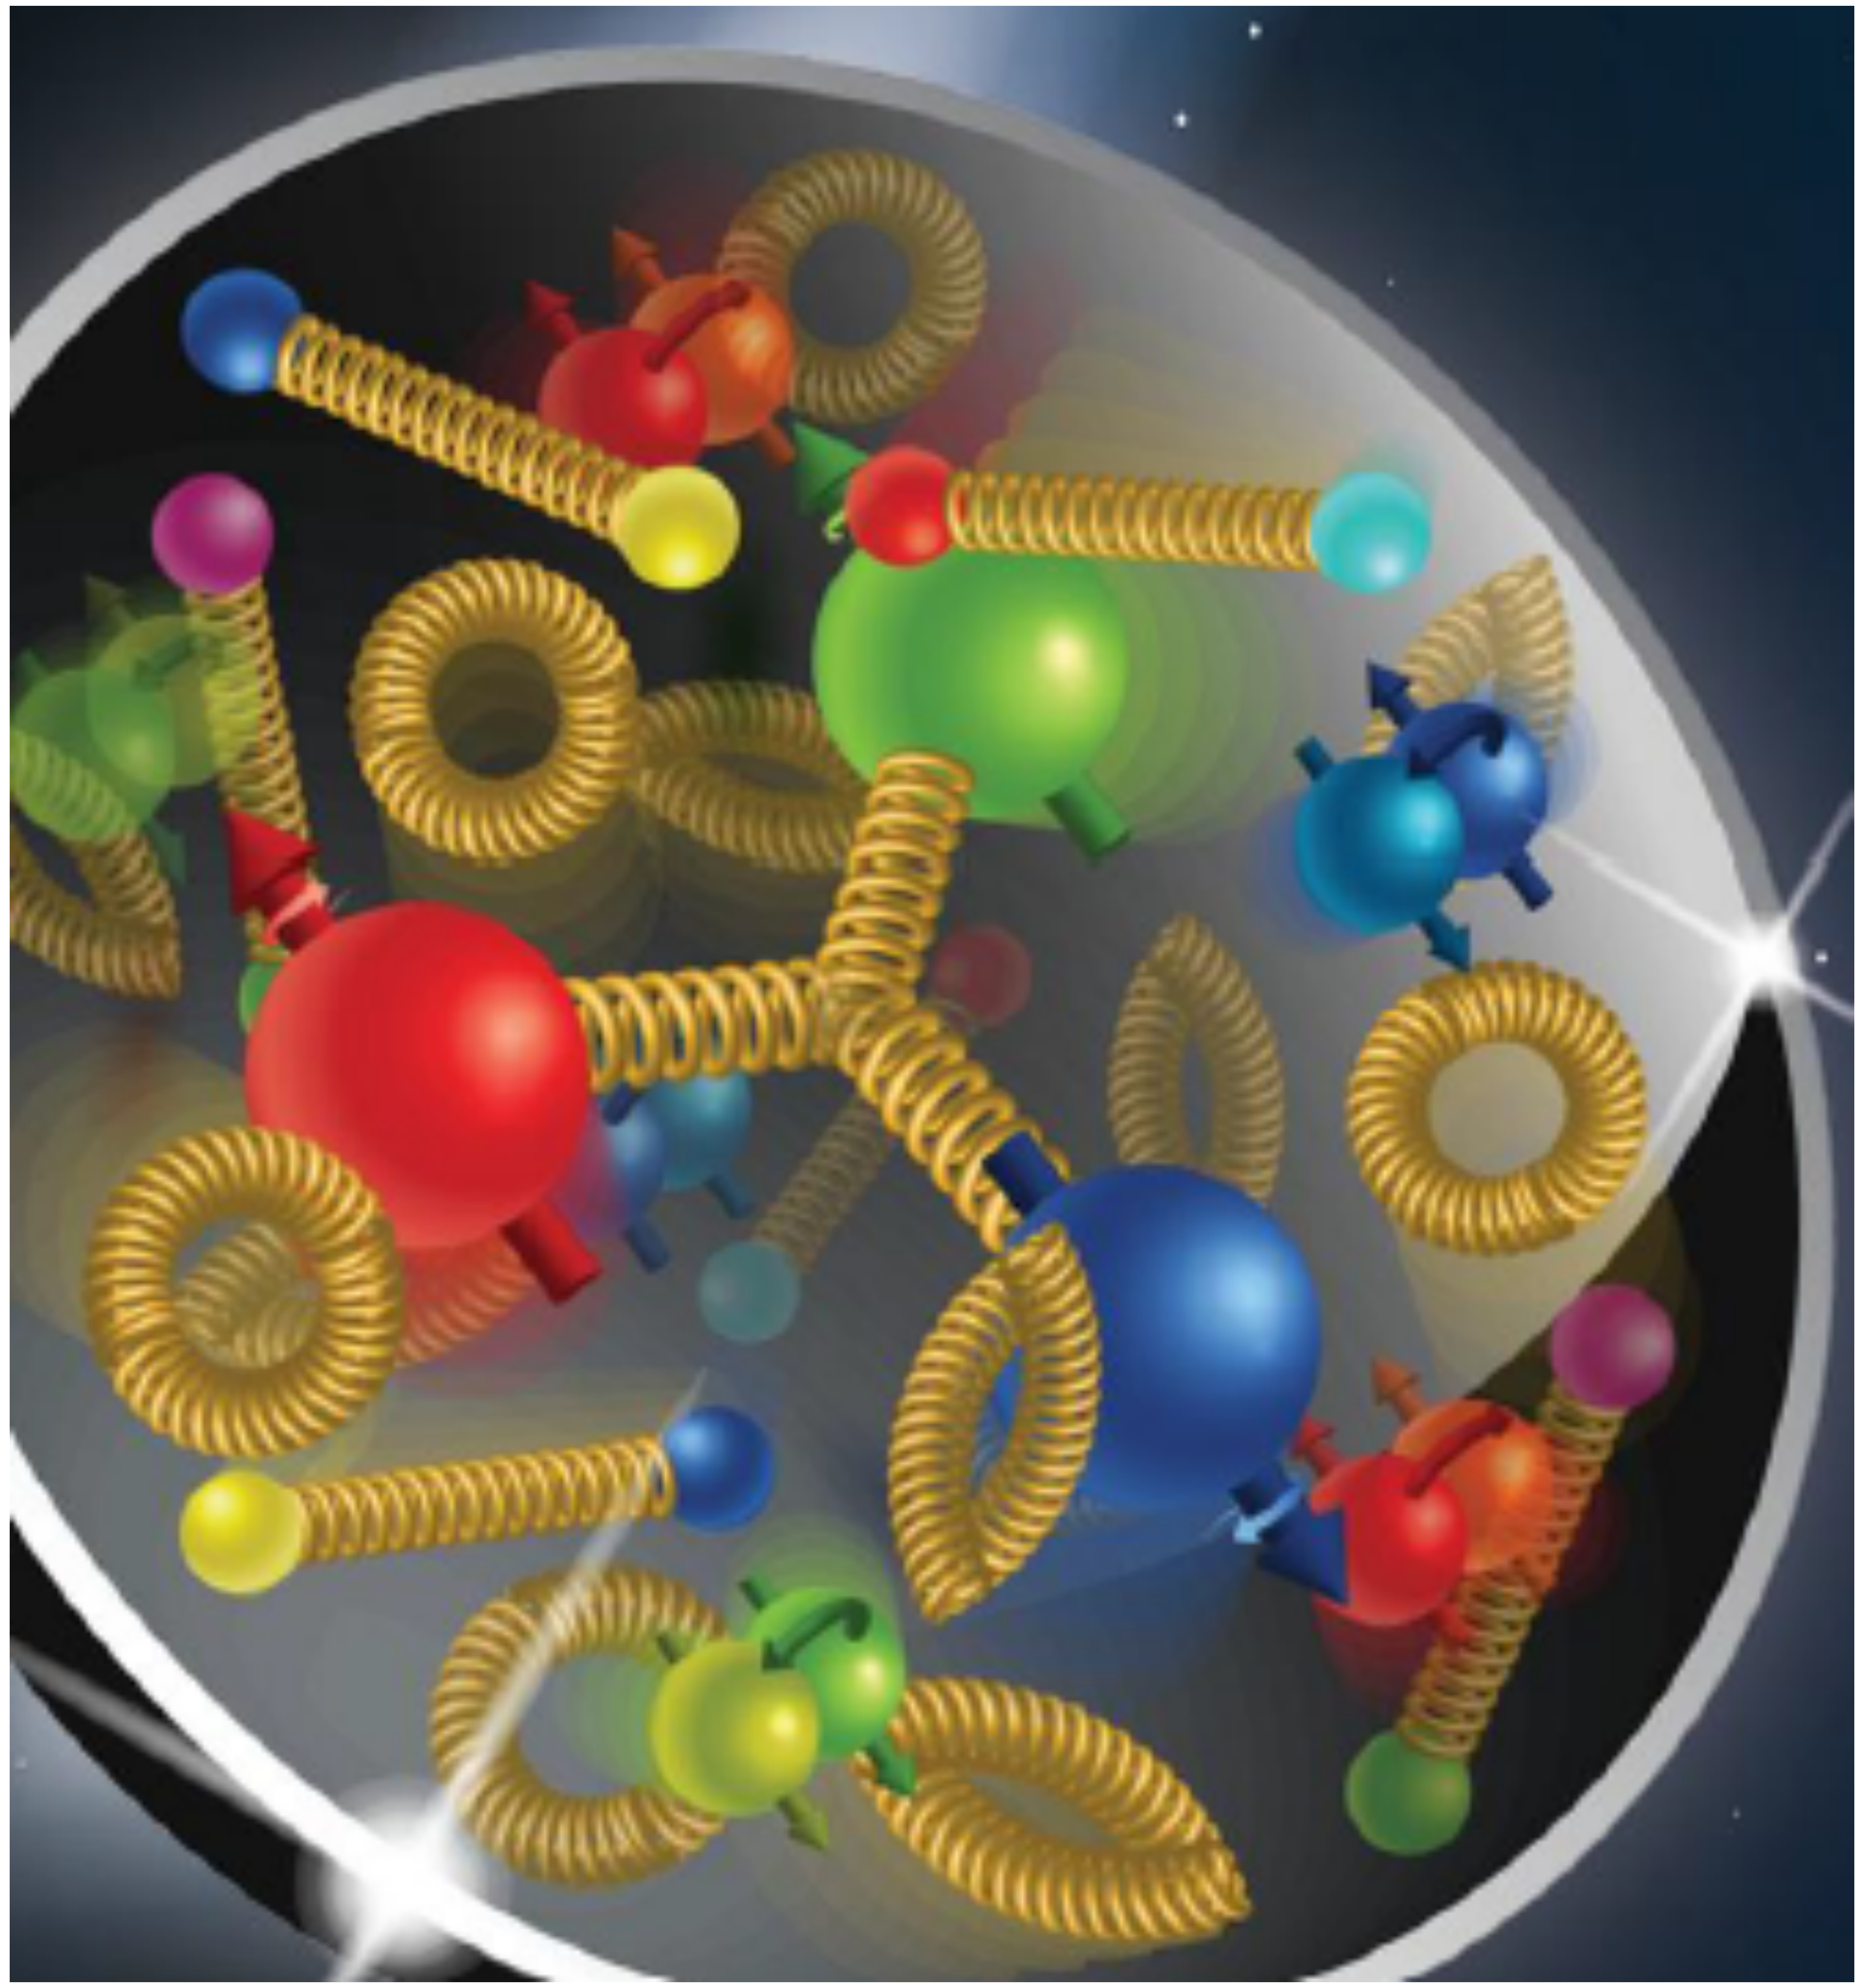

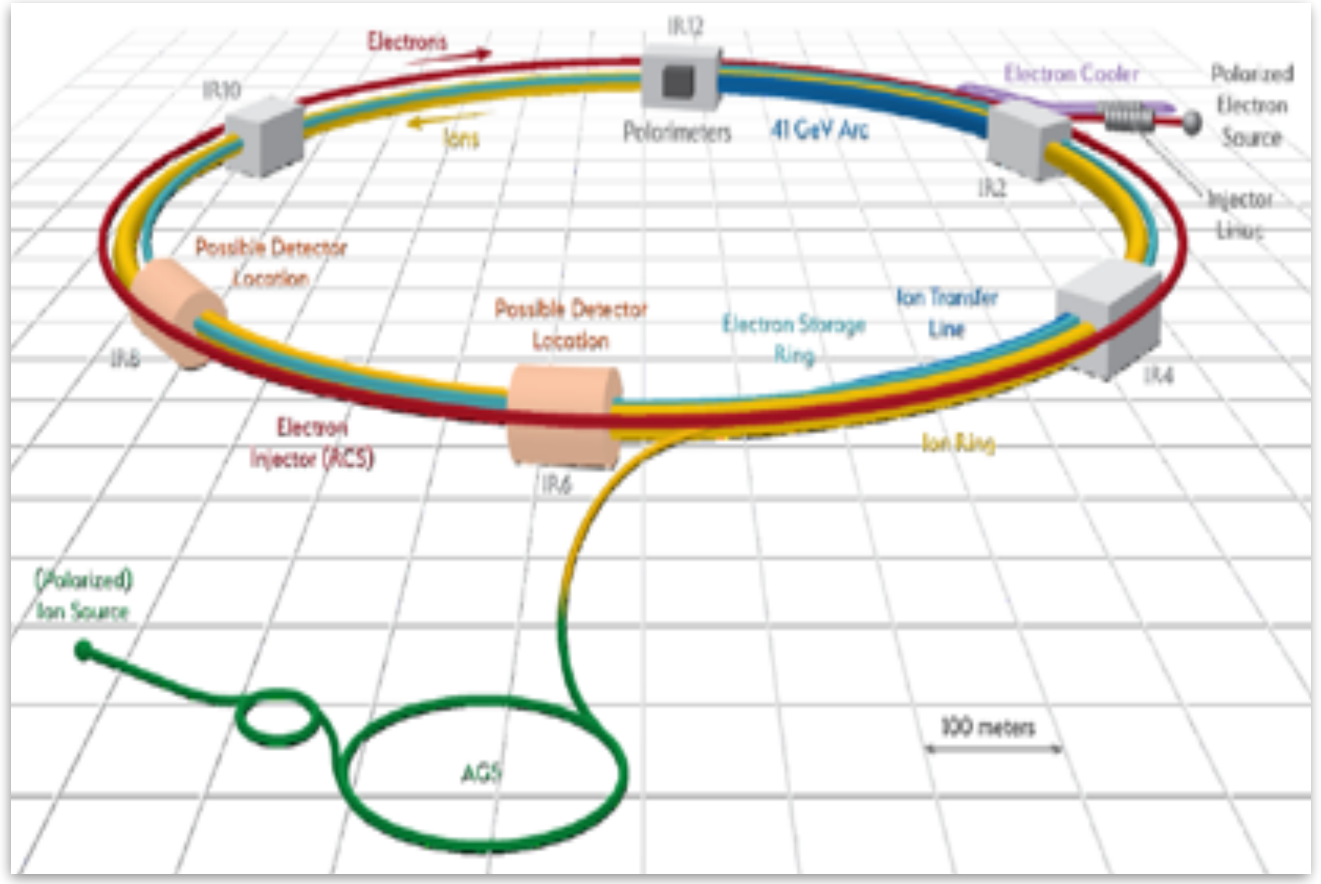

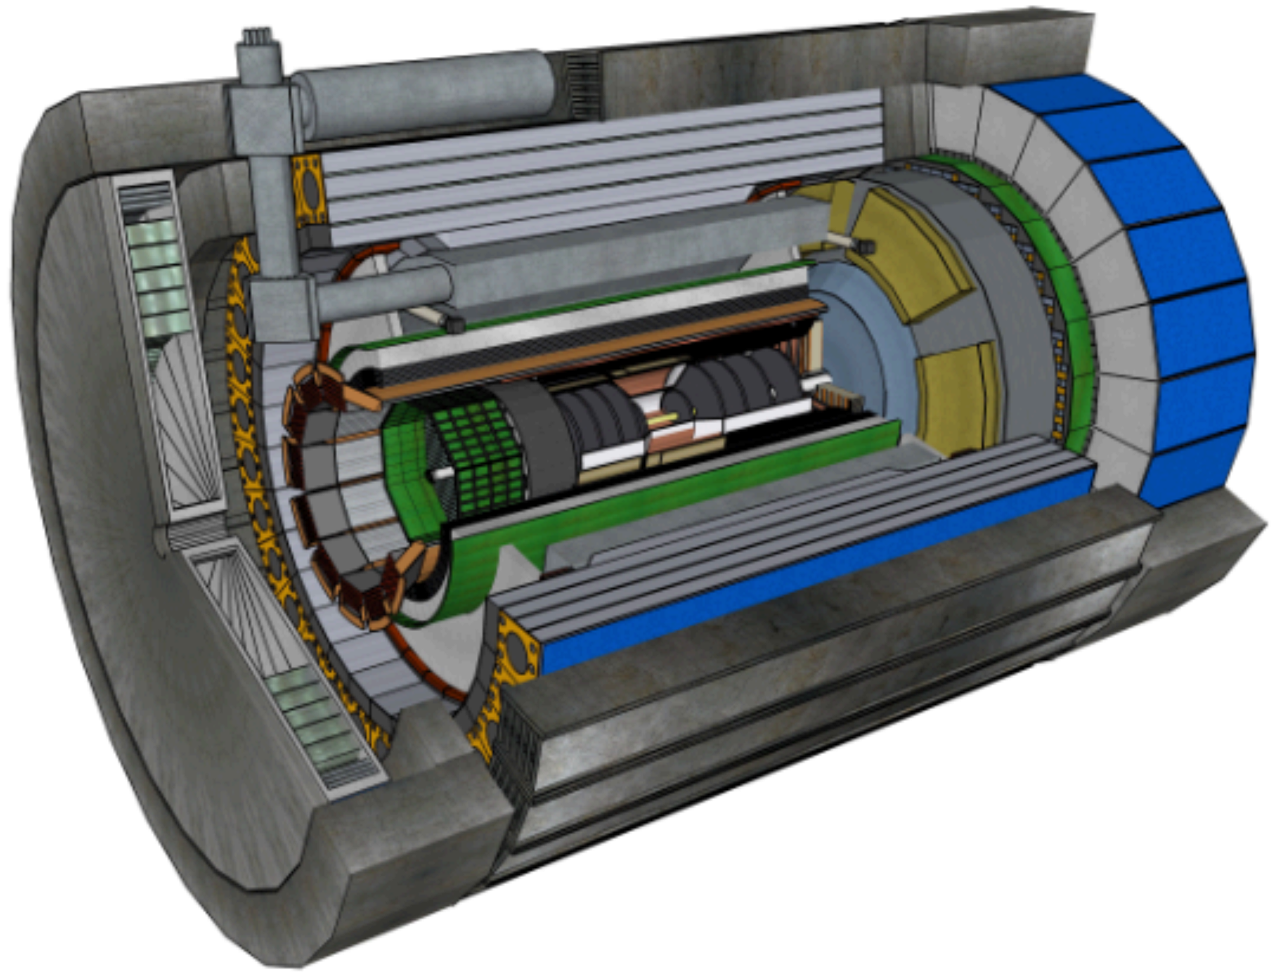

SRO Project update

M.Battaglieri (INFN), M.Diefenthaler (JLab), J.Huang (BNL), J.Landgraf (BNL), T.Wenaus (BNL)

Streaming readout (SRO) WG

ePIC Streaming Computing Model Working Group

Kick-off meeting - July 11, 2023

Software and Computing Coordinator (Markus)
+ Deputy Coordinator **Operations (Wouter)**
+ Deputy Coordinator **Development (Sylvester)**
+ Deputy Coordinator **Infrastructure (Torre)**

Guiding Principles: DE&I, Software Principles, Sustainability

Operation WGs:
- **Production (CD)**
- **User Learning**
- **Validation (CD)**

Development WGs (CI):
- **Physics and Detector Simulation**
- **Reconstruction**
- Analysis Tools

Infrastructure WGs:
- **Streaming Computing Model**
- Multi-Architecture Computing
- Distributed Computing

Cross-cutting WG:
- Data and Analysis Preservation

- Structure of Software and Computing in ePICS
- Presented by Markus on May 2nd 2023 kickoff-meeting

- ePIC SRO Computing Model WG belongs to 'infrastructure'
- Co-conveners: M.Battaglieri (INFN), J.Huang (BNL) + J.Landgraf (BNL)

Marco Battaglieri
Co-convener

Jin Huang
Co-convener

Jeffery Landgraf
Co-convener for electronics & DAQ WG
Kindly helping SRO group organization during next months when Jin focuses or sPHENIX commissioning

Jin Huang (BNL) replaced by Taku Gunji (TokyoU)

- Weekly Tuesday Zoom meeting at 9AM ET https://indico.bnl.gov/event/28514/
- Resources: https://docs.google.com/document/d/1t5vBfgro8Kb6MKc-bz2Y67u3cOCpHK4dfepbJX-nEbE/edit?tab=t.0#heading=h.y3evqgz3sc98

ePIC Streaming Computing Model Working Group

Compute-Detector Integration to Maximize and Accelerate Science

- **Maximize Science** **Capture every collision signal, including background.**
  - Event selection using all available detector data for **holistic reconstruction**:
    - **Eliminate trigger bias** and provide accurate estimation of uncertainties during event selection.
  - Streaming background estimates ideal to **reduce background** and related systematic uncertainties.
- **Accelerate Science** **Rapid turnaround of 2-3 weeks for data for physics analyses.**
  - Timeline driven by alignment and calibration.
  - Preliminary information from DSCs indicates that 2-3 weeks are realistic.
- **Technologies** Compute-detector integration using:

Streaming readout for continuous data flow of the full detector information.	**AI** for autonomous alignment and calibration as well as autonomous validation for rapid processing.	**Heterogeneous computing** for acceleration (CPU, GPU).

ePIC Collaboration Meeting, January 20, 2025. 11

ePIC

Use Case	Echelon 0	Echelon 1	Echelon 2	Echelon 3
Streaming Data Storage and Monitoring	✓	✓		
Alignment and Calibration		✓	✓	
Prompt Reconstruction		✓		
First Full Reconstruction		✓	✓	
Reprocessing		✓	✓	
Simulation		✓	✓	
Physics Analysis		✓	✓	✓
AI Modeling and Digital Twin		✓	✓	

Why SRO is important for ePIC?

