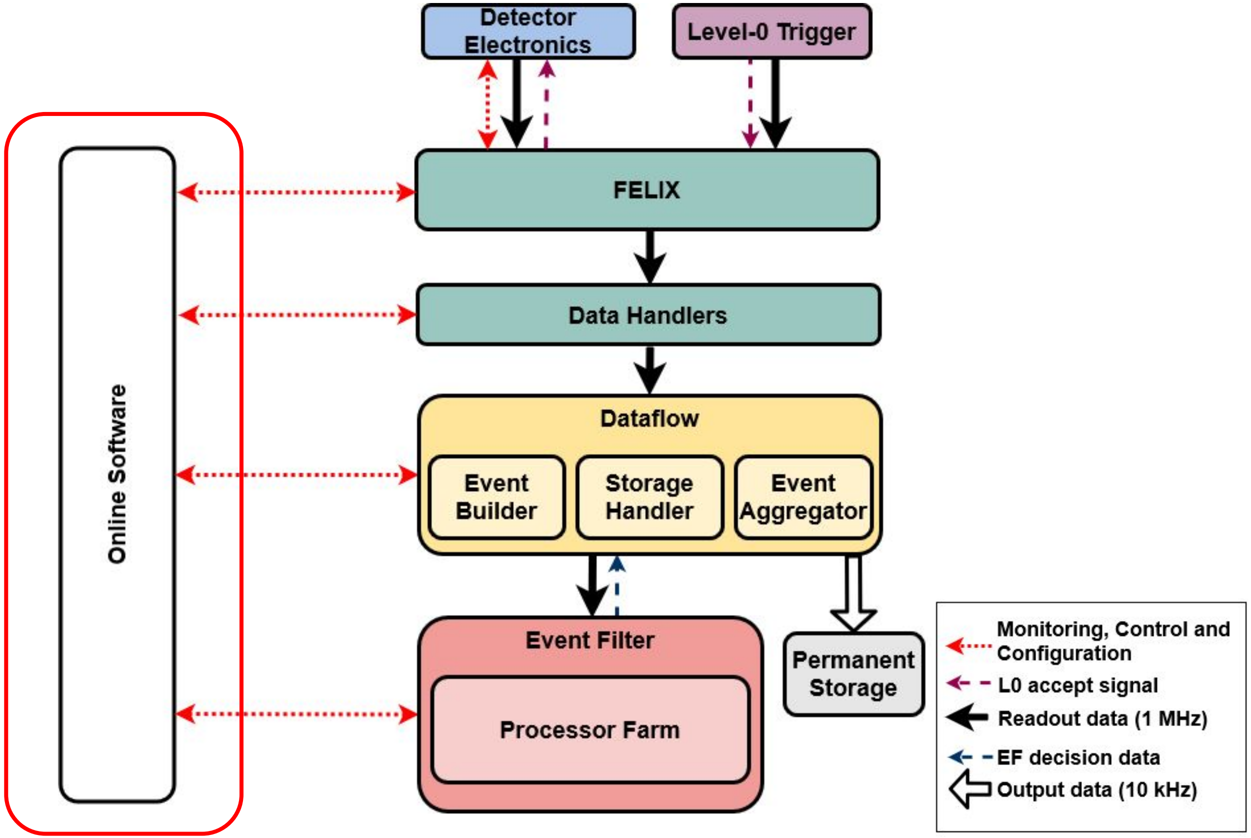

EF Orchestration
Requirements

- Support different types of application life-cycles: always-running, run-to-completion and cron-like services
- Allow dynamic and static allocation of processes to computing nodes
- Able to dynamically handle cluster resources (enabling/disabling computation units at runtime, efficient exploitation of the available CPU power and memory)
- Scale to thousands of hosts
- Control and monitor the status of all active processes

A robust and reliable mechanism for the **management of all processes running** in the Event Filter (EF) farm is a requirement to guarantee stable and effective execution of the EF service