- ✳ **Expedite ePIC Science outcome**
  - Computing/detector integration for a rapid turnaround
- ✳ **SRO breaks barriers between DAQ and computing**
  - Data collected at Echelon0 and processed at BNL and JLab
  - Important role of Echelon2 and Echelon3

The ePIC Streaming Computing Model

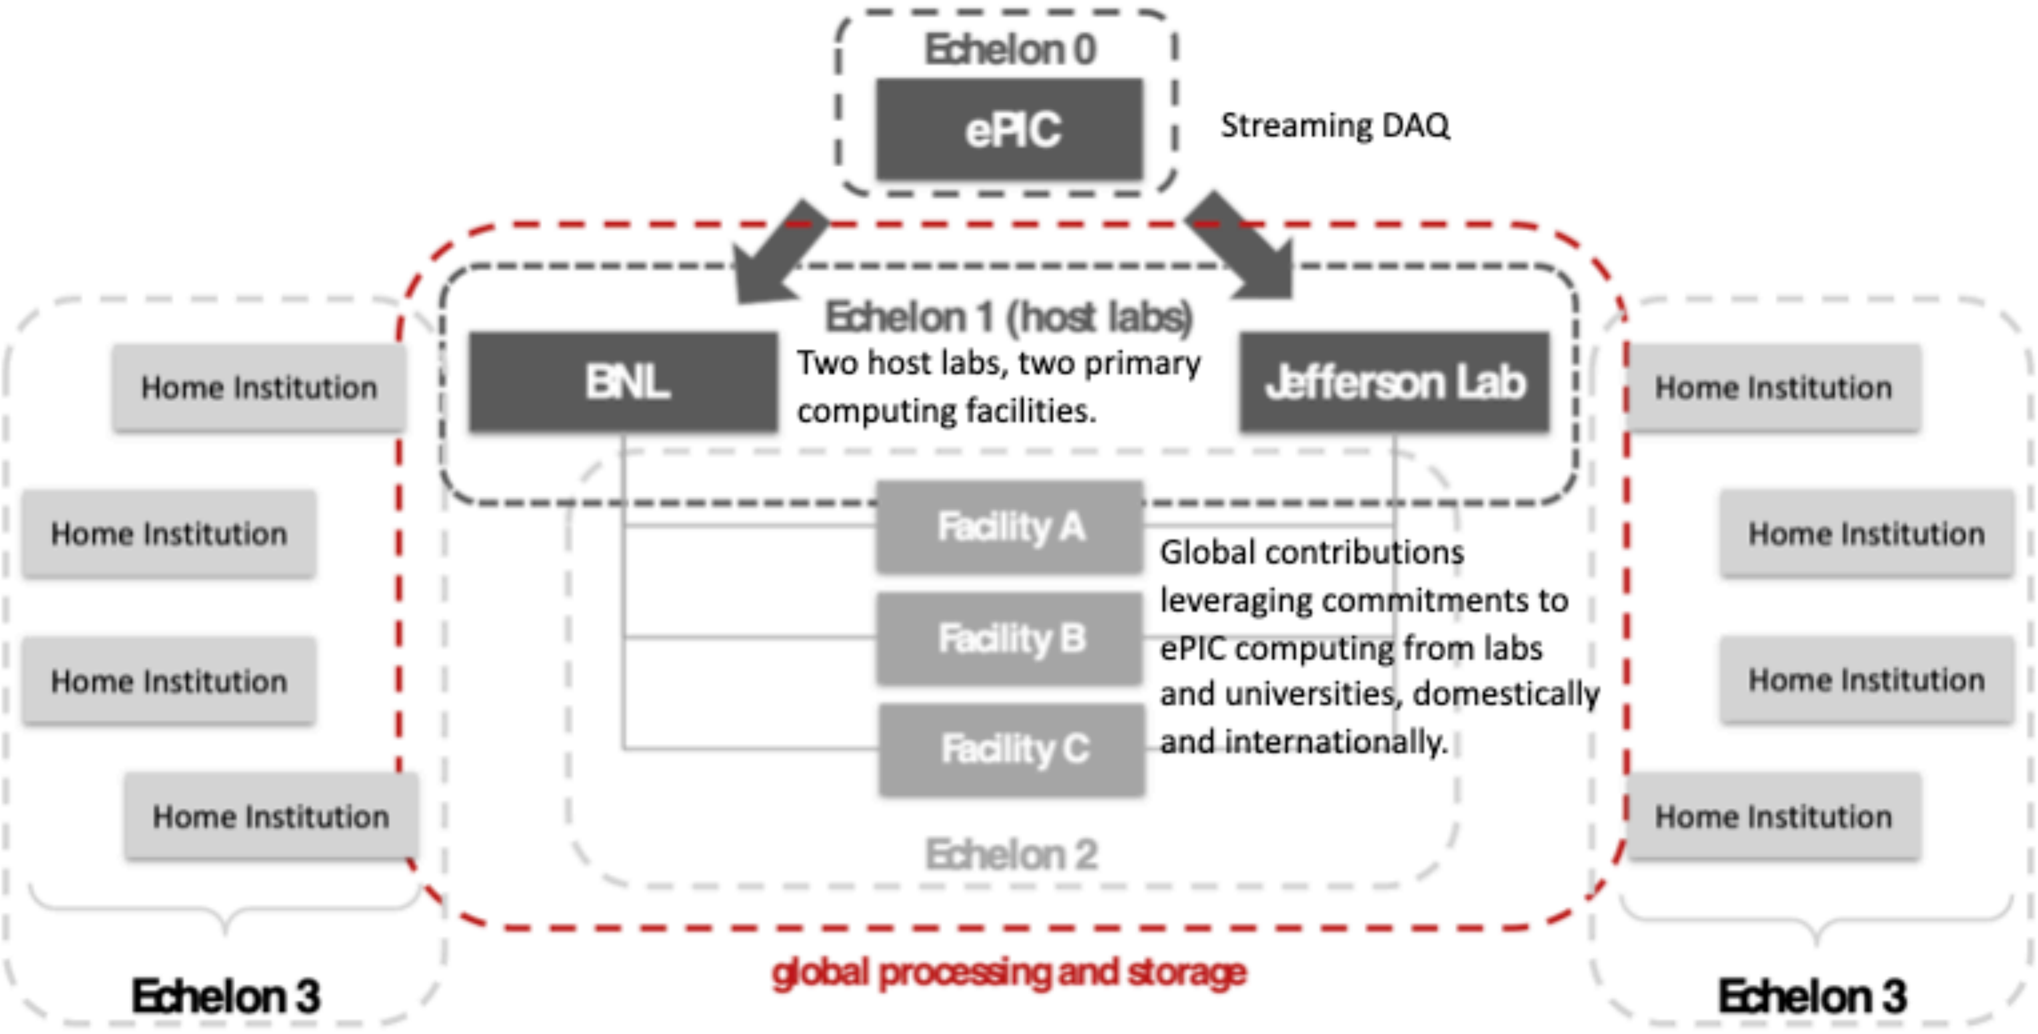

Supporting the analysis community *where they are* at their home institutes, primarily via services hosted at Echelon 1 and 2.

ePIC Collaboration Meeting, January 20, 2025. 12

ePIC

Streaming RO in ePIC

ePIC DAQ uses streaming architecture

DAQ

Definition of streaming is "No L0 trigger"

- All data is zero suppressed by the front end electronics
- No system wide deadtime in normal operation
- Collaboration should have the full ability to make data selection cuts on the widest possible criteria
  - Flexible event selection, data selection and background characterization
- But subject to an overall throughput budget of ~100Gb/sec

ePIC Streaming will include

- Capabilities for hardware and firmware based triggering
- Capability for flow control
- Zero suppression & aggregation within data packets

Greater sensitivity to noise

Greater sensitivity to backgrounds

Computing

Definition of streaming is "Process data as it arrives"

- Fast Analysis (~3 weeks not months or years) using automation of calibration and reconstruction.

Requires some overlap between DAQ/ computing

- Automation of calibrations
  - Mapping calibration dependencies
  - Mapping application of calibrations
  - Mapping evolution of calibrations
- QA and monitoring can make use of full offline structure
- Consistent schemes and language for data/metadata
- Event selection / tagging / and accounting

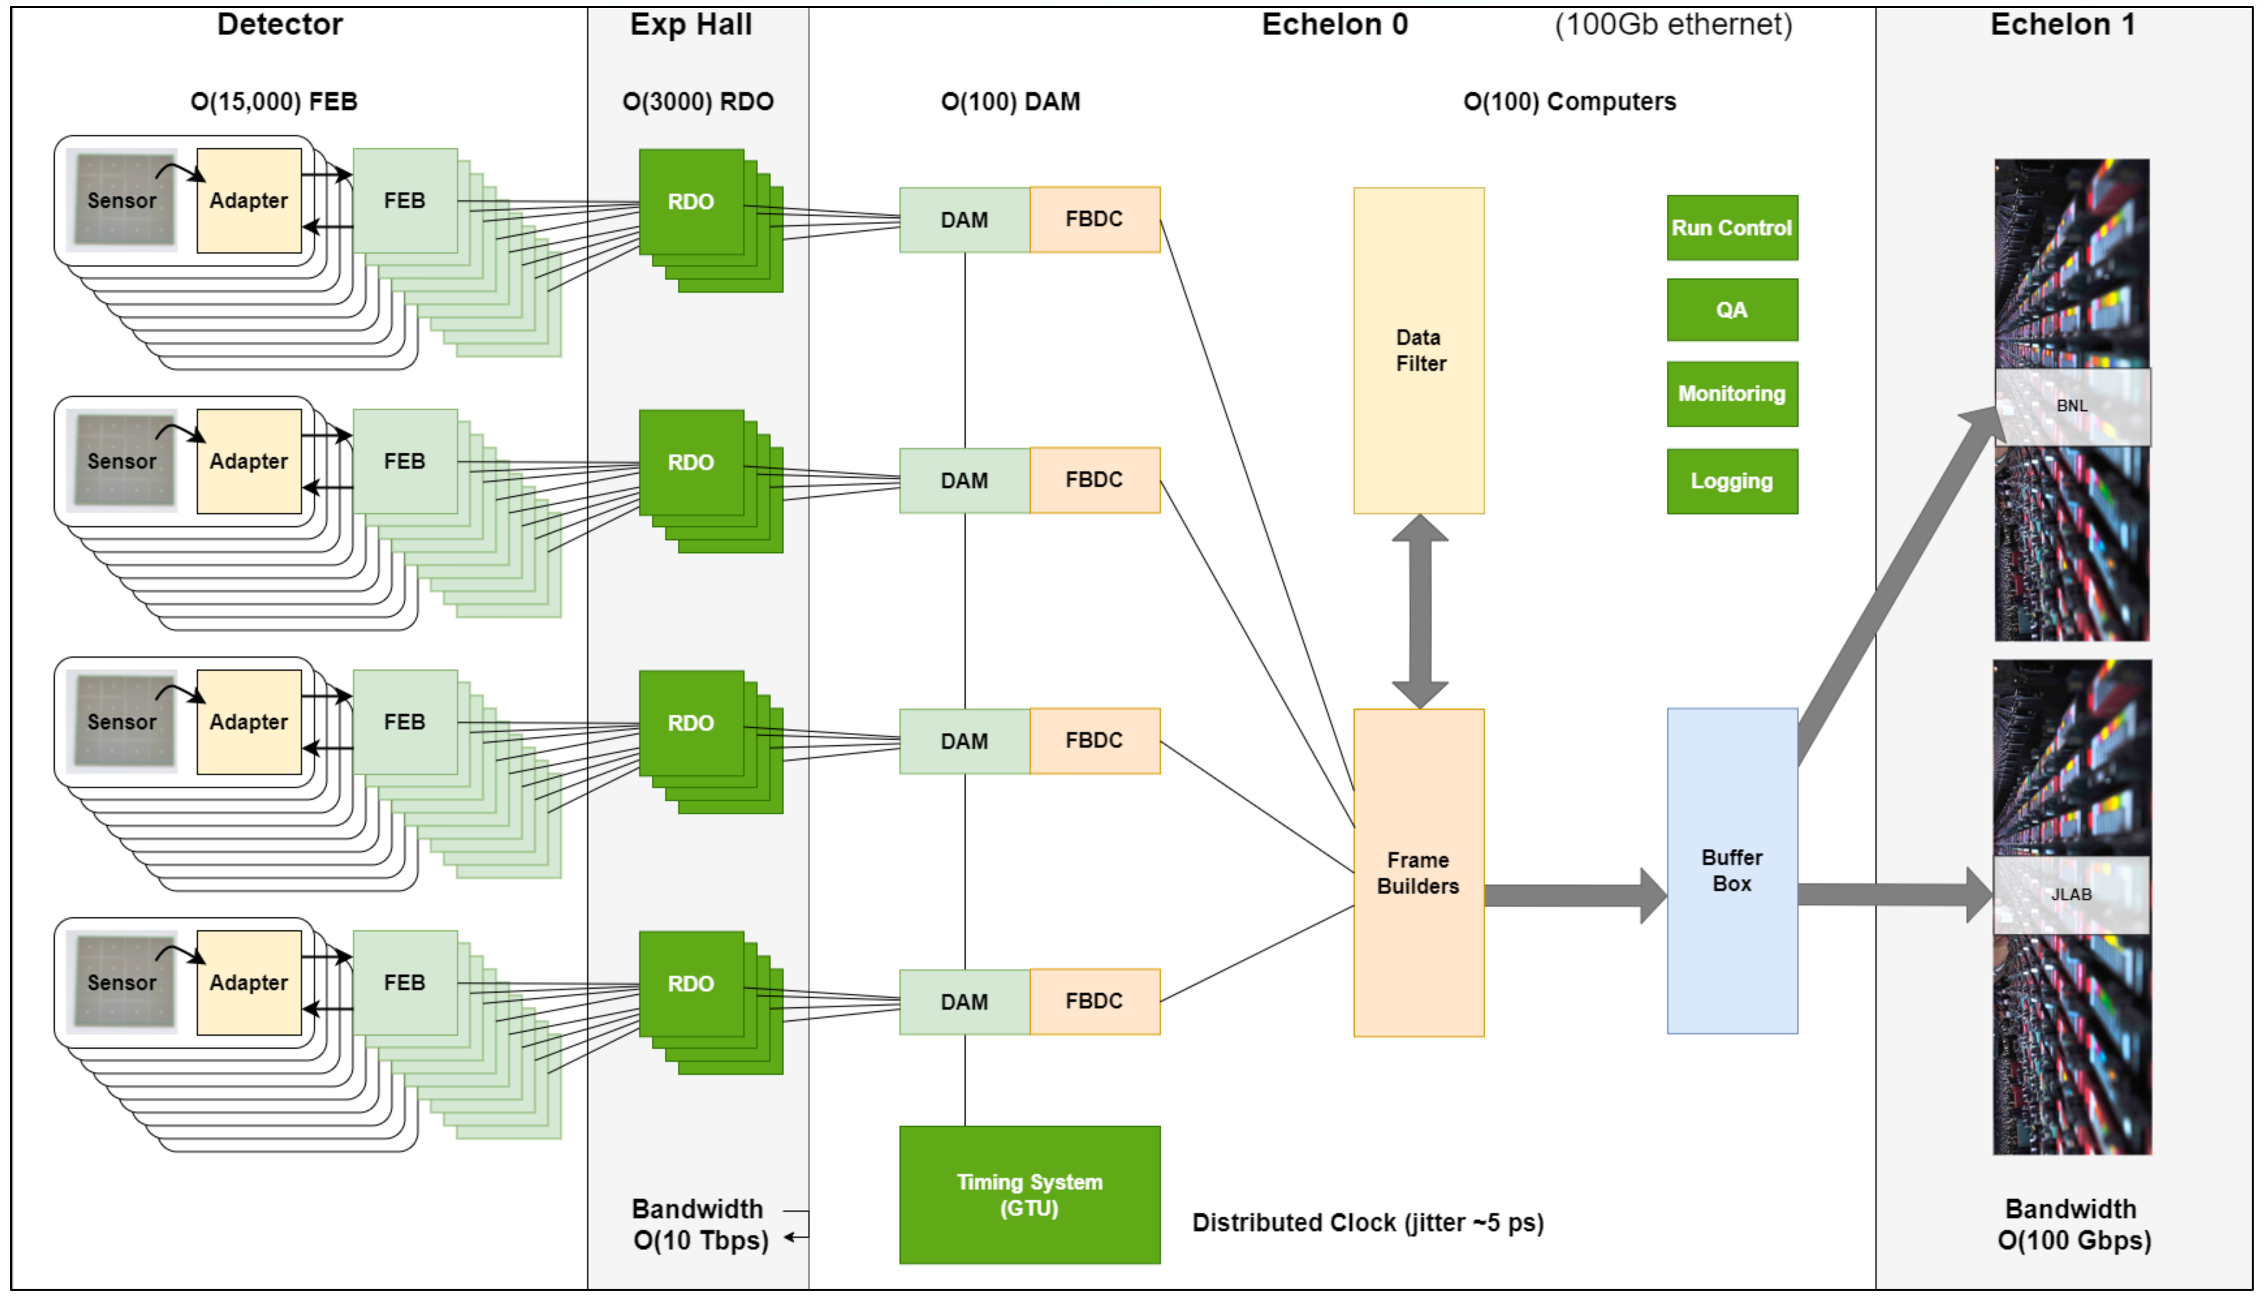

Worst case Data Flow

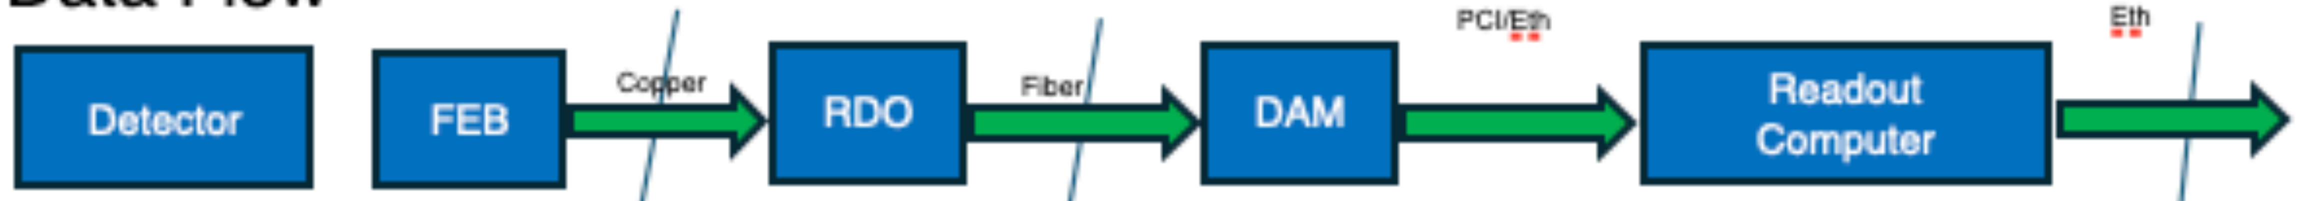

Aggregate	2.0 Tb/sec
Noise	1.6 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	2.0 Tb/sec
Per RDO (Avg)	0.7 Gb/sec

Aggregate	109 Gb/sec
Collision Signal	62 Gb/sec
Synchrotron Rad	Under design
Electron Beam	4.5 Gb/sec
Hadron Beam	1.0 Gb/sec
Noise	41 Gb/sec

"Worse Case Data Flow"

- $\mathcal{L}_{max} \Rightarrow 10^{34}$ cm^{-2} s^{-1}
- 10 x 275 GeV DIS
- DIS rate ~500kHz
- After 9 x 10^{9} (1-MeV n_{eq} / cm^{2}) delivered fluence
  - 300kHz DCR in dRICH
  - 200x rate with no damage

At Worse Case

- Need a factor of 20 reduction

Key components: FE aggregator

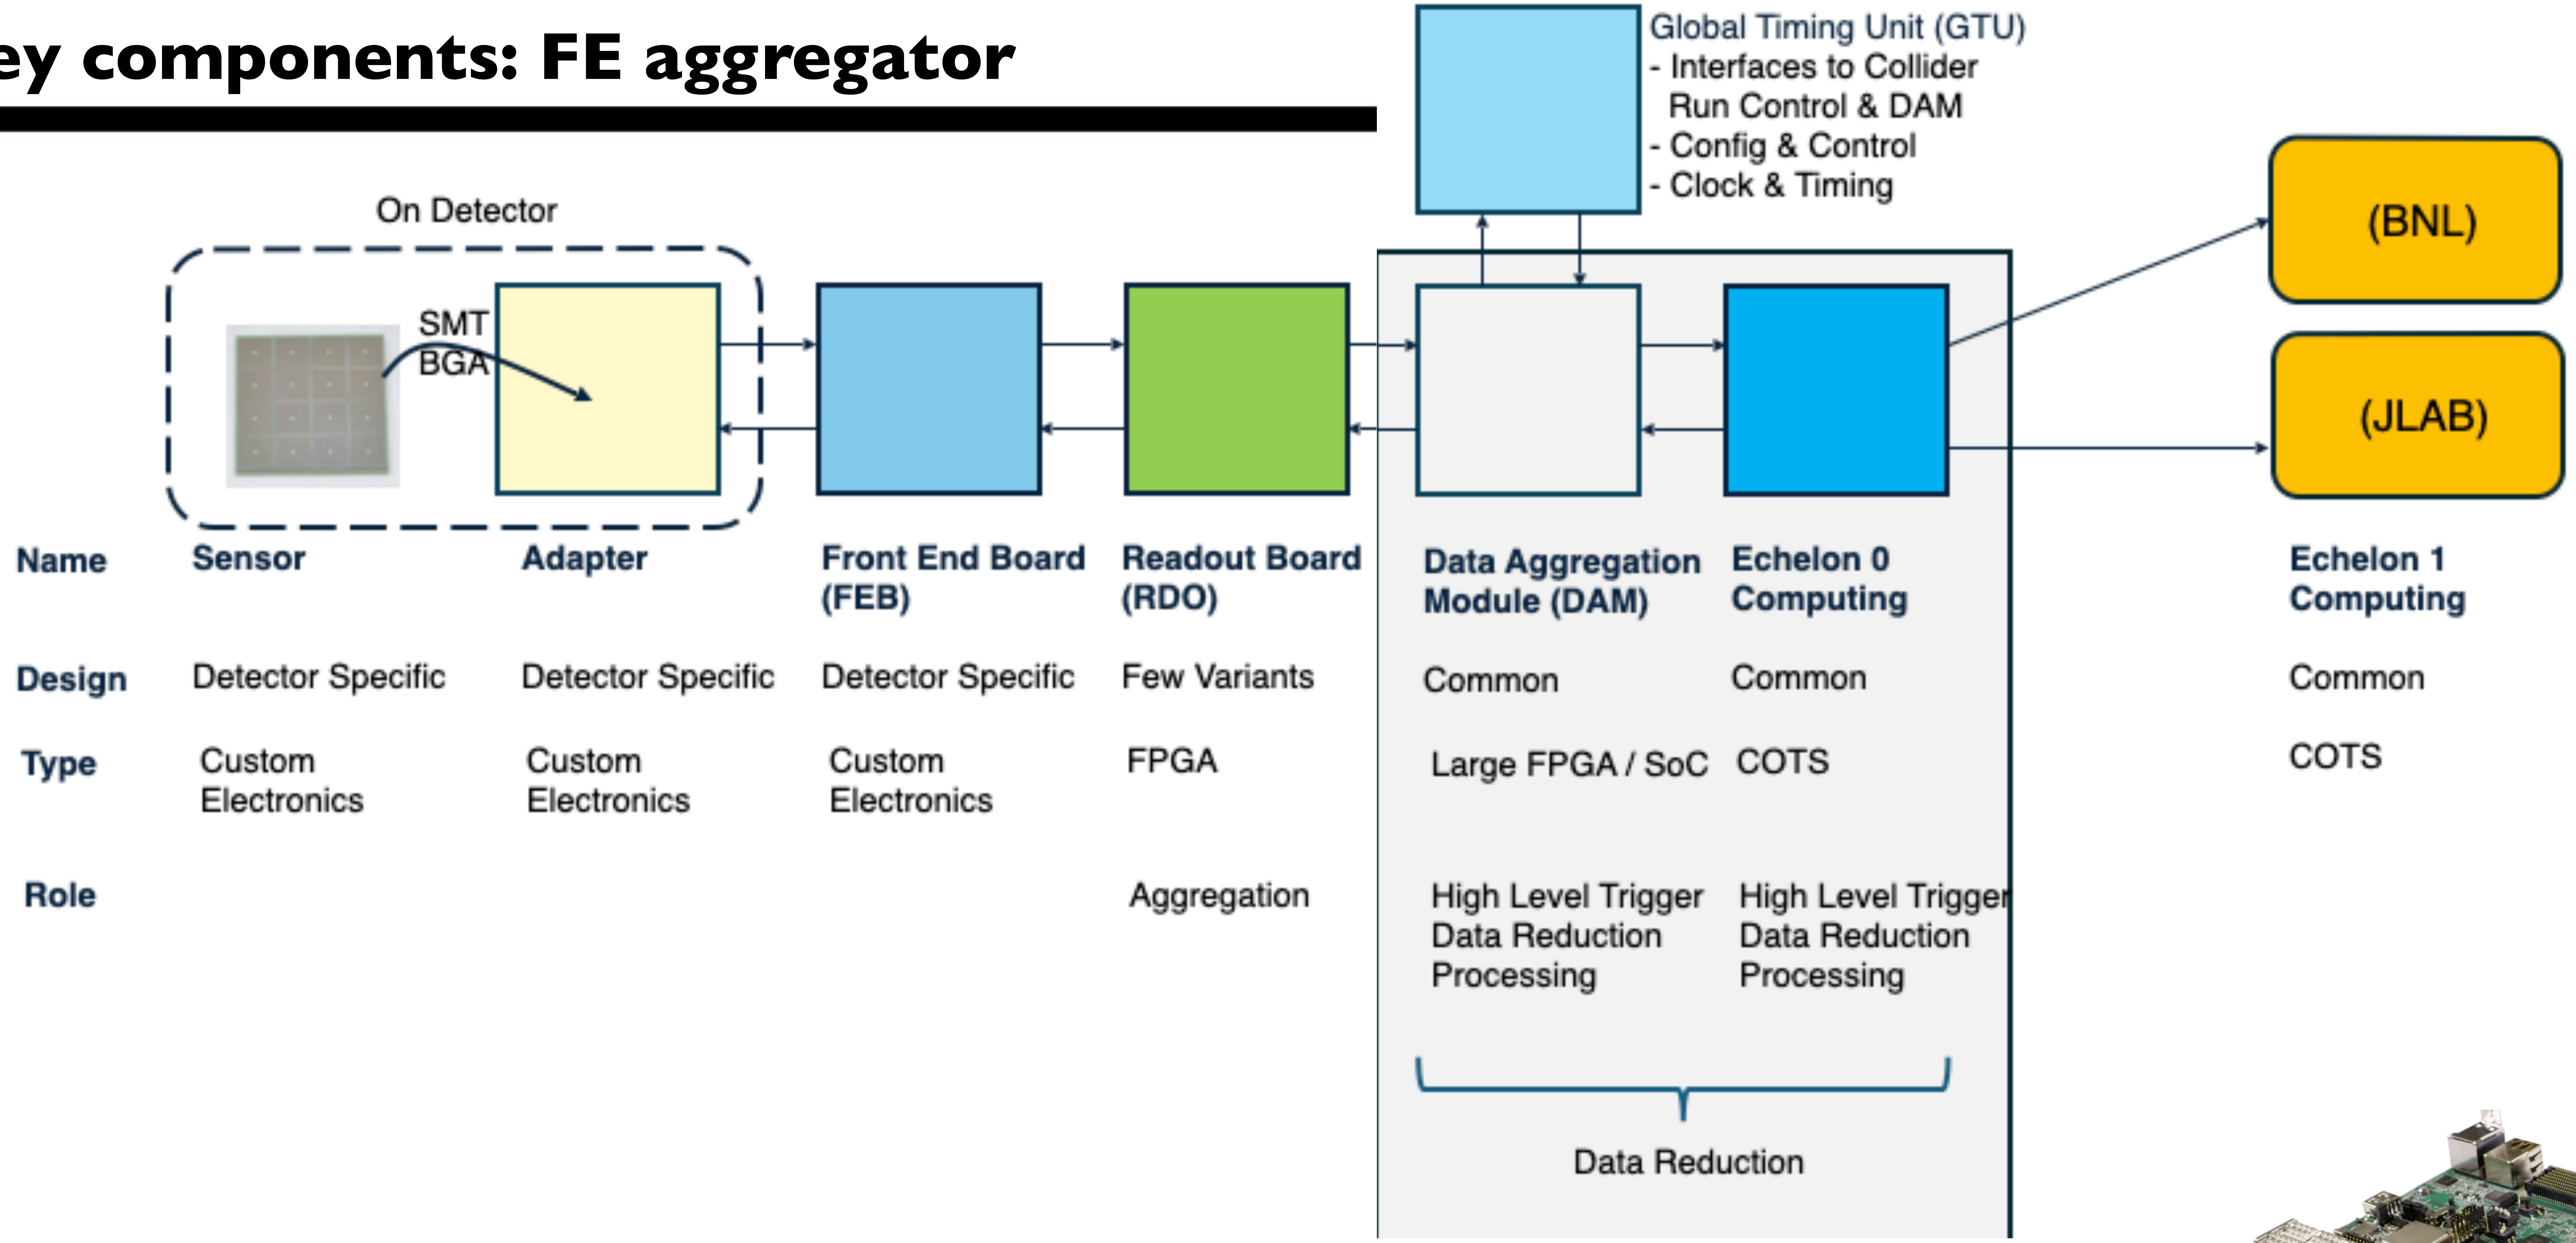

DAM Board – FLX-155

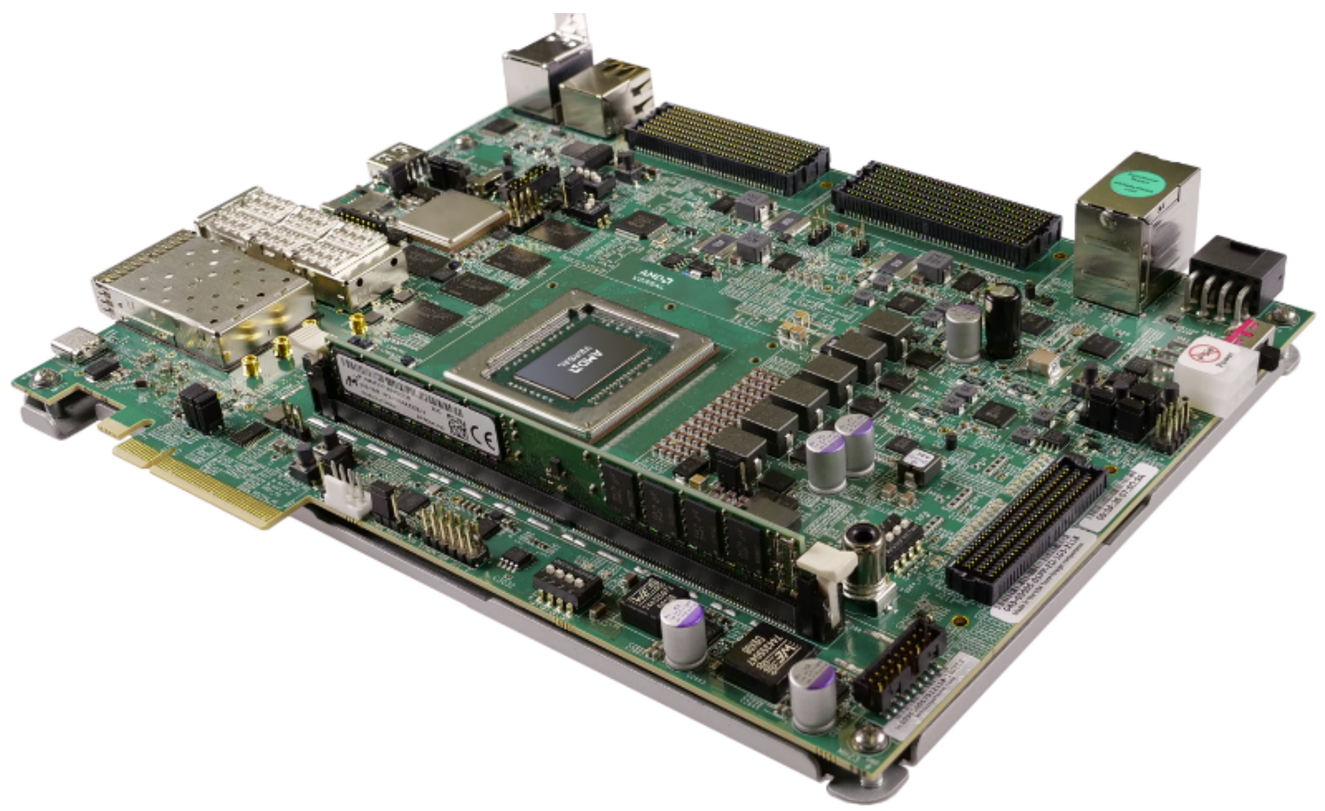

- FELIX (Front-End Link Exchange) hardware for future ATLAS experiments at HL-LHC is being developed at BNL (Omega Group)
- First boards under evaluation at BNL by Omega group

Key components: data reduction

General Principles

1. Don't remove any information unless required by throughput / storage requirements
2. If any information is removed, write out a fraction of data without removing the information (in conjunction with the reduced data)
3. Any filtering / data reduction code needs to be documented and available for access by the collaboration

Non-destructive reduction techniques:

- Compression (Standard or ML techniques)
- Reduction of ASIC data inefficiencies
- Aggregation of headers

Techniques that drop information:

- (1-d) cluster finding for ADC data
- (2-d) cluster finding
- Feature finding
- High Level Noise Rejection
- Averages / Histograms
- Simple Analysis
- Arbitrary Deadtime

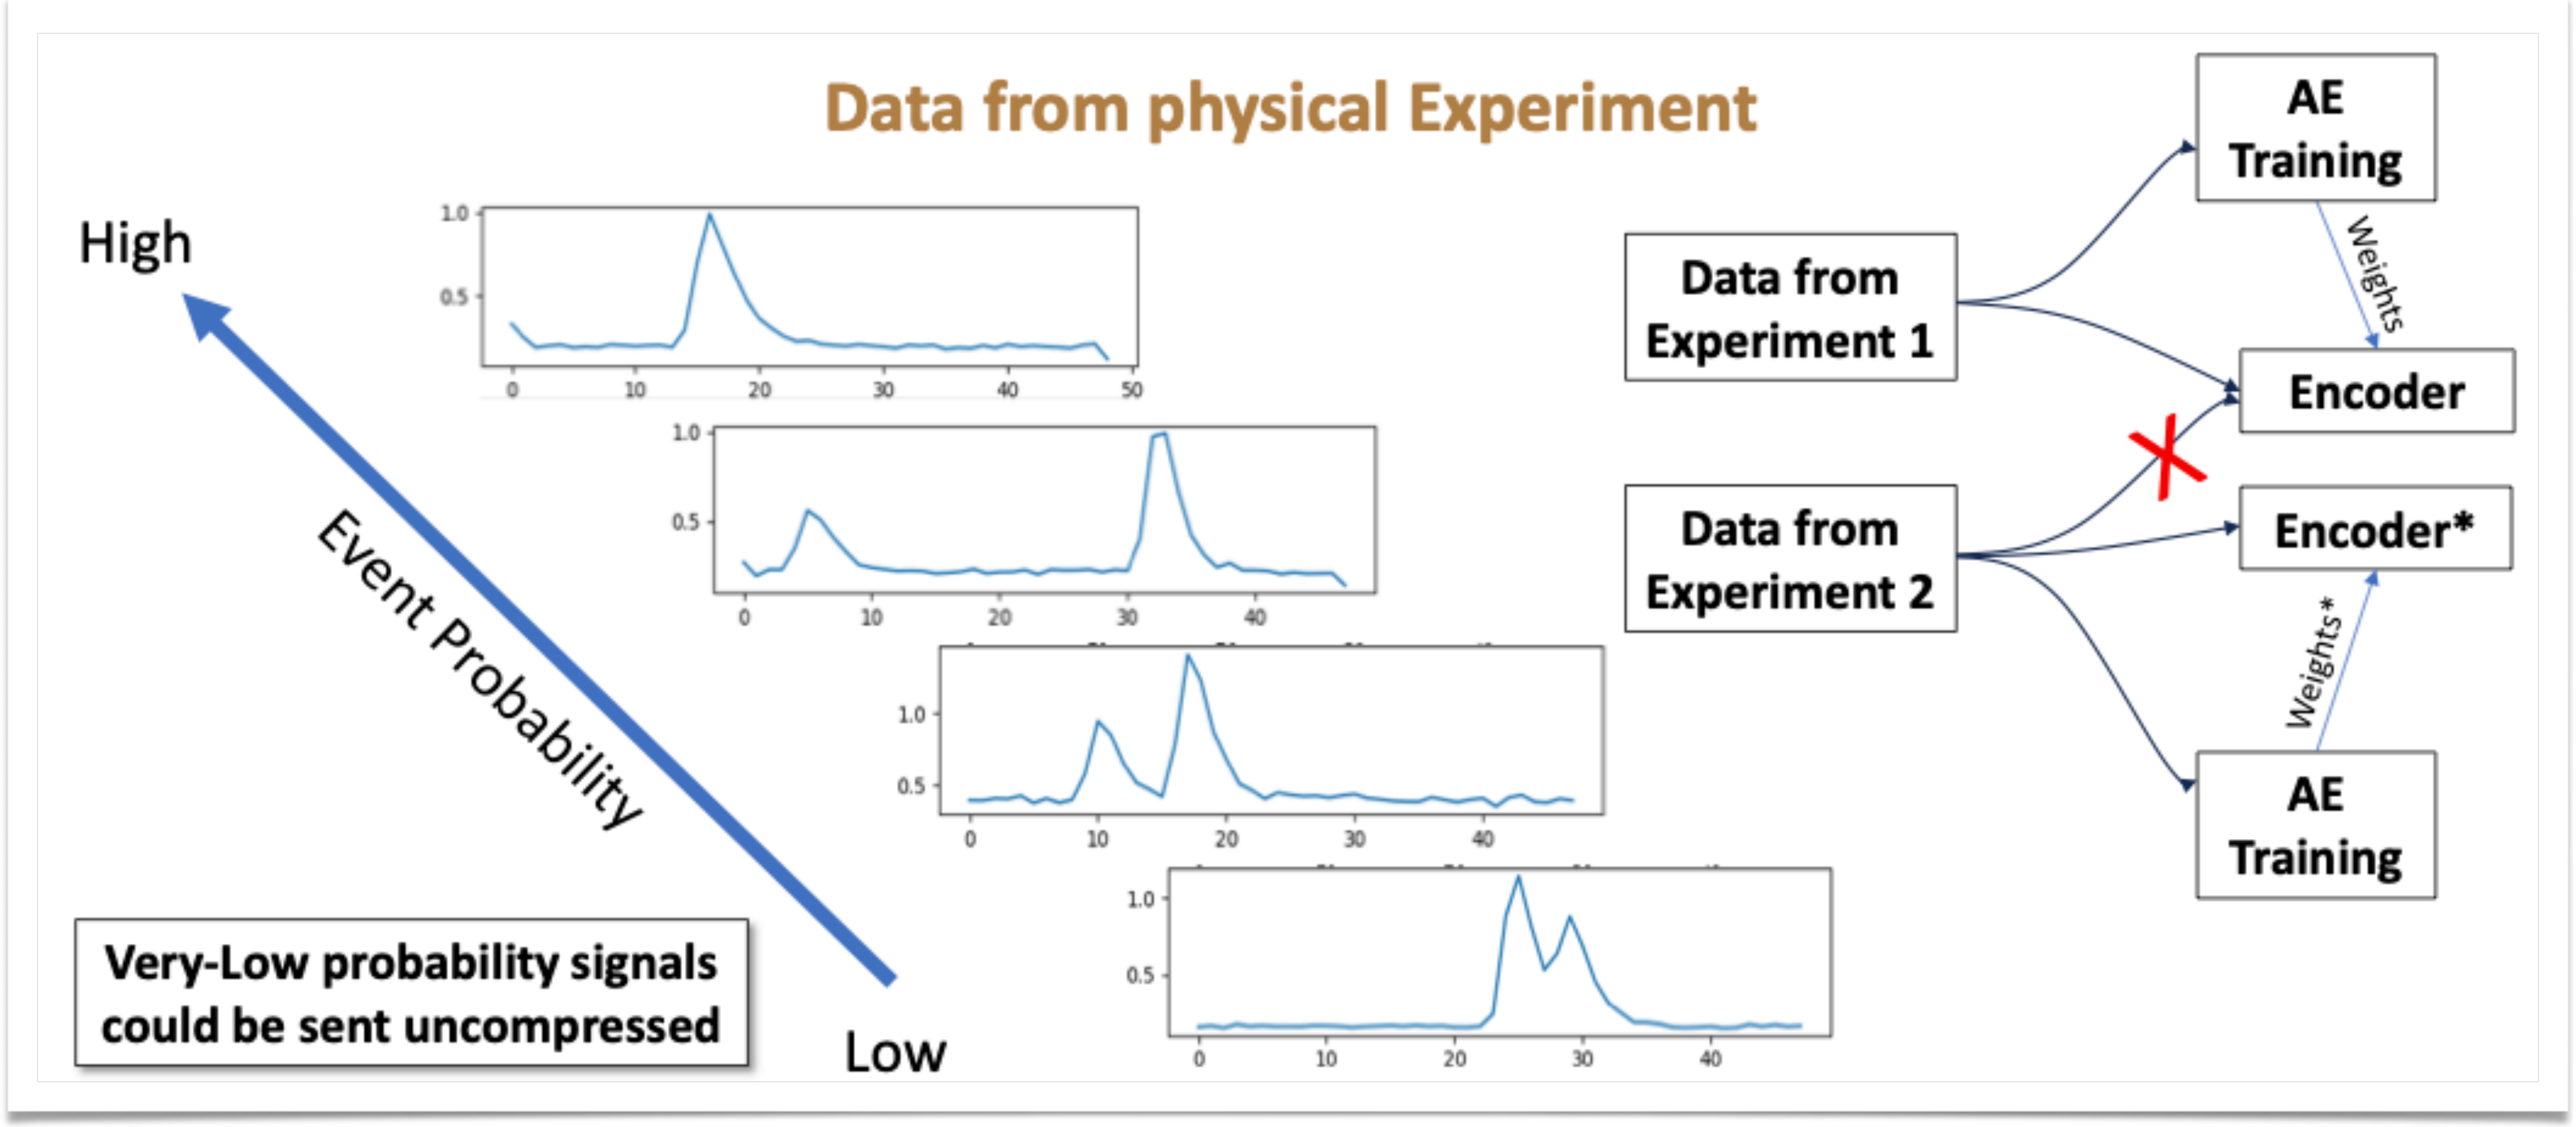

- Real-time data processing using an autoencoder to compress the digitised waveform