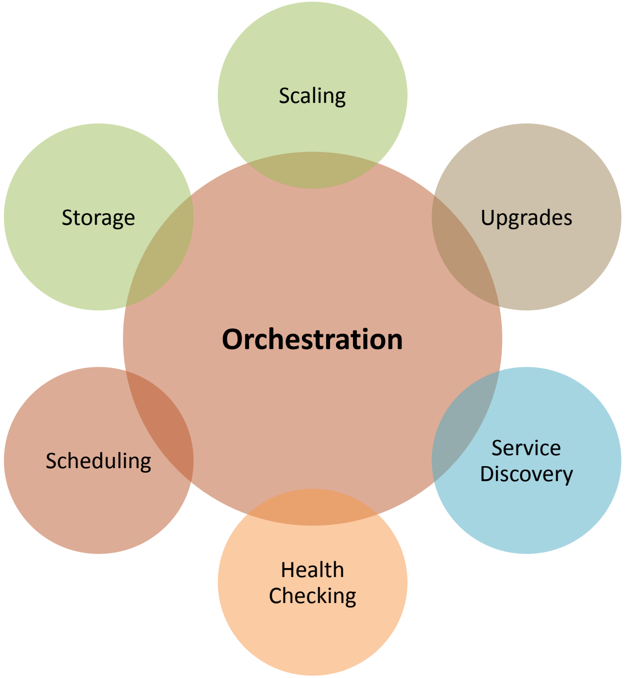

Kubernetes

An open-source system for automating deployment, scaling, and management of containerized applications

Has been identified as an **excellent candidate** for the EF farm orchestration

Features and services for application management

- Scheduling of applications based on required resources and other constraints
- Automatic re-scheduling of applications when the application itself fails or the node where the application is running dies
- Built-in support for service discovery and load-balancing
- Management of several storage back-ends, allowing transparent mounting of both local and network storage volumes
- Easy (via command line tools or GUIs) and automated (based on CPU usage) scaling of the number of application instances
- Full exploitation of containerization technologies
- Live project, large community, comprehensive eco-system

Kubernetes

Current Status

Installation at the ATLAS experimental site
Kubernetes could successfully orchestrate a 2500 servers farm

Performance Evaluation
Container start and stop times for large deployments
Simple “pause” containers pre-loaded on the worker nodes

Managed to start about 25000 containers in less than 80 seconds

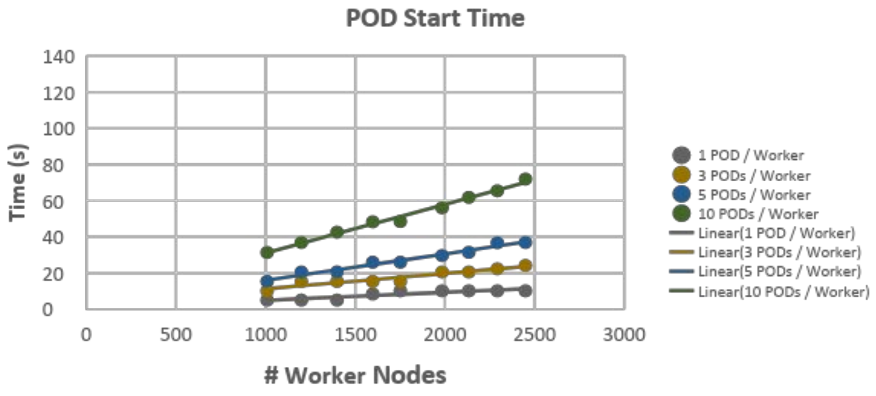

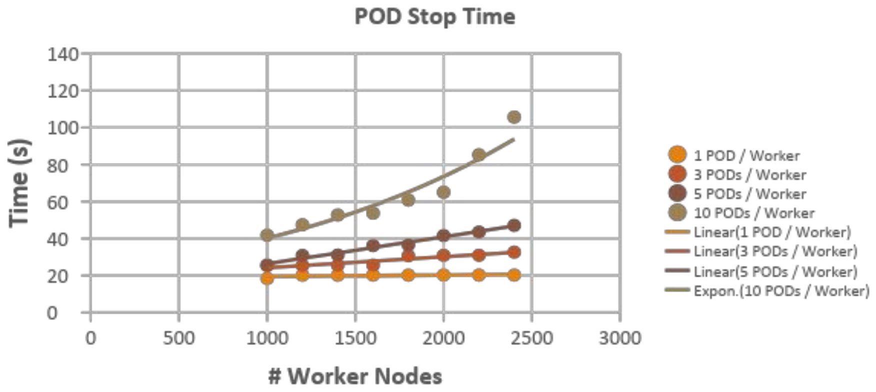

Conclusions

- HL-LHC upgrade requires a major upgrade of all components of the ATLAS TDAQ

- Impact of technology evolution
  - Full event building thanks to network improvements
  - Persistent storage design on hold due to increased storage cost and endurance updates
  - RDMA-based readout network
  - Management of processing farm with Kubernetes orchestration

- Upcoming challenges
  - Close interaction with industry to track technology updates
  - Incorporate experience from run 3 operations

Thanks

Questions?