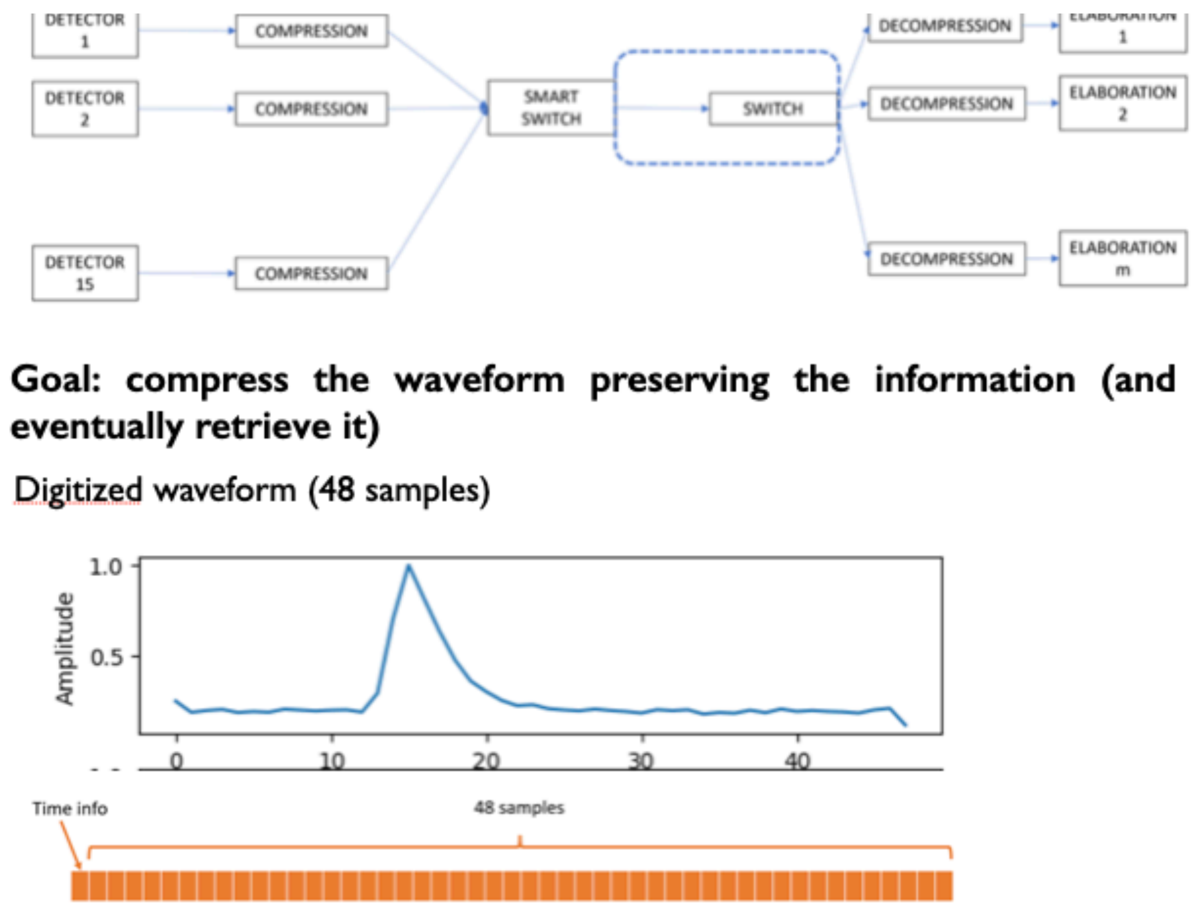

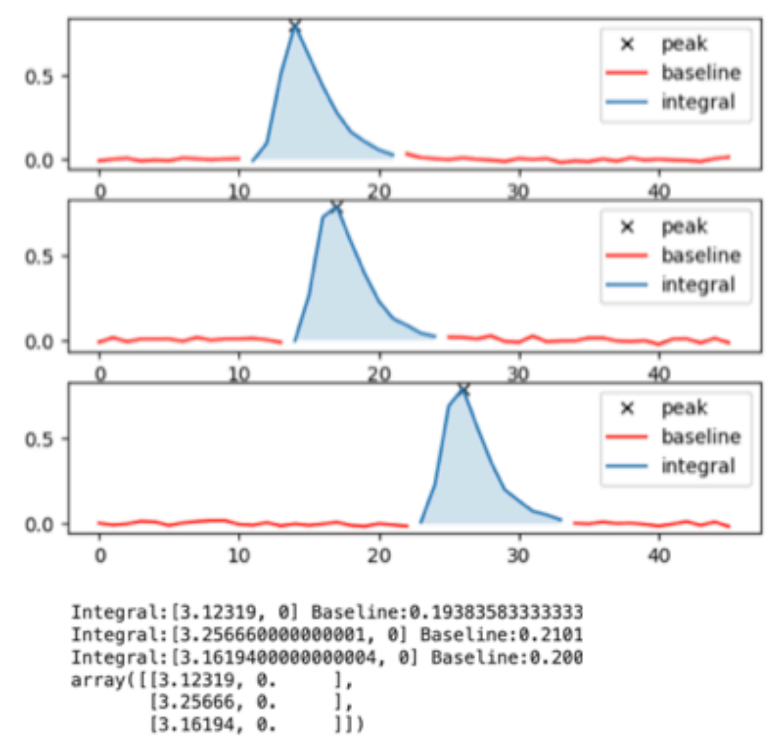

Integral and peak position is a possible data reduction (lossy) algorithm

Key components: data reduction

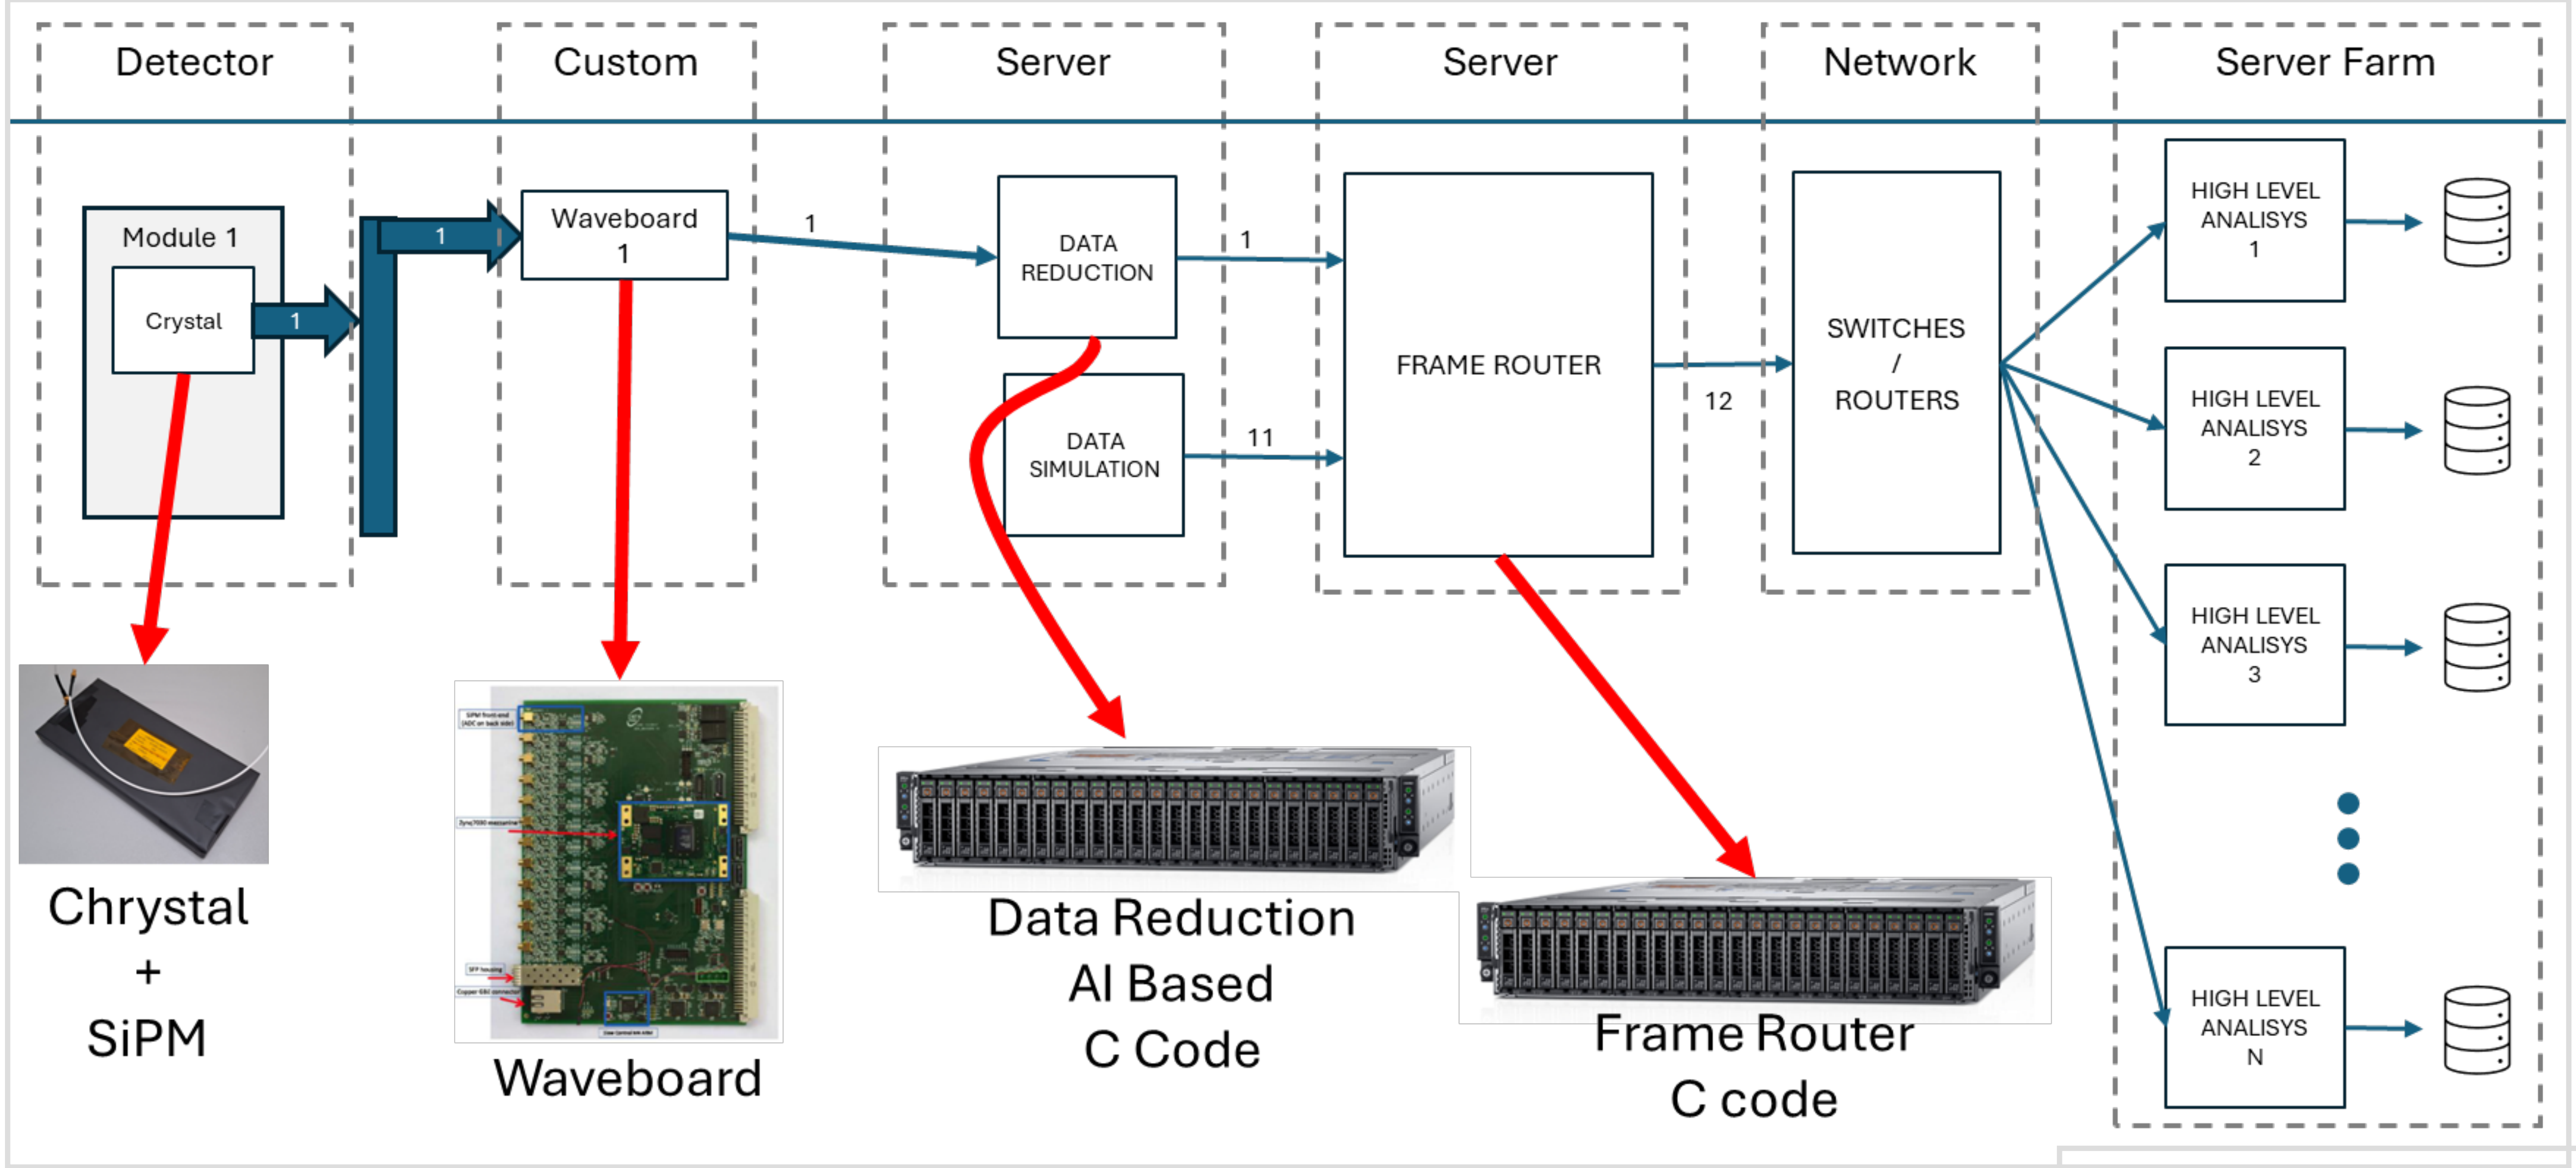

Test Setup in Genova

- Sample signals from JLAB-CLAS12 detector (similar to PbWO ECal crystal + PMT signals)
- fADC 250MHz, 16bits, discrite
- dedicated (heterogeneous) hardware for data reduction
  - GPU: Tesla V100
  - FPGA: Alveo V70
  - CPU: Dell C6400

Results

Best performance achieved with a CPU-based server

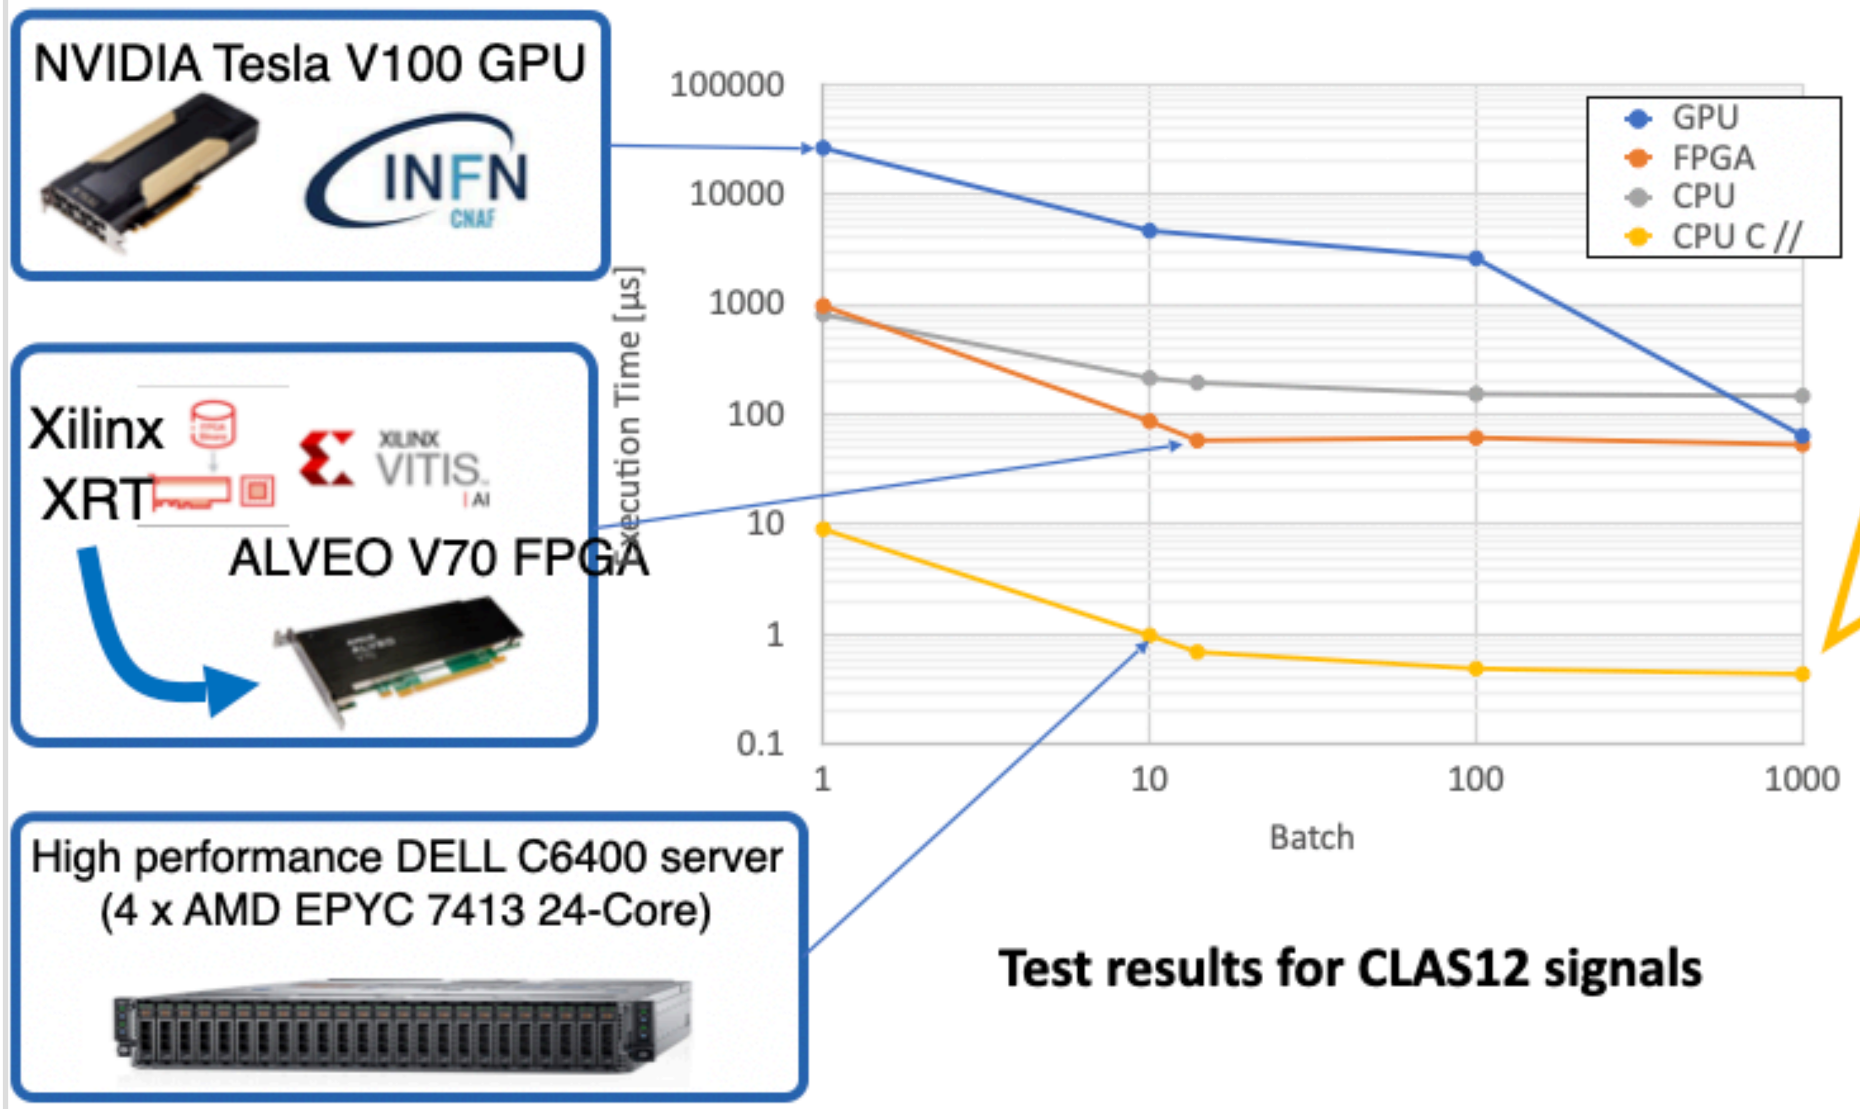

Test results for CLAS12 signals

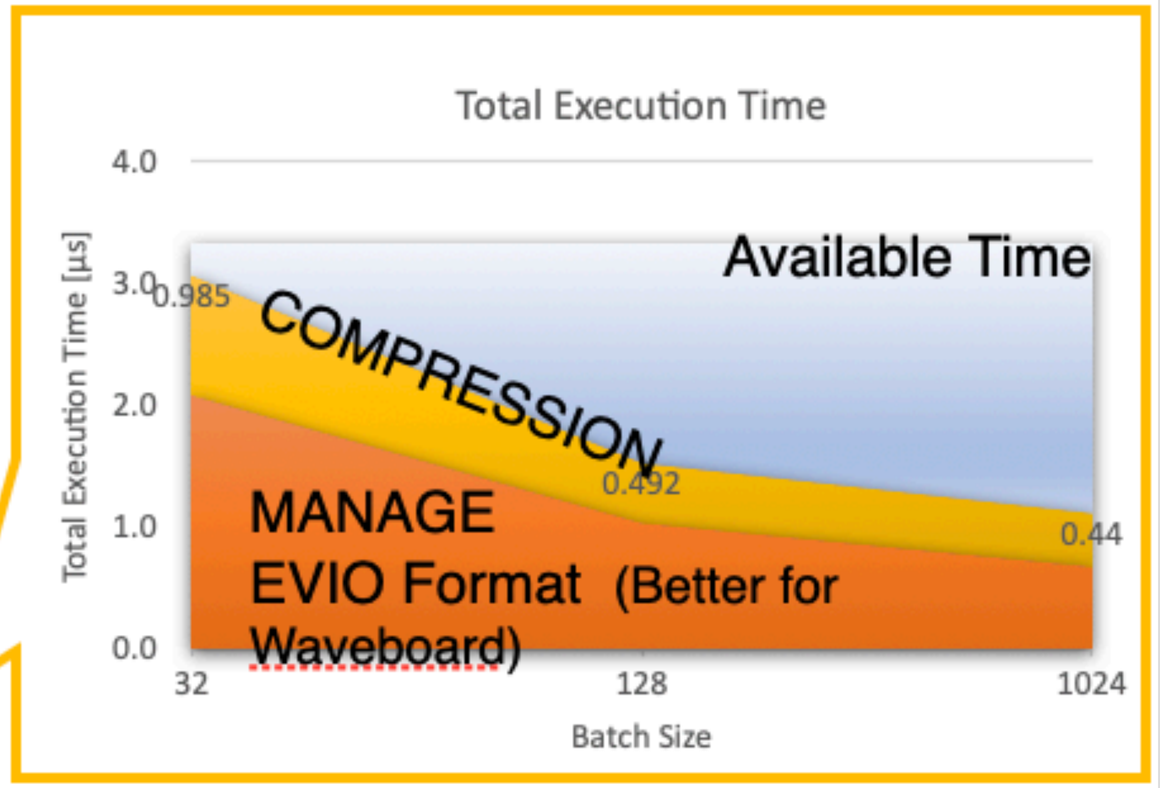

Next Step

Hardware level

HLS4ML

Dedicated connectivity

Alinx VD 100

Special case: dRICH

Using Graph Neural Networks for data reduction in the ePIC dRICH detector

Luca Pontisso
INFN Roma, APE Lab
for the ePIC Roma1/2 team

More in Alessandro's and Cristian's talks

Analysis of dRICH Output Bandwidth

The dRICH DAQ chain in ePIC ➔ **the throughput issue**

SiPMS ~ 320k
↓
RDO 1248
↓
DAM ~ 30
↓ **100 GbE x 30**
ePIC processing and storage system

dRICH DAQ parameters	
RDO boards	1248
ALCOR64 x RDO	4
dRICH channels (total)	319488
Number of DAM L1	27
Input link in DAM L1	47
Output links in DAM L1	1
Number of DAM L2	1
Input link to DAM L2	27
Link bandwidth [Gb/s] (assumes VTRX+)	10
Interaction tagger reduction factor	1
Interaction tagger latency [s]	2,00E-03
EIC parameters	
EIC Clock [MHz]	98,522
Orbit efficiency (takes into account gap)	0,92

Bandwidth analysis		Limit
Sensor rate per channel [kHz]	300,00	4.000,00
Rate post-shutter [kHz]	55,20	800,00
Throughput to serializer [Mb/s]	34,50	788,16
Throughput from ALCOR64 [Mb/s]	276,00	
Throughput from RDO [Gb/s]	1,08	10,00
Input at each DAM I [Gbps]	50,67	470,00
Buffering capacity at DAM I [MB]	12,97	
Output from every DAM	50,67	10,00
Total throughput	1.368,14	270,00

- Sensors DCR: 3-300 kHz (increasing with radiation damage ➔ with experiment lifetime).
- Full detector throughput (FE): 14-1400Gbps
- **A reduction is needed** to match 30 channels aggregated bandwidth (and safety margin)
- EIC beams bunch spacing: 10 ns ➔ bunch crossing rate of 100 MHz
- For the low interaction cross-section (DIS) ➔ one interaction every ~100 bunches ➔ interaction rate of ~1MHz.
- **A system tagging the (DIS) interacting bunches** could solve the issue reducing down to ~1/100 the data throughput

Two complementary approaches are possible:

1. **Develop a dedicated sub-detector tagging relevant interactions.**
2. **This proposal.**

Streaming Data Structure

Time Frames are unit of building data within DAQ system

- At most 2^{16} bunch crossings (660 us)
- O(5MB)
- Time frames will contain data from all detectors
- Time frames are not tied to revolutions of the beam
  - ➢ Do not want persistent effects due to time frame borders being correlated with spin states

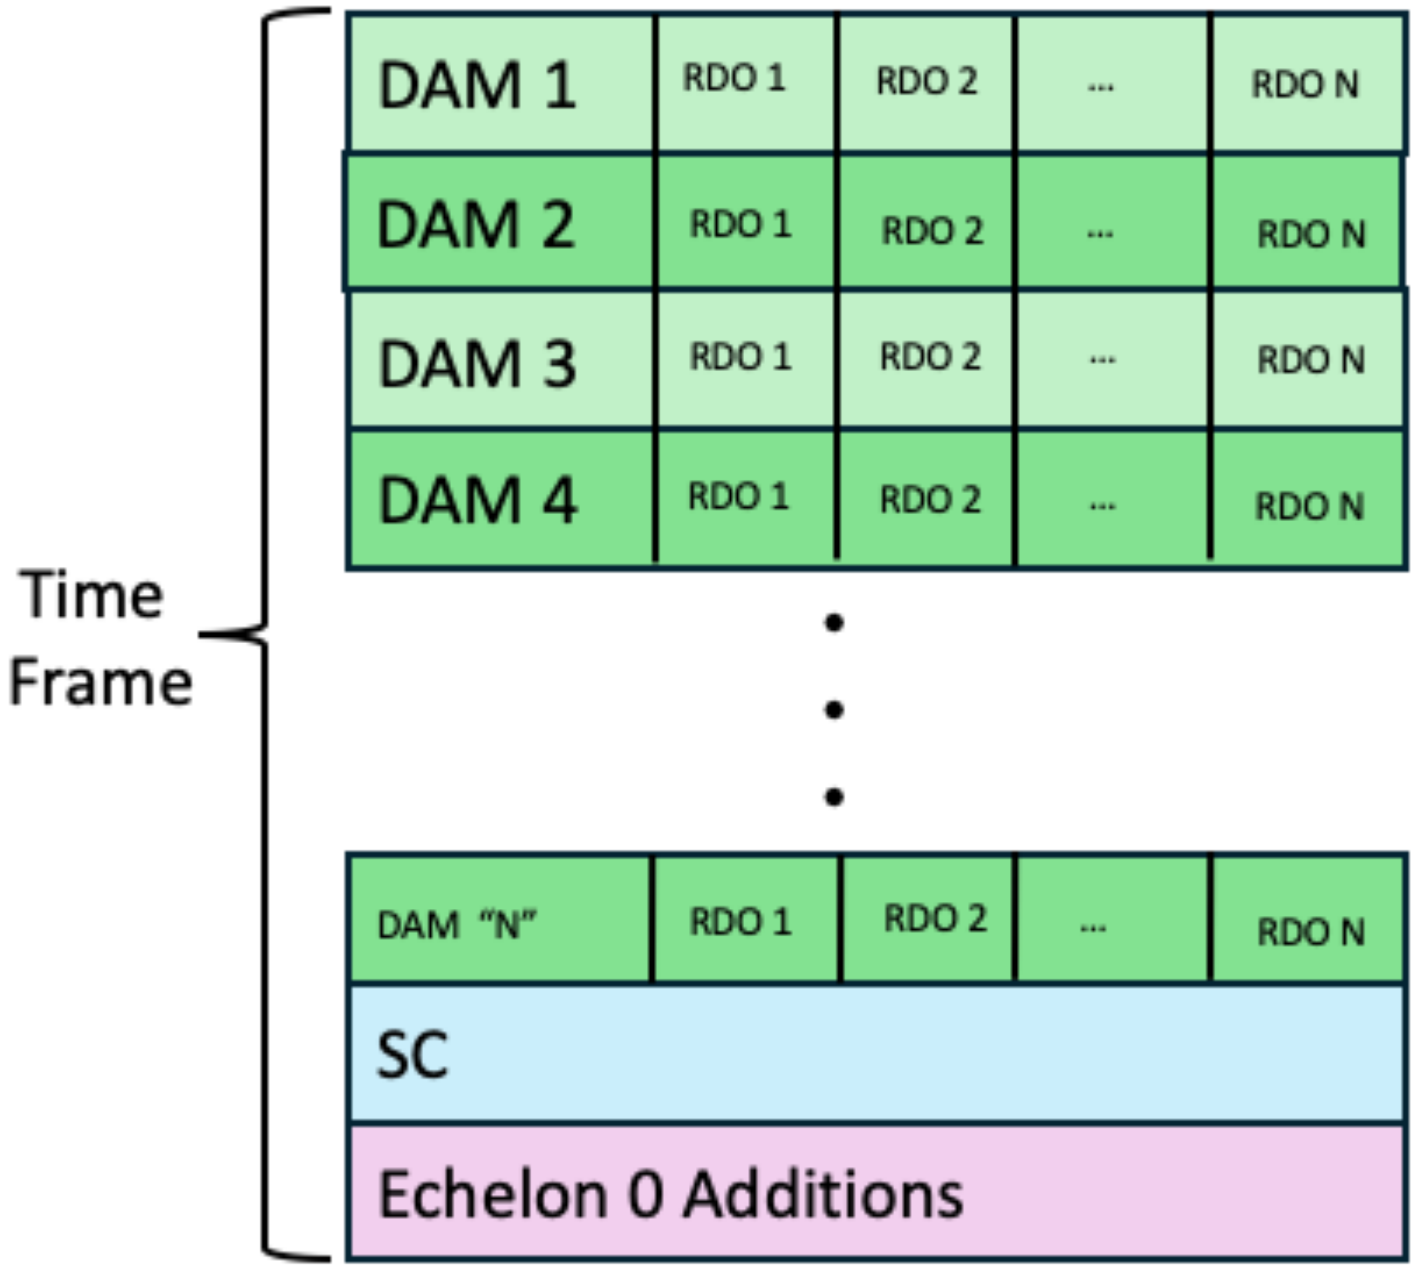

Super Time Frames are unit of management of data within S&C

- And extended number of Time Frames O(1000)
- Time Ordered
- Dense (non-overlapping)

ePIC State Machine

- Will have continuously running subsystems (24/7) [scaler / luminosity]
- Will have subsystems turned on/off and set to different modes at different times [beam sensitive detectors / detectors with complex configuration / calibration]
- Will have subsystem failures (down to RDO level at least), and also intend to have capability for auto-recovery of electronics where possible.
- Will have low level failures / timeouts / overruns. The strategy is to mark effected channels and write out unaffected channels

Key components: data management

Data Transfer Between Echelon 0 and Echelon 1

Buffers in DAQ and **Echelon 1** sites ensure latency tolerance to avoid deadtime, smooth streaming operation and robustness against data flow interruptions.

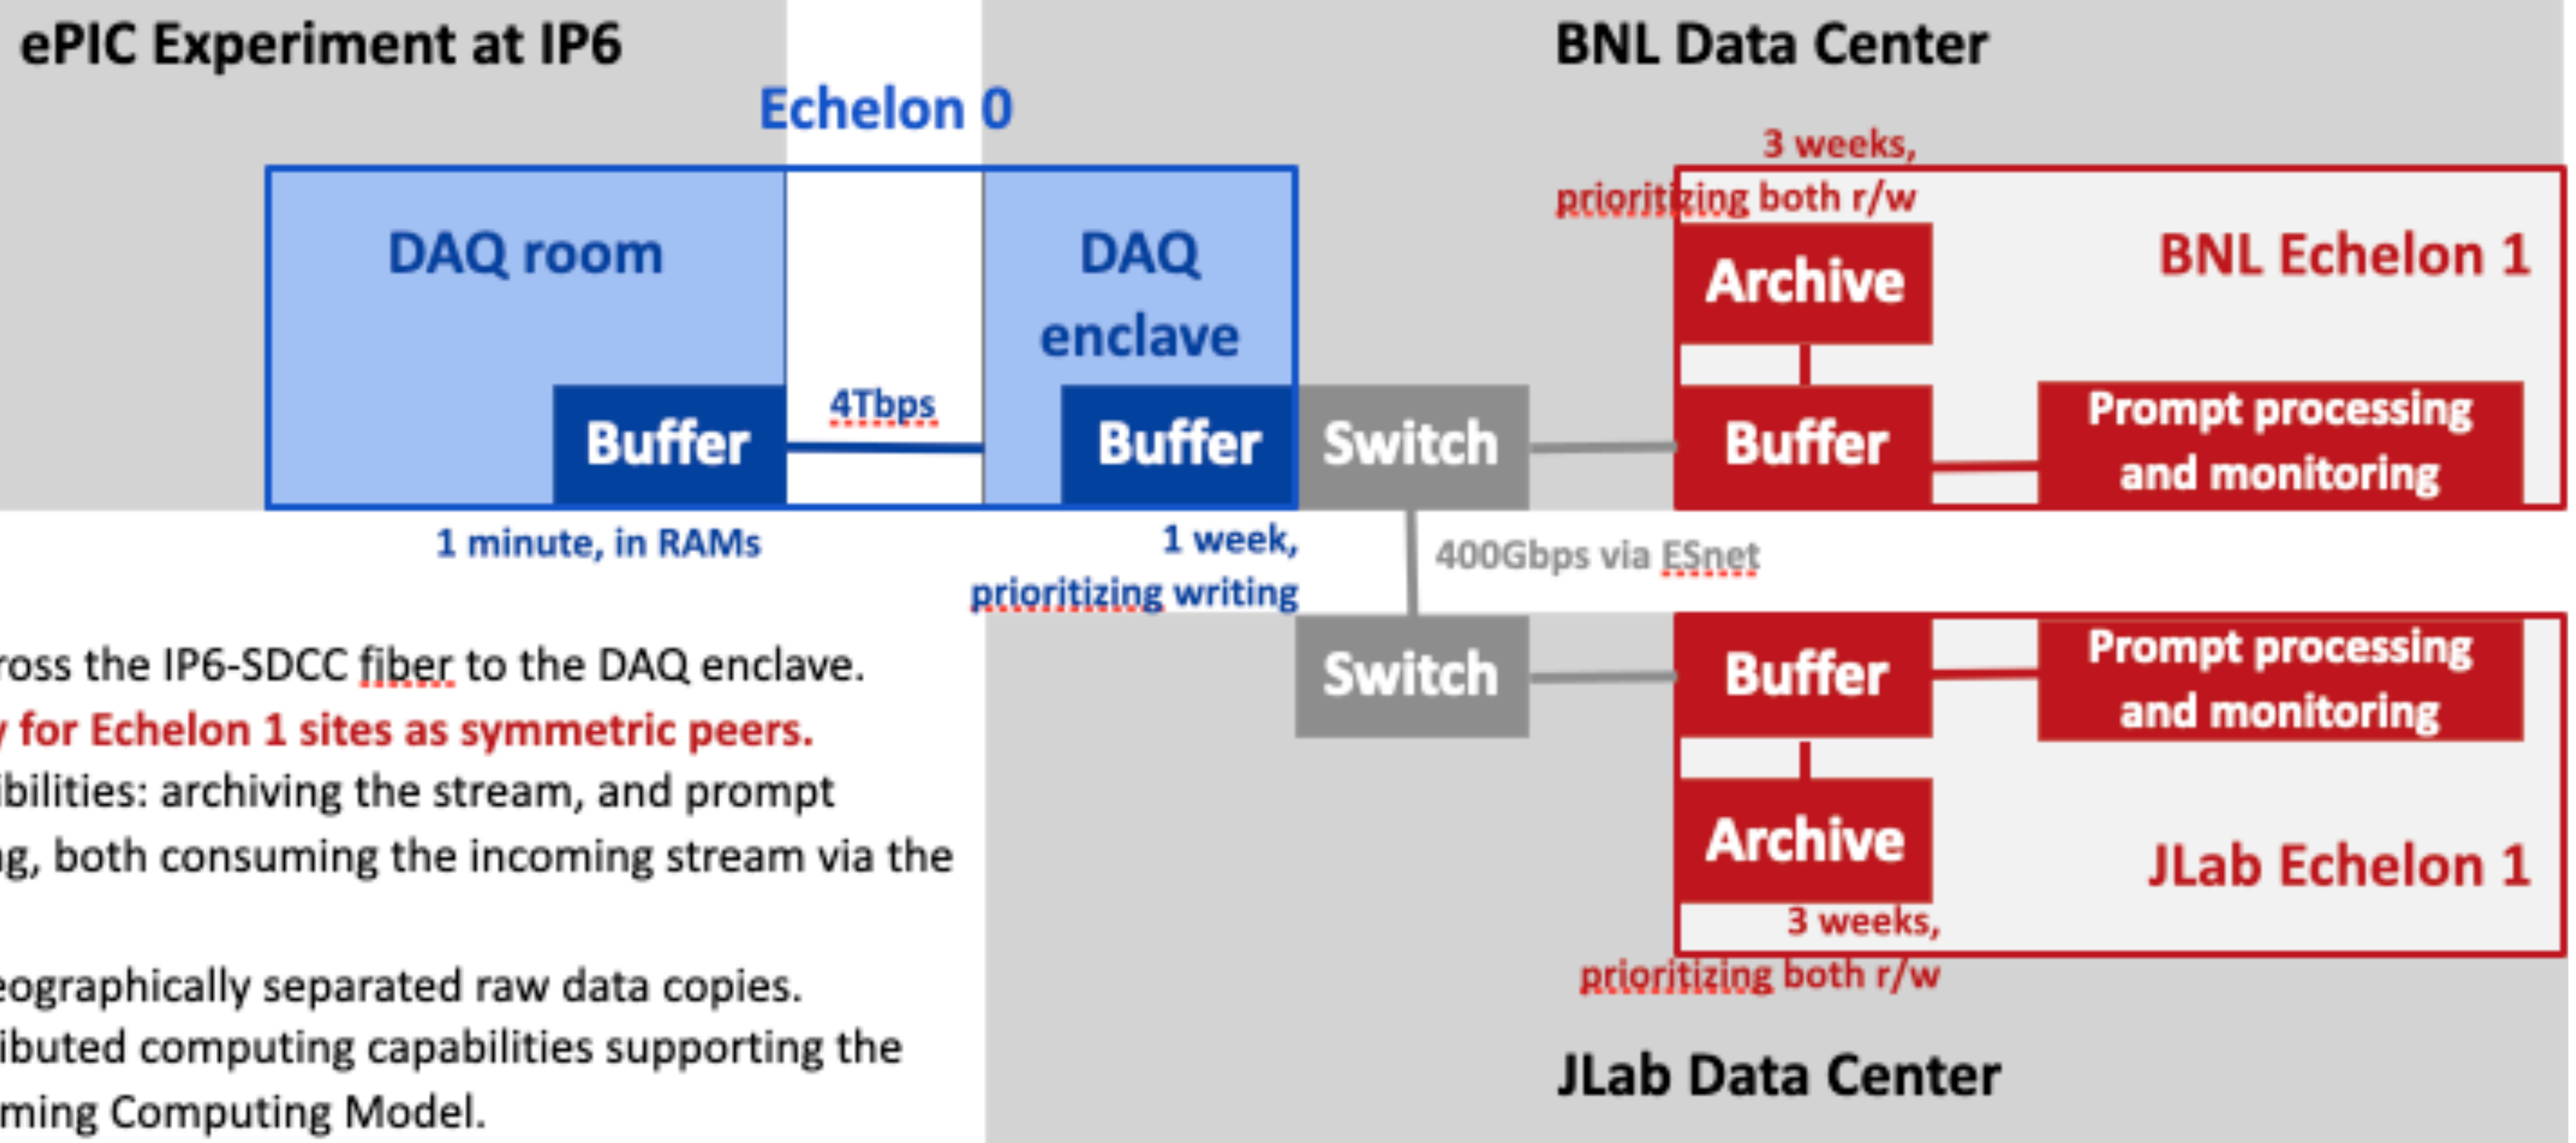

- Echelon 0 extends across the IP6-SDCC fiber to the DAQ enclave.
- **Deliver the capability for Echelon 1 sites as symmetric peers.**
- Foremost E1 responsibilities: archiving the stream, and prompt processing/monitoring, both consuming the incoming stream via the buffer:
  - Delivers two geographically separated raw data copies.
  - Uses ePIC distributed computing capabilities supporting the E0/1/2/3 Streaming Computing Model.
  - Will be up to the ePIC collaboration together with sites to determine the E1 roles in detail.

Key components: SRO orchestrator

SRO DAQ components need to be ‘orchestrated’

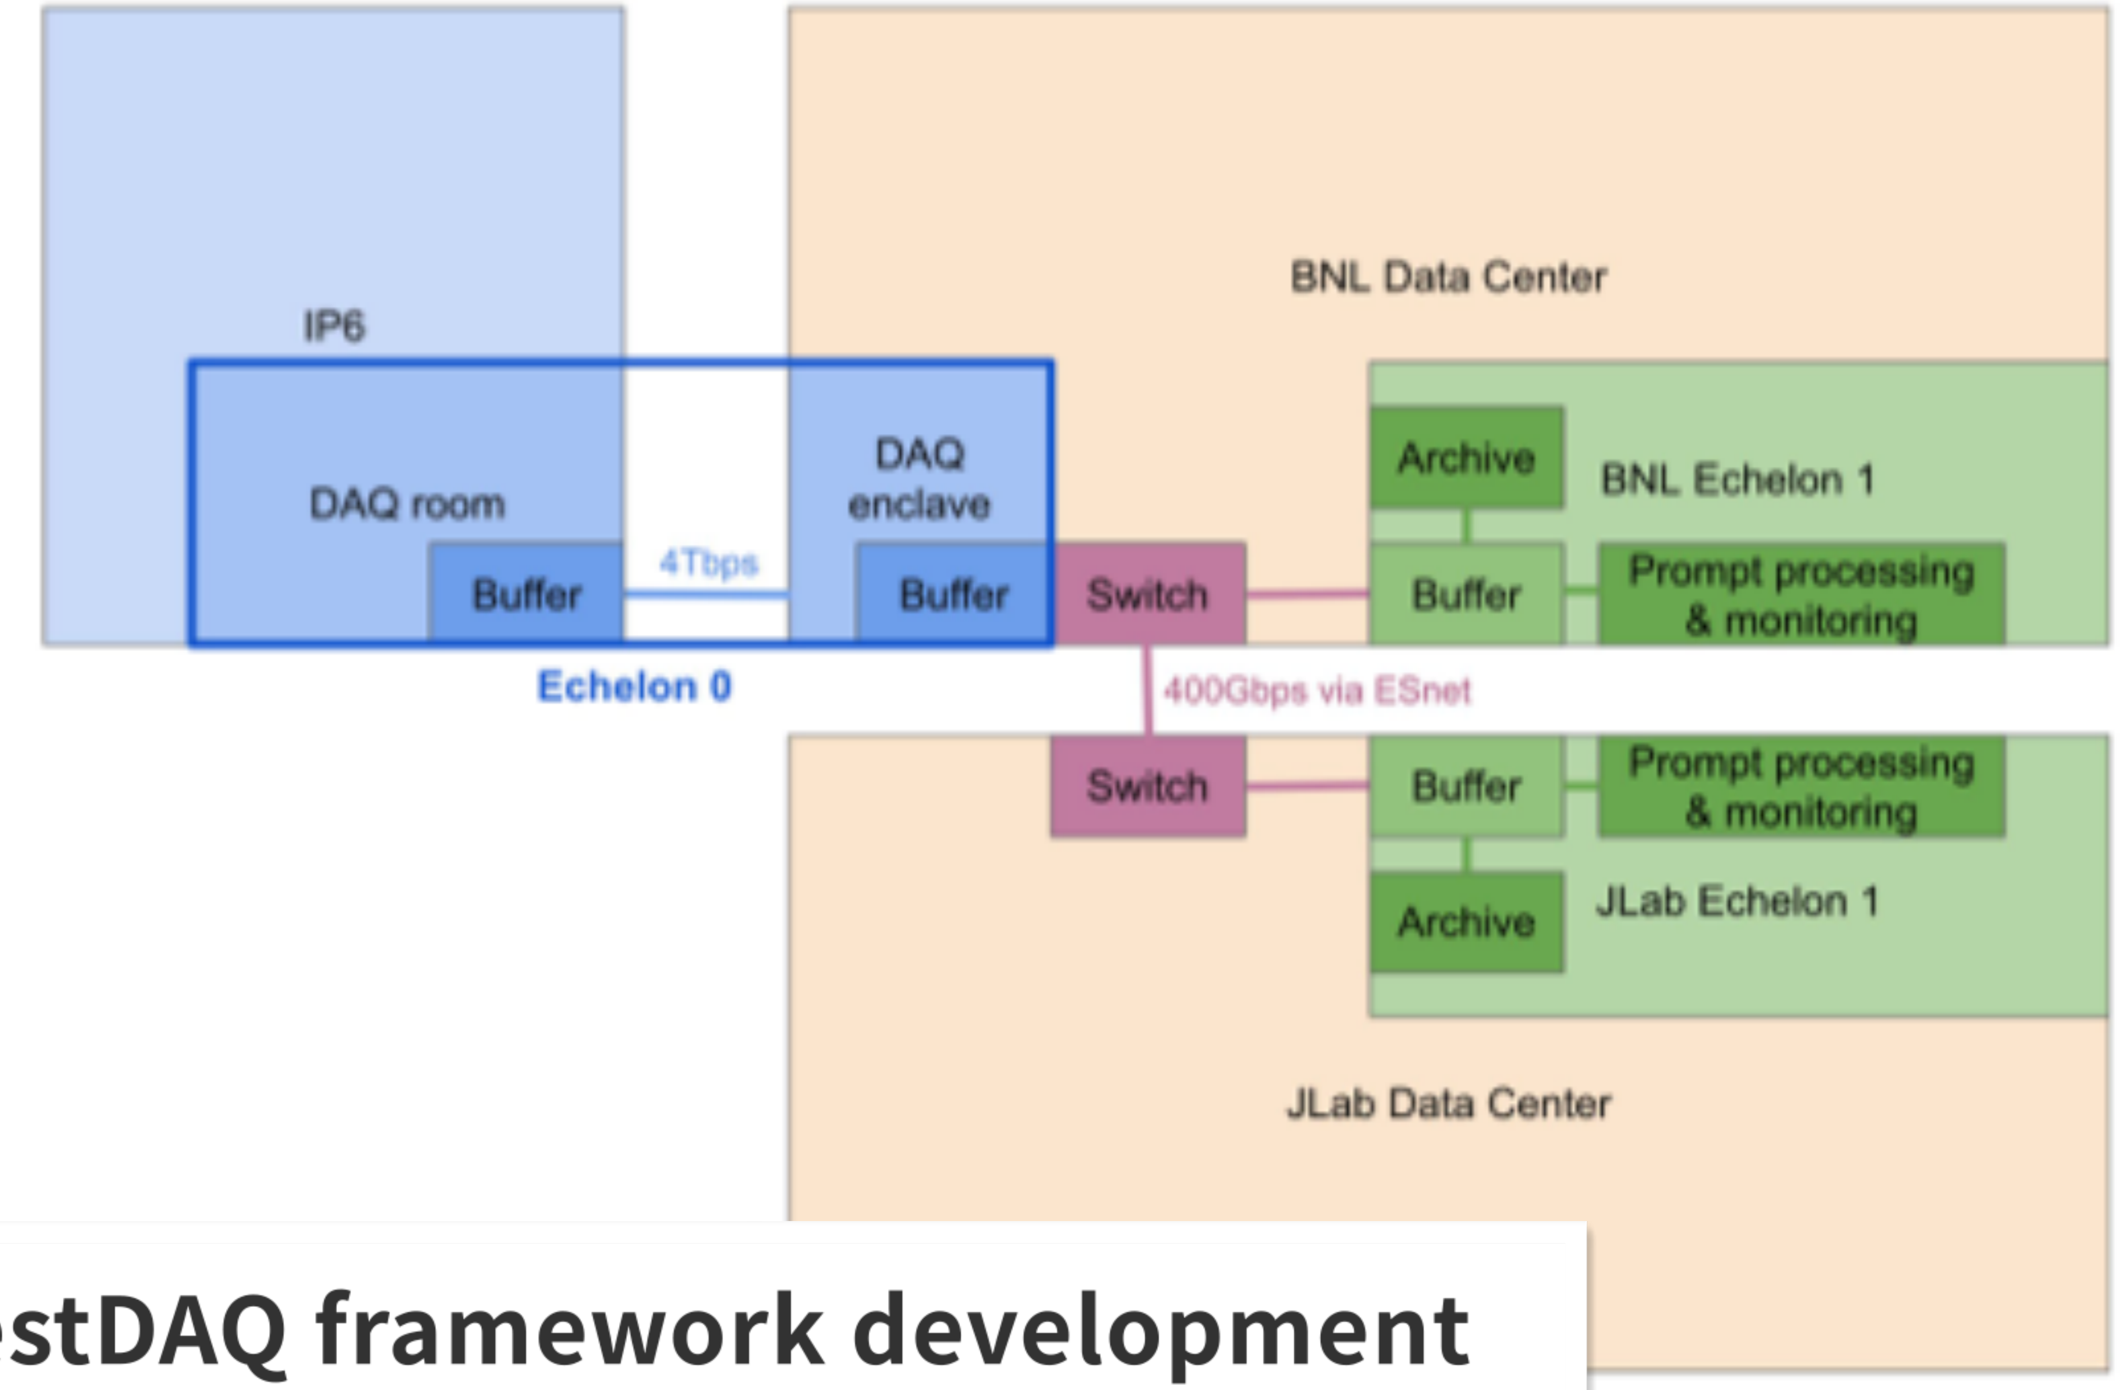

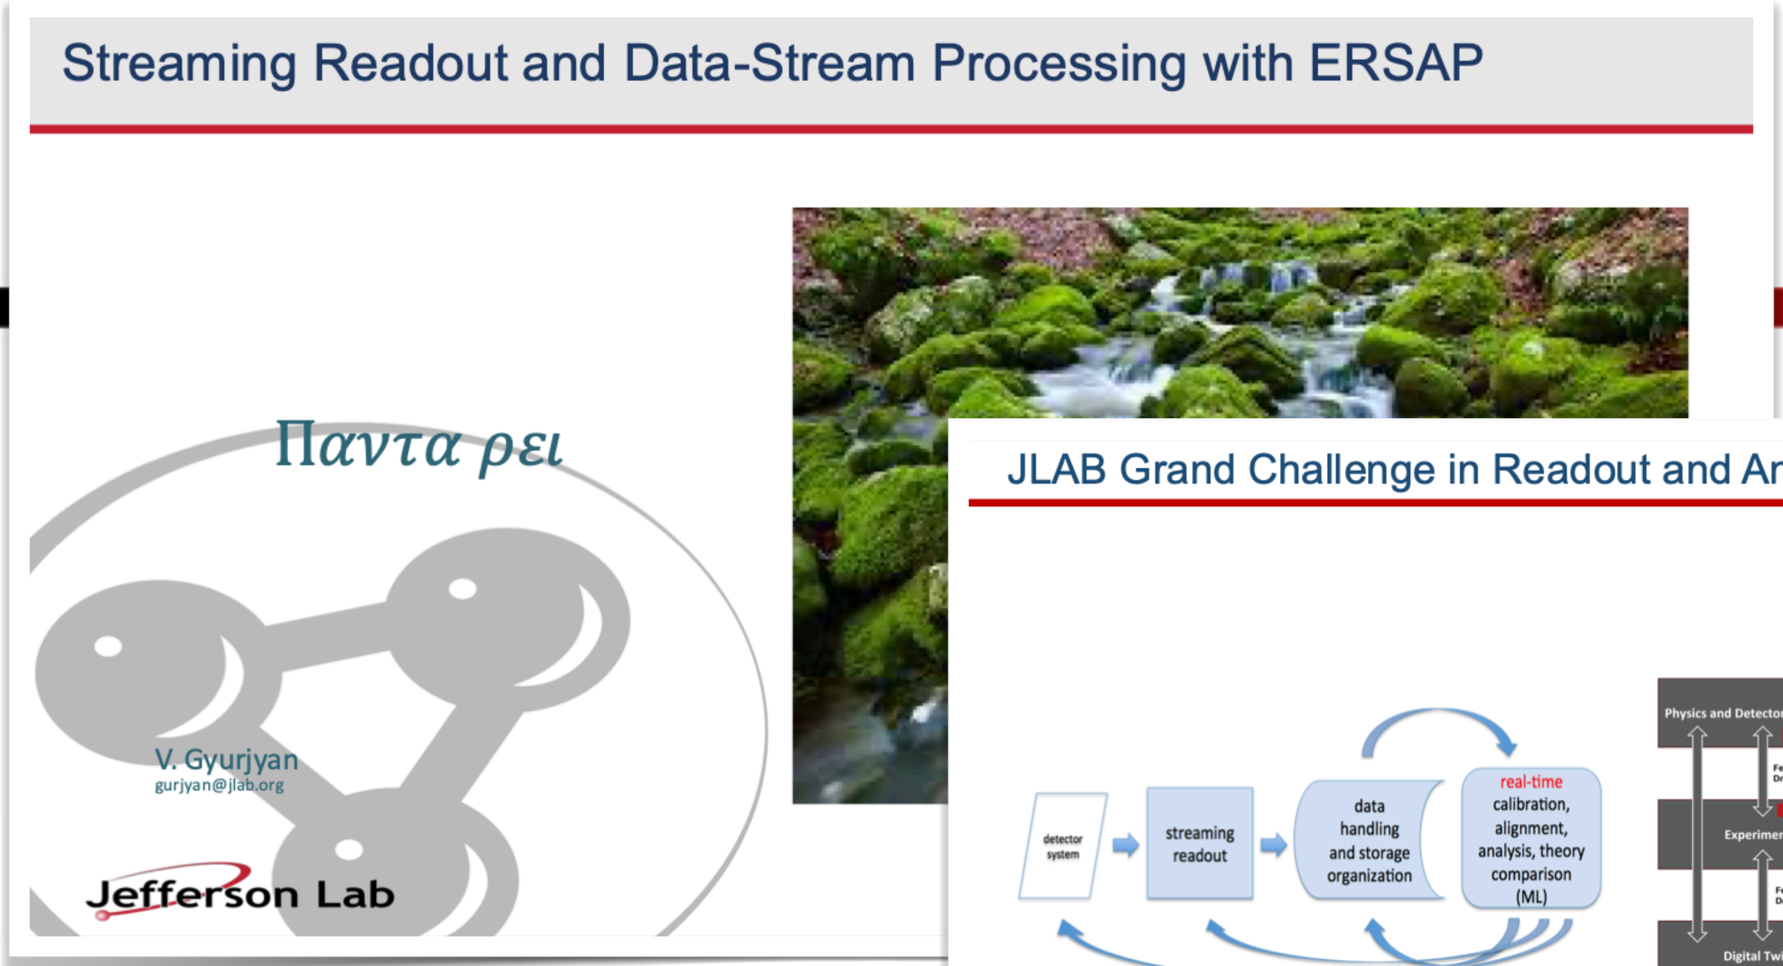

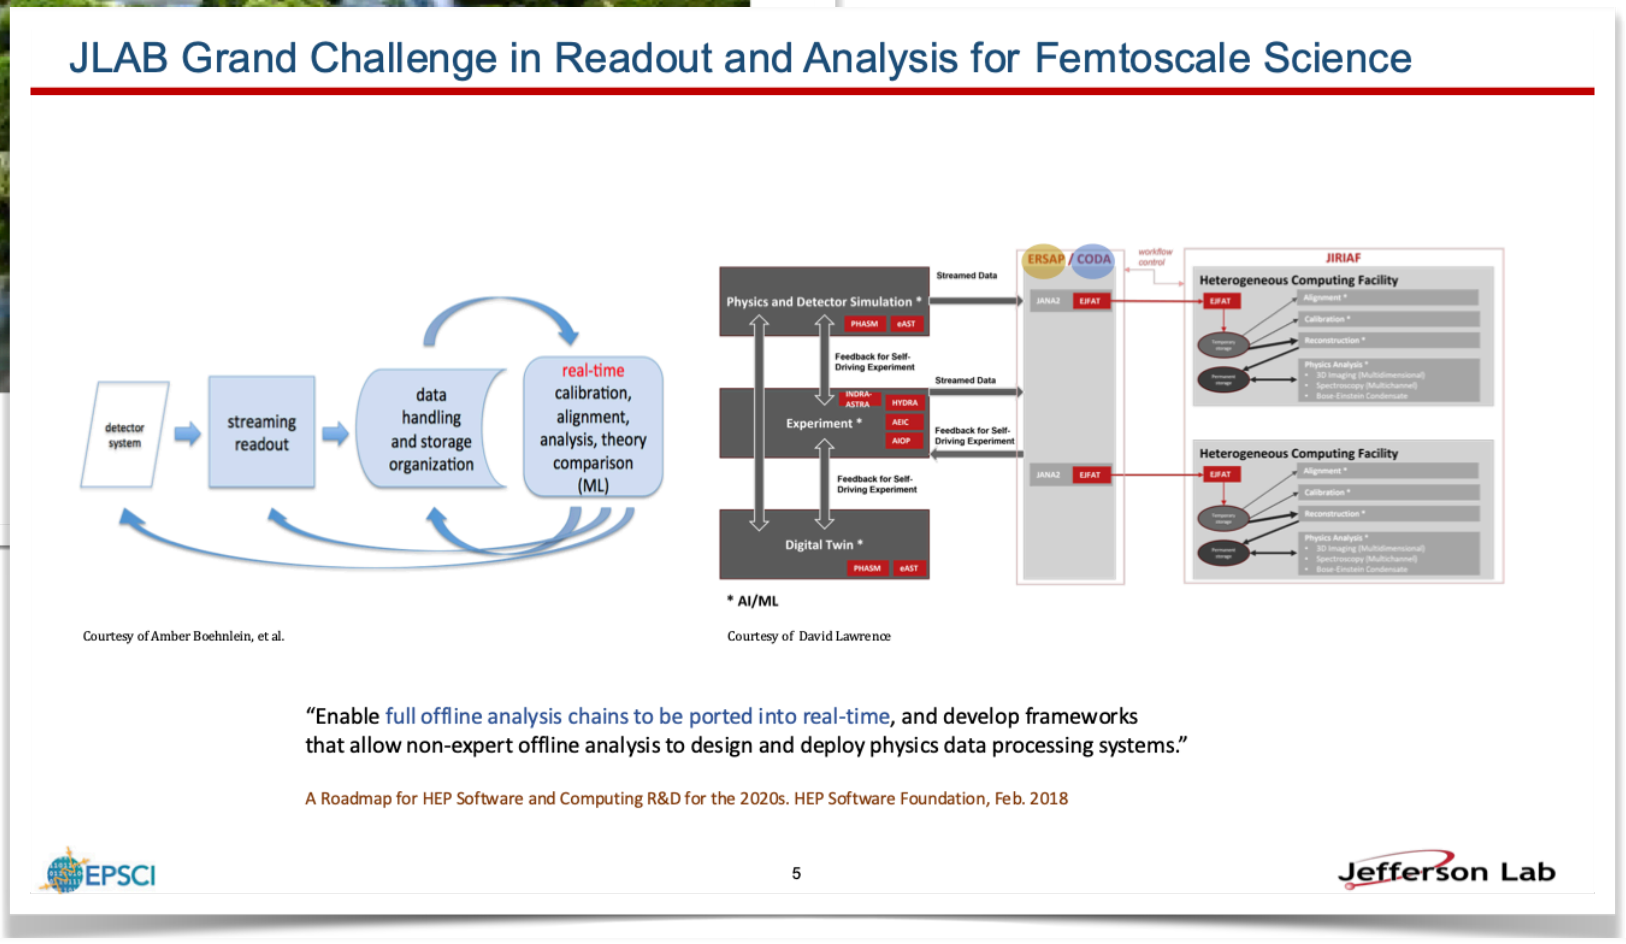

JLab-ERSAP
- micro-service based
- distributed
- heterogeneous hw ready

RCDAQ – Some of the High Points

sPHENIX

- Each interaction with RCDAQ is a network connection that transmits the action to be taken and a response coming back
- The most-often used implementation is an atomic shell command. There is no “starting an application and issuing internal commands” (think of your interaction with, say, root)
  - That makes everything in RCDAQ inherently scriptable in standard bash or your other favorite shell (or python)
  - We start a sPHENIX DAQ run by pressing one button that fires off a script that takes care of it all
  - In test beams and tests in your lab you can script entire measurement campaigns and run them “on autopilot” – think bias voltage scans, position scans etc
- RCDAQ out of the box doesn’t know about any particular hardware. All knowledge how to read out something, say, the FELIX board, comes by way of a **plugin** that teaches RCDAQ how to do that.
- All RCDAQ control interfaces are network-transparent
- There is no practical limit for concurrent control connections for RCDAQ

7

NestDAQ framework development

T.N. Takahashi
Research Center for Nuclear Physics, Osaka University

Outline

- Motivation
- Implementation
- Future prospect
- Summary

SPADI
Alliance

BNL-RCDAQ
- hybrid (triggered/SRO) DAQ
- used in sPHENIX

SPADI alliance
- Japan-based but aim in g to expand in the US (e.g. FRIB)

Key components: alignment/calibrations

- **ePICS SRO DAQ aims for a rapid turnaround from data to full calibrated/reconstructed data**
- **Data reconstruction time scale driven by calibrations (2-3 weeks max)**

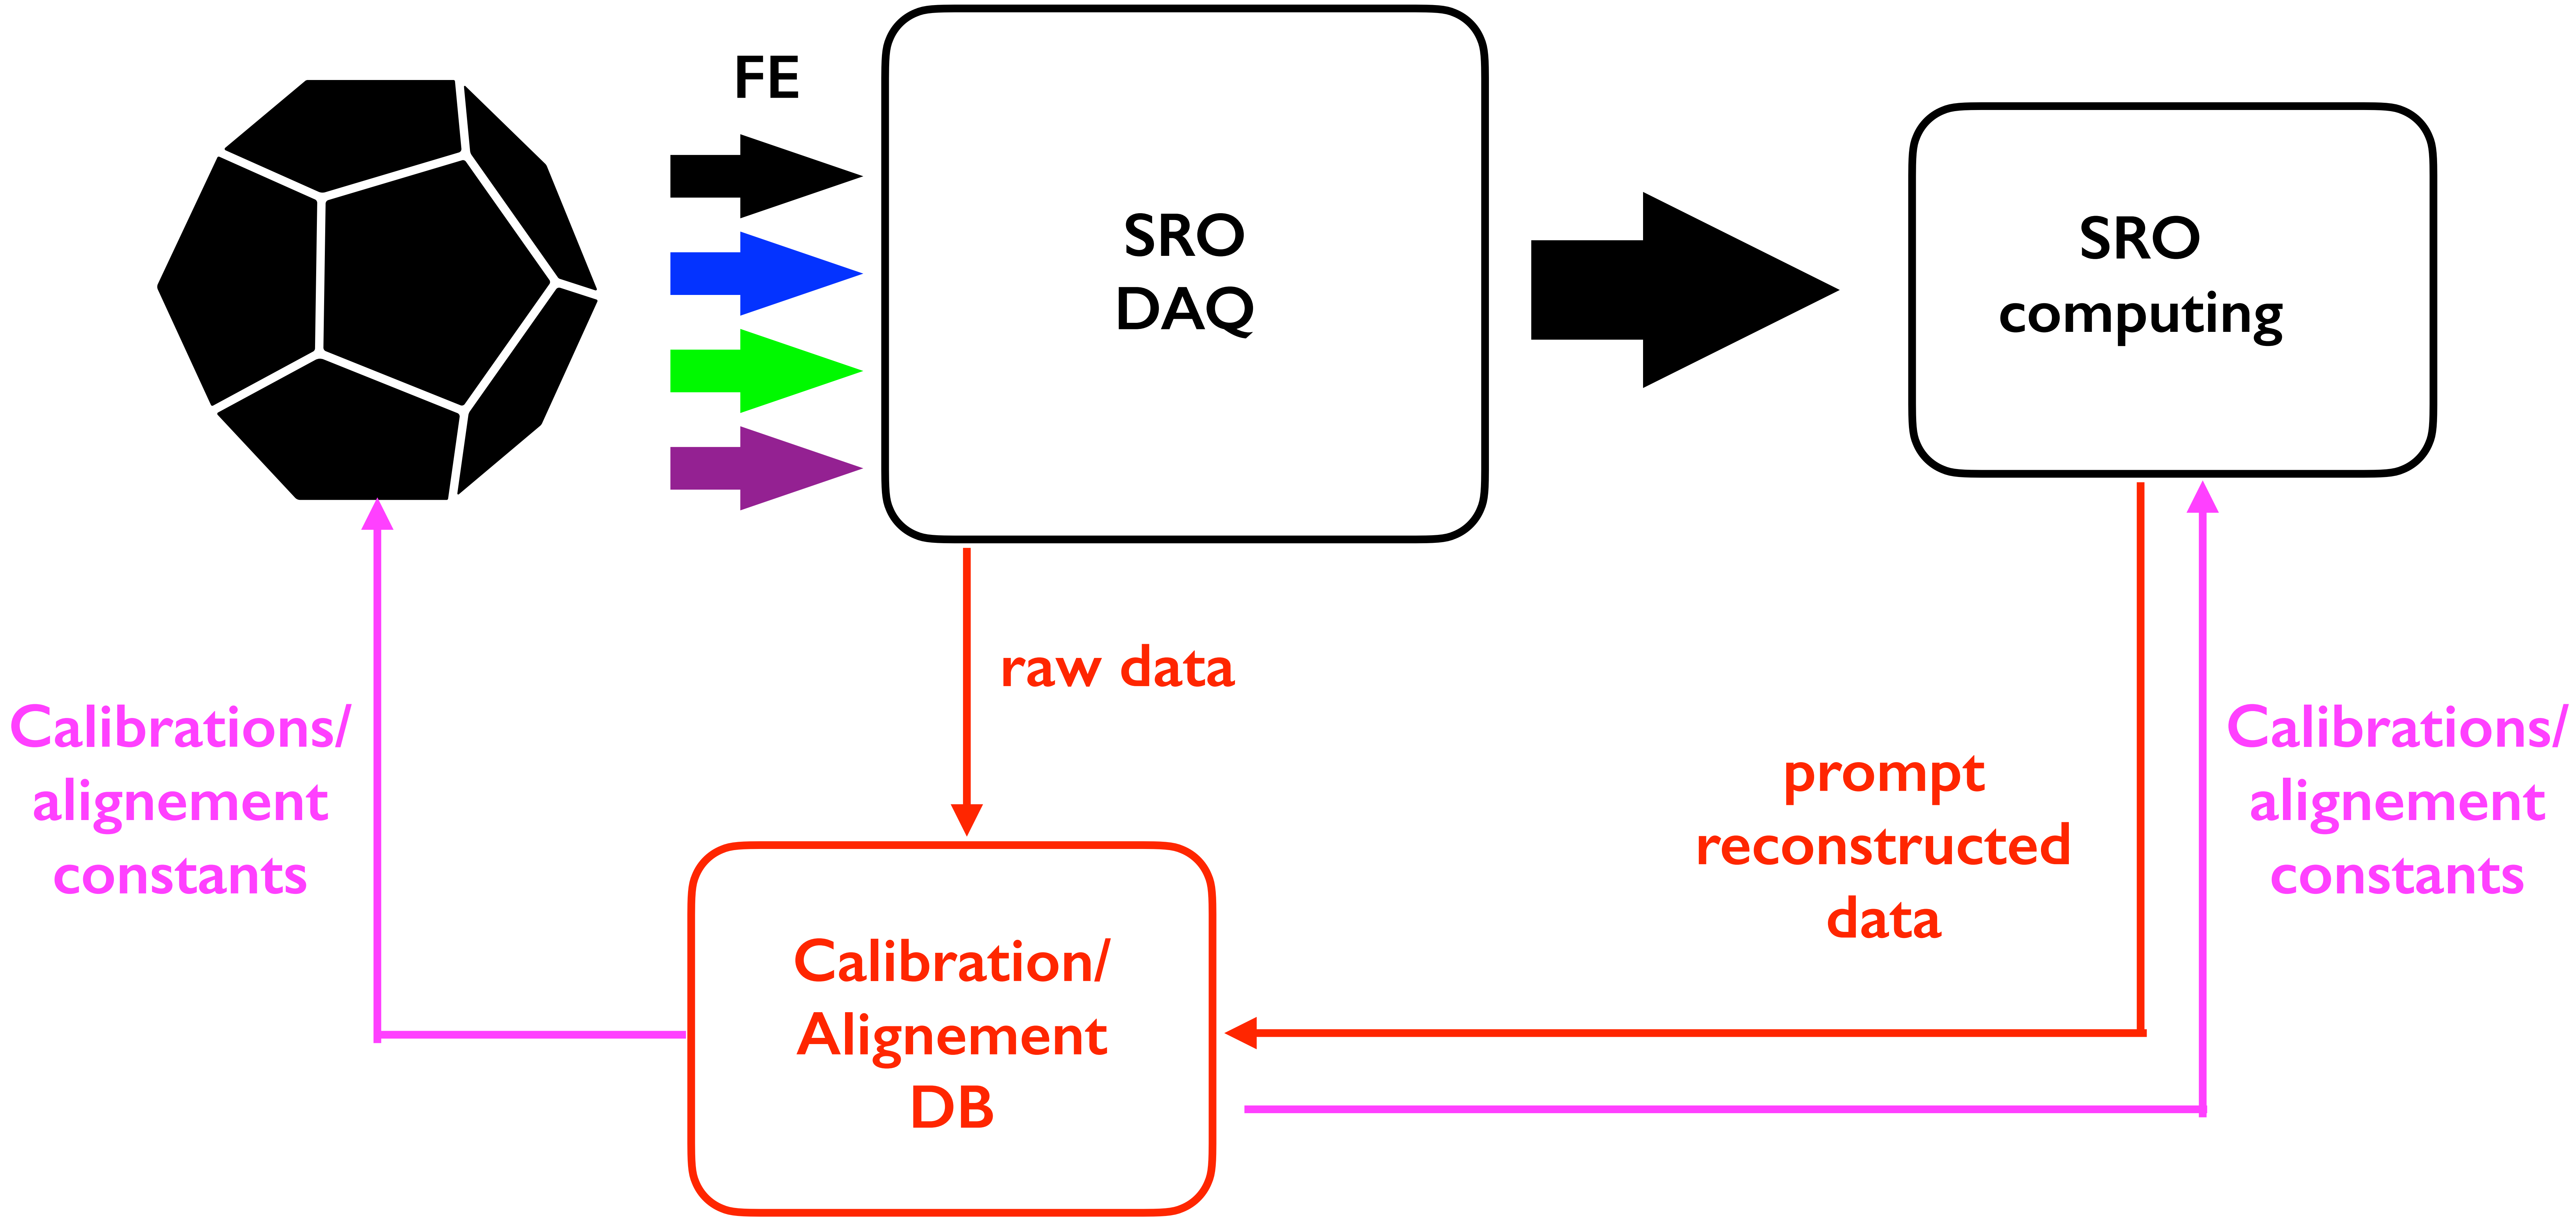

- Specific to each sub-detector system
- Multiple iterations
- Different computing levels (Echelon's)
- Automatic/autonomous (AI?)
- FAST!!!!

Defining calibrations and alignment procedures is essential and urgent!

Everybody agrees but ….

Rosi

Key components: alignment/calibrations

Questions to be answered

- Needs to define calibration requirements from each sub-detector
  - how much data is needed?
  - when often?
  - how/where to apply corrections to data?
- Correction should be autonomous (AI/ML algorithms as a 2nd iteration or from the start?)
- At which level (Echelon 0 and/or Echelon 1 or 2)
- Calibration and simulation framework
- **Implementation**
  - iterative procedure in (semi) real-time (some detectors may need info from others)
  - Are calibration parameters biasing the data set we will write on disk?
  - Are calibration procedures background-aware and how to reliably estimate that bg?
  - Identify required Infrastructures (e.g monitoring tools)

    ASIC level: e.g. 0 suppression

    DAQ level: eg. clustering

    SRO level: final physics extraction: how it propagates back to the FE?
  - Identify dependencies from other subsystems
  - Identify calib procedures requiring/not-requiring dedicated runs
  - Is an RND-trigger (fully unbiased) data stream in parallel to the prod runs stream?
  - Alignment/tracking may require special procedures need to be defined upfront

Special task force in the SRO WG (M.Battaglieri, T.Gunji) for alignment/calibrations

Is it NOT premature, even if,

- Detectors are still being designed
- Procedures require feedback from real implementation
- WG assumption
  - A lot of work can/should be done in advance to understand details of alignment and calibration procedures
- Strategy
  - A (living/updated) survey of different subsystem procedures
  - Engage sub-detector leaders to develop a shared workflow

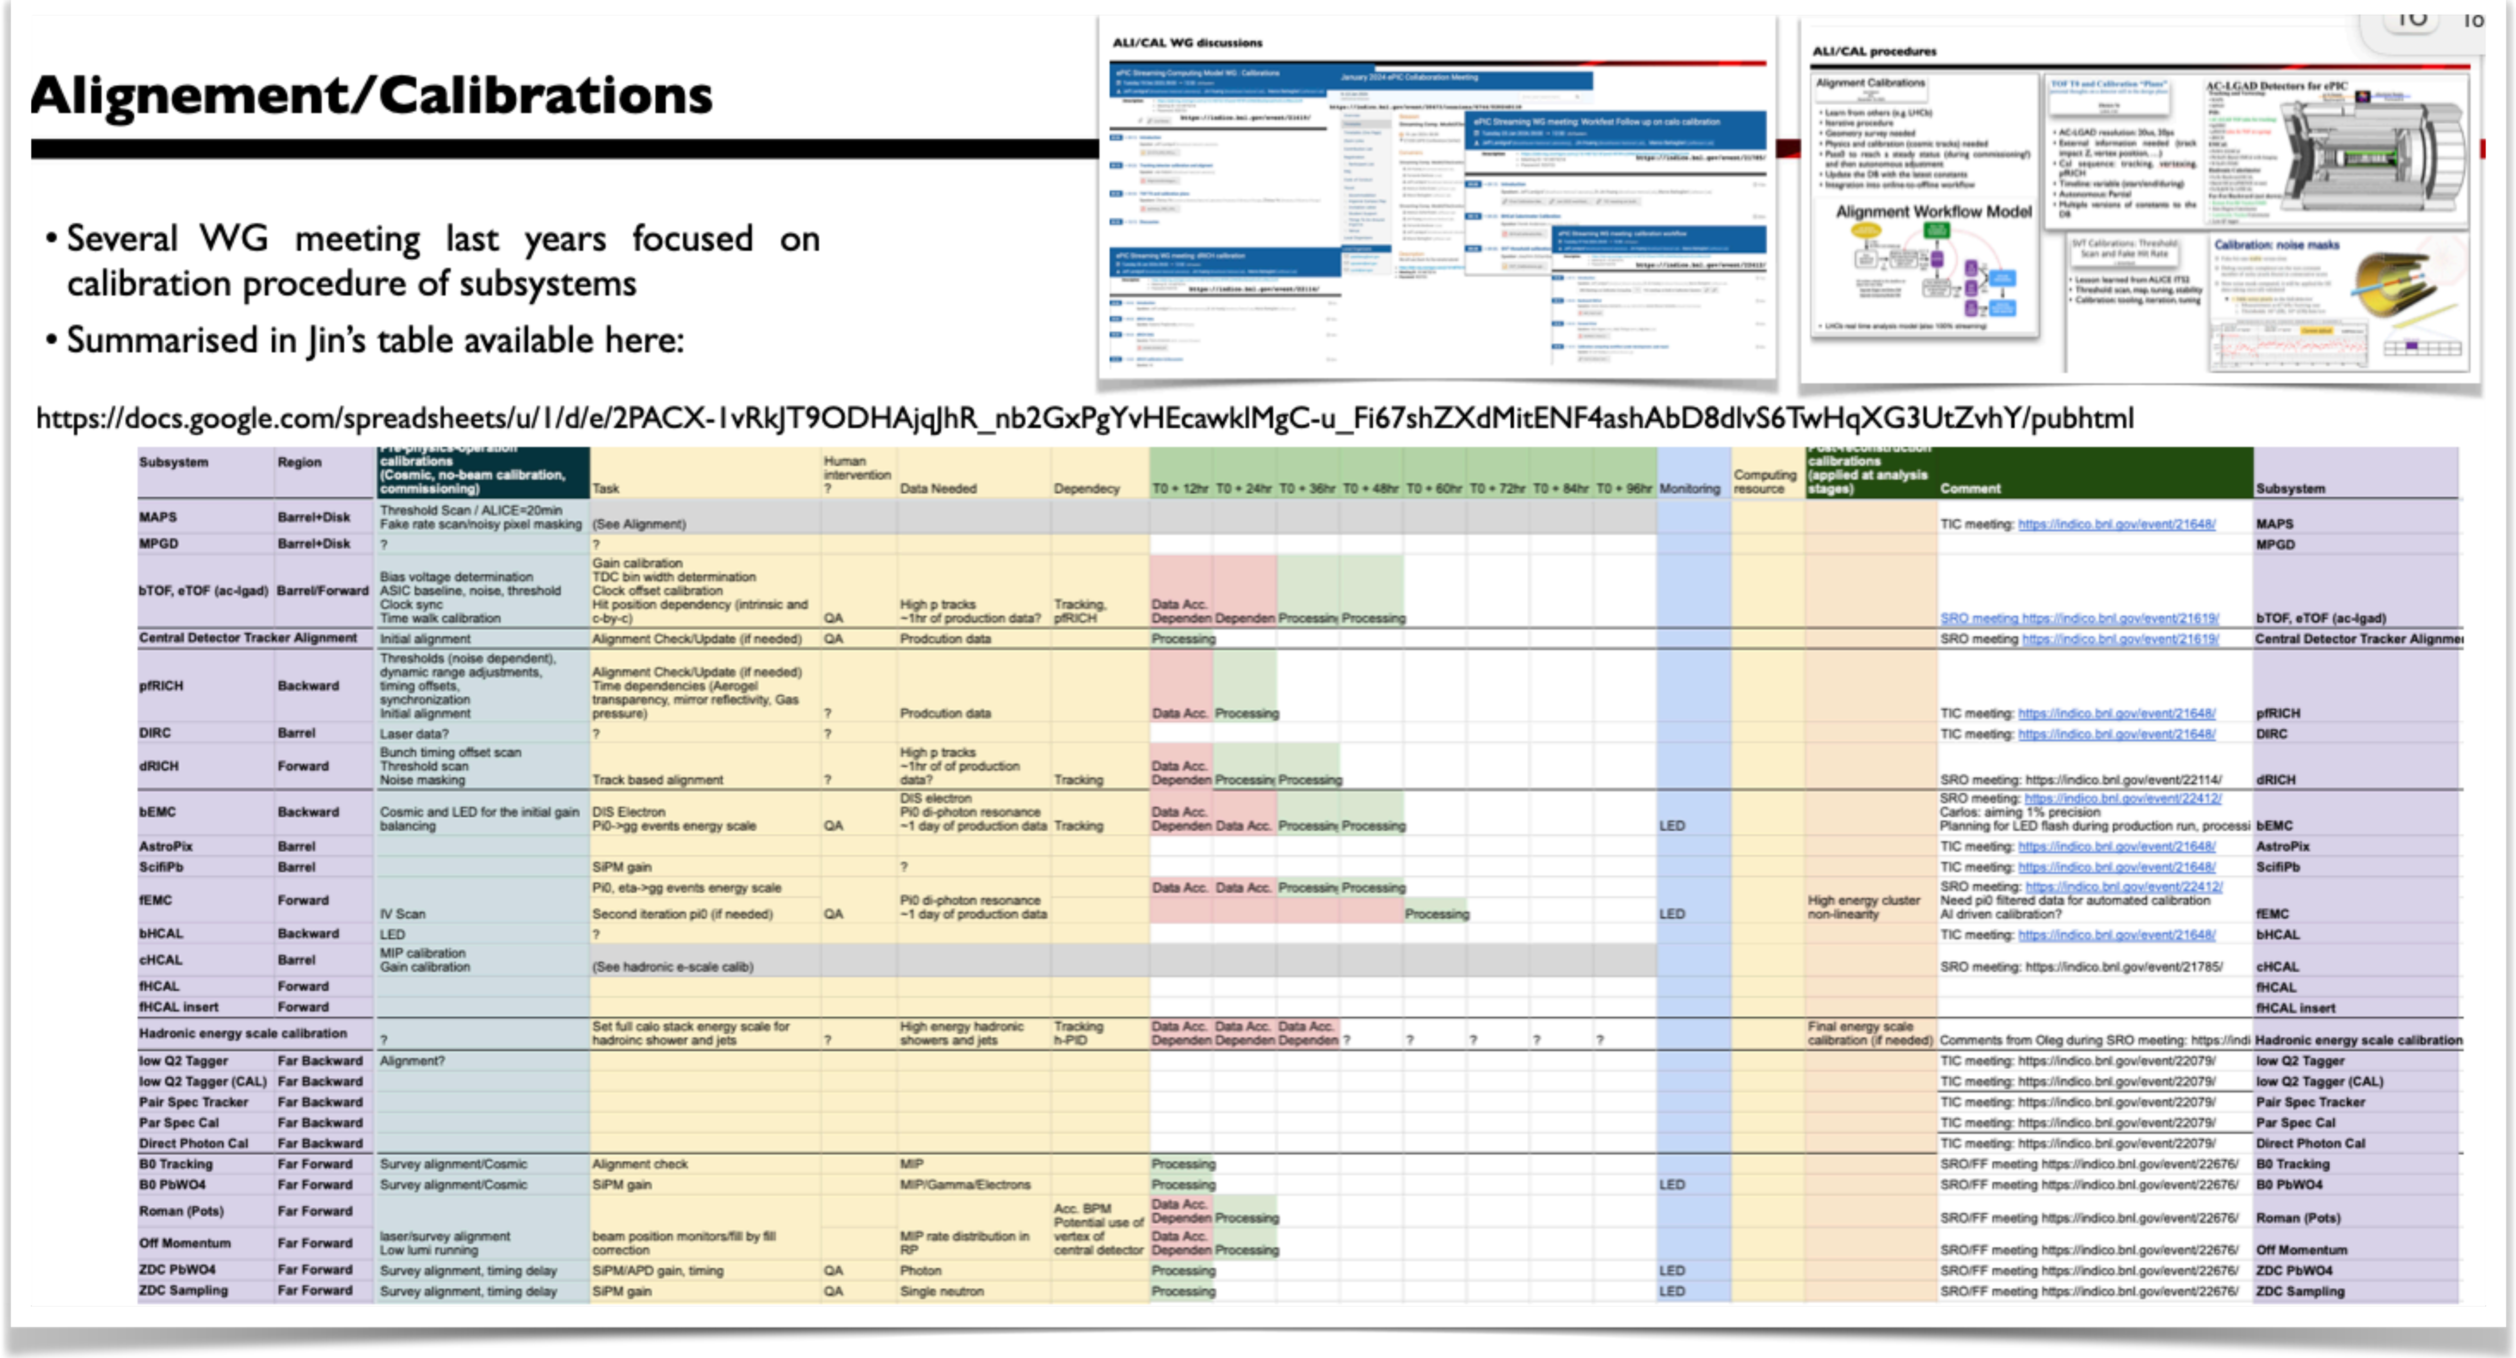

ePIC SRO Milestones

Scope

Components and functionalities of the system to define and implement in a first prototype:

- stream definition, structure and associated metadata
  - bulk data stream
  - monitoring stream
- E0 - E1 data flows
  - how bulk data moves from E0 to E1
  - how fast monitoring data moves from E0 to E1
- E1 bulk data orchestration
  - How and where bulk data lands at E1s
  - Triggering actions on the data (archiving, prompt processing) as it arrives
- E1 fast monitoring infrastructure
  - How fast monitoring data moves from E0 to E1
  - How and where it is received and processed by monitoring agents/workers
- E1 data processing orchestration
  - Workflows for processing physics data at the E1s
- E1 calibration orchestration
  - Workflows, from simple to complex, for performing prompt calibration/ alignment at E1s
  - Tools to define and automatically execute complex workflows with dependencies down the processing chain
- Detector/data state machine
  - Infrastructure to manage and interface to detector/data state in an E1 resident service
  - cataloging of detector/data states, clients served and their needs
- E2 extension
  - An extension of the testbed will prototype the inclusion of Echelon 2 (E2) computing facilities. Volunteer prospective E2 sites are welcome and invited to join the testbed effort and enable us to begin this extension.

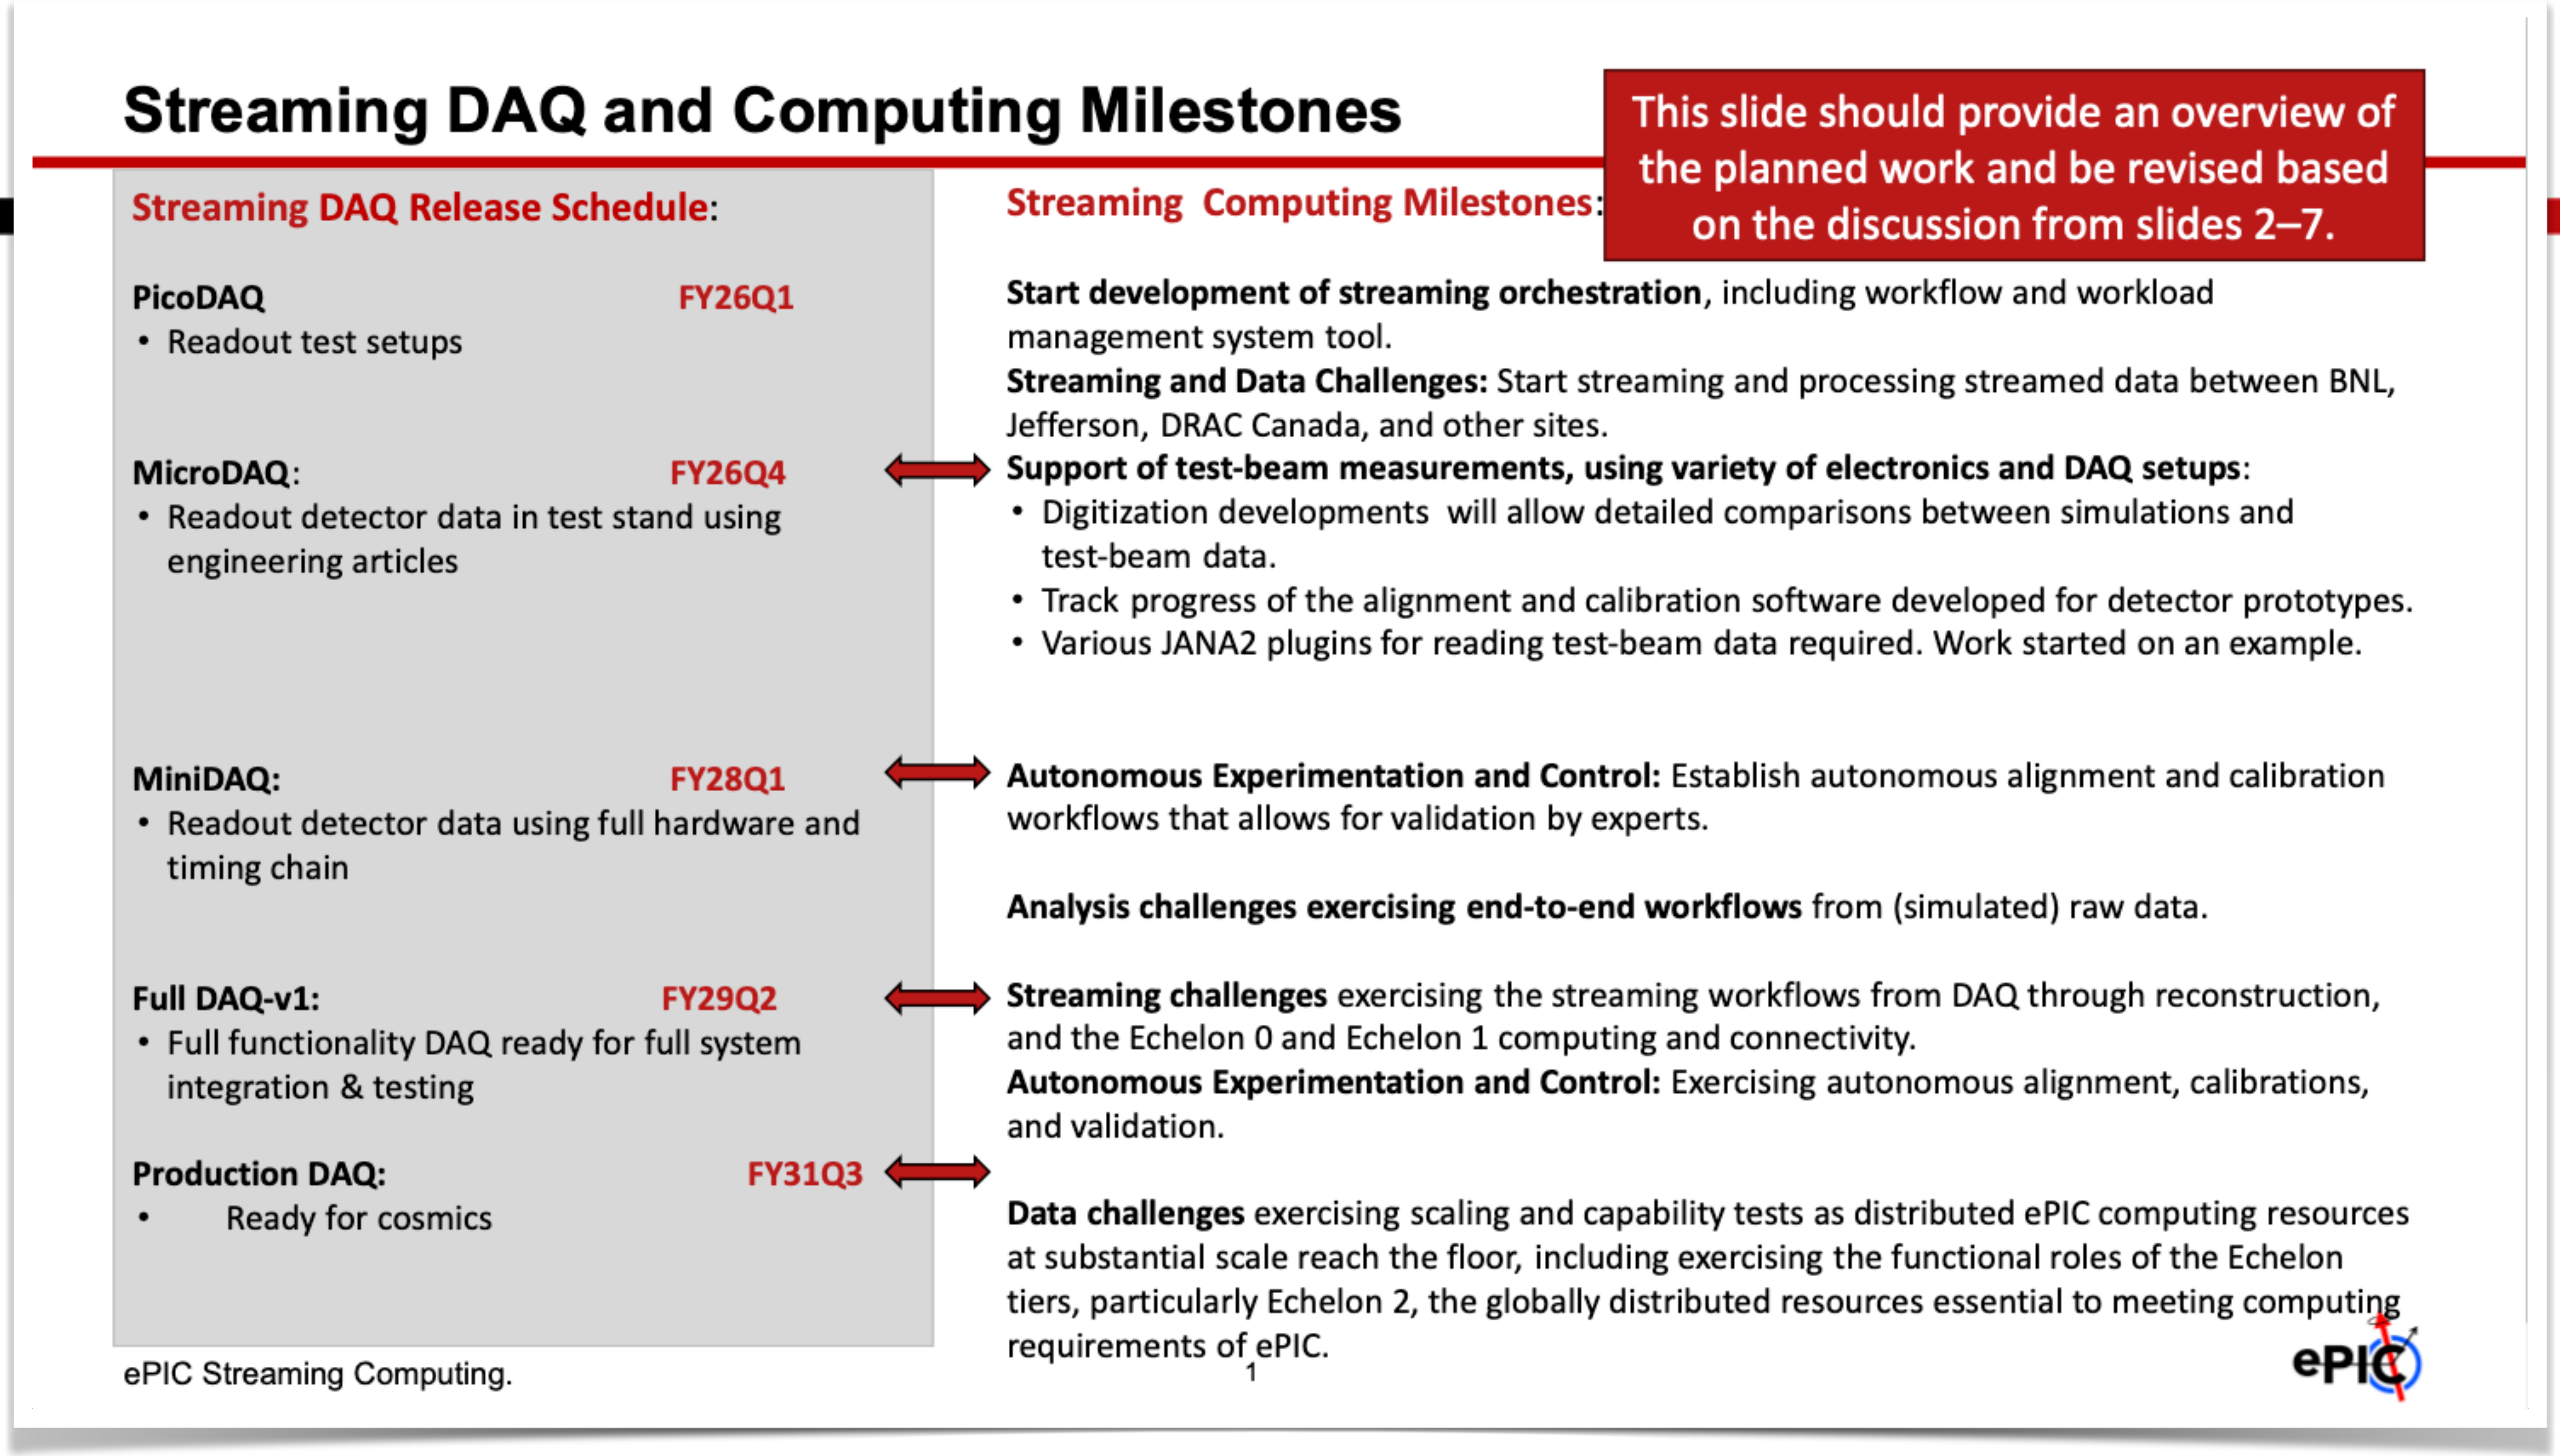

Timeline

1. Establish the **service infrastructure** needed - PanDA, Rucio and associated components (done, at BNL)
2. Update the **E0-E1 schematic** diagram with new elements
3. As the testbed progresses, implement **monitoring of all things** from the beginning, see in real time what the testbed is doing
4. Basic **STF data processing workflow from DAQ egress to E1** archiving & processing
   1. Rucio registration, E1 transfer, downstream triggers
   2. PanDA based batch processing
5. **Stream-based processing**
   1. STF content - metadata describing the data
   2. interfaces and workflow to implement streaming processing based on the the physically stored STF based data
6. **TF stream (fast stream)** processing
7. Basic **calibration workflows**
8. Complex calibration workflows
9. Explore orchestration of streaming processing based on **detector/data state model**

SRO WG activity and documentation


ePIC Software & Computing Report

The ePIC Streaming Computing Model
Version 2, Fall 2024

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Dmitry Kalinkin[6], Jeffery Landgraf[4], David Lawrence[3] and Torre Wenaus[4]
for the ePIC Collaboration

[1]Istituto Nazionale di Fisica Nucleare - Sezione di Genova, Genova, Liguria, Italy.
[2]University of Manitoba, Winnipeg, Manitoba, Canada.
[3]Jefferson Lab, Newport News, VA, USA.
[4]Brookhaven National Laboratory, Upton, NY, USA.
[5]Argonne National Laboratory, Lemont, IL, USA.
[6]University of Kentucky, Lexington, KY, USA.

Abstract

This second version of the ePIC Streaming Computing Model Report provides a 2024 view of the computing model, updating the October 2023 report with new material including an early estimate of computing resource requirements; software developments supporting detector and physics studies, the integration of ML, and a robust production activity; the evolving plan for infrastructure, dataflows, and workflows from Echelon 0 to Echelon 1; and a more developed timeline of high-level milestones. This regularly updated report provides a common understanding within the ePIC Collaboration on the streaming computing model, and serves as input to ePIC Software & Computing reviews and to the EIC Resource Review Board. A later version will be submitted for publication to share our work and plans with the community. **New and substantially rewritten material in Version 2 is dark green. The present draft is preliminary and incomplete and is yet to be circulated in ePIC for review.**

https://doi.org/10.5281/zenodo.14675920

Publication on "ePIC Streaming Computing Model"

Version 1: Prepared for ECSAC review in 2023.
Version 2: Prepared for ECSAC review in 2024.

- An early estimate of computing resource requirements
- Software developments supporting detector and physics studies, the integration of ML
- PlanS for infrastructure, dataflows, and workflows from Echelon 0 to Echelon 1
- A developed timeline of high-level milestones

✳ Several reviews focused on SRO and/or on frontend electronics/backend reconstruction software

✳In progress: ECSAC Review 2025

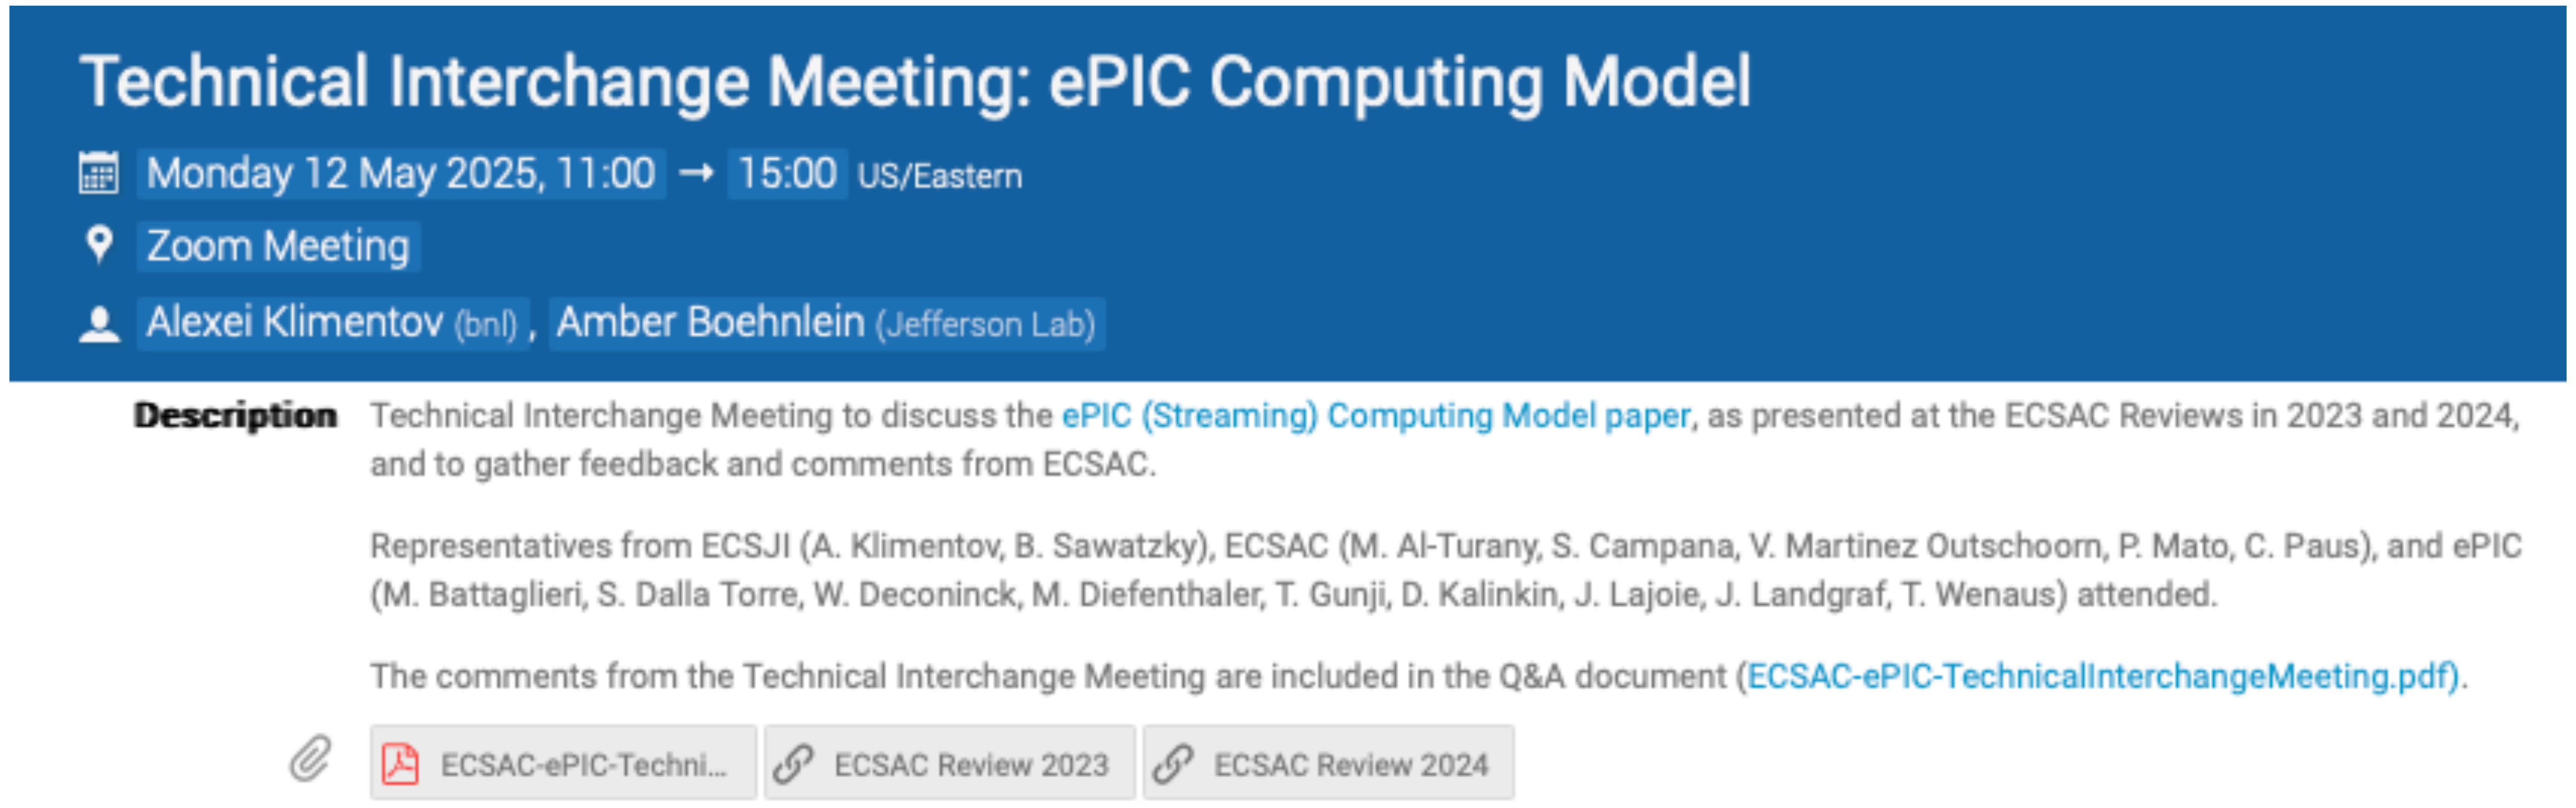

✳ Next date: EICUG/ ePIC Collaboration meeting: Streaming readout workfest

Streaming readout (SRO) workshop series

✳ **Last edition SRO XII in Tokyo in december 2024**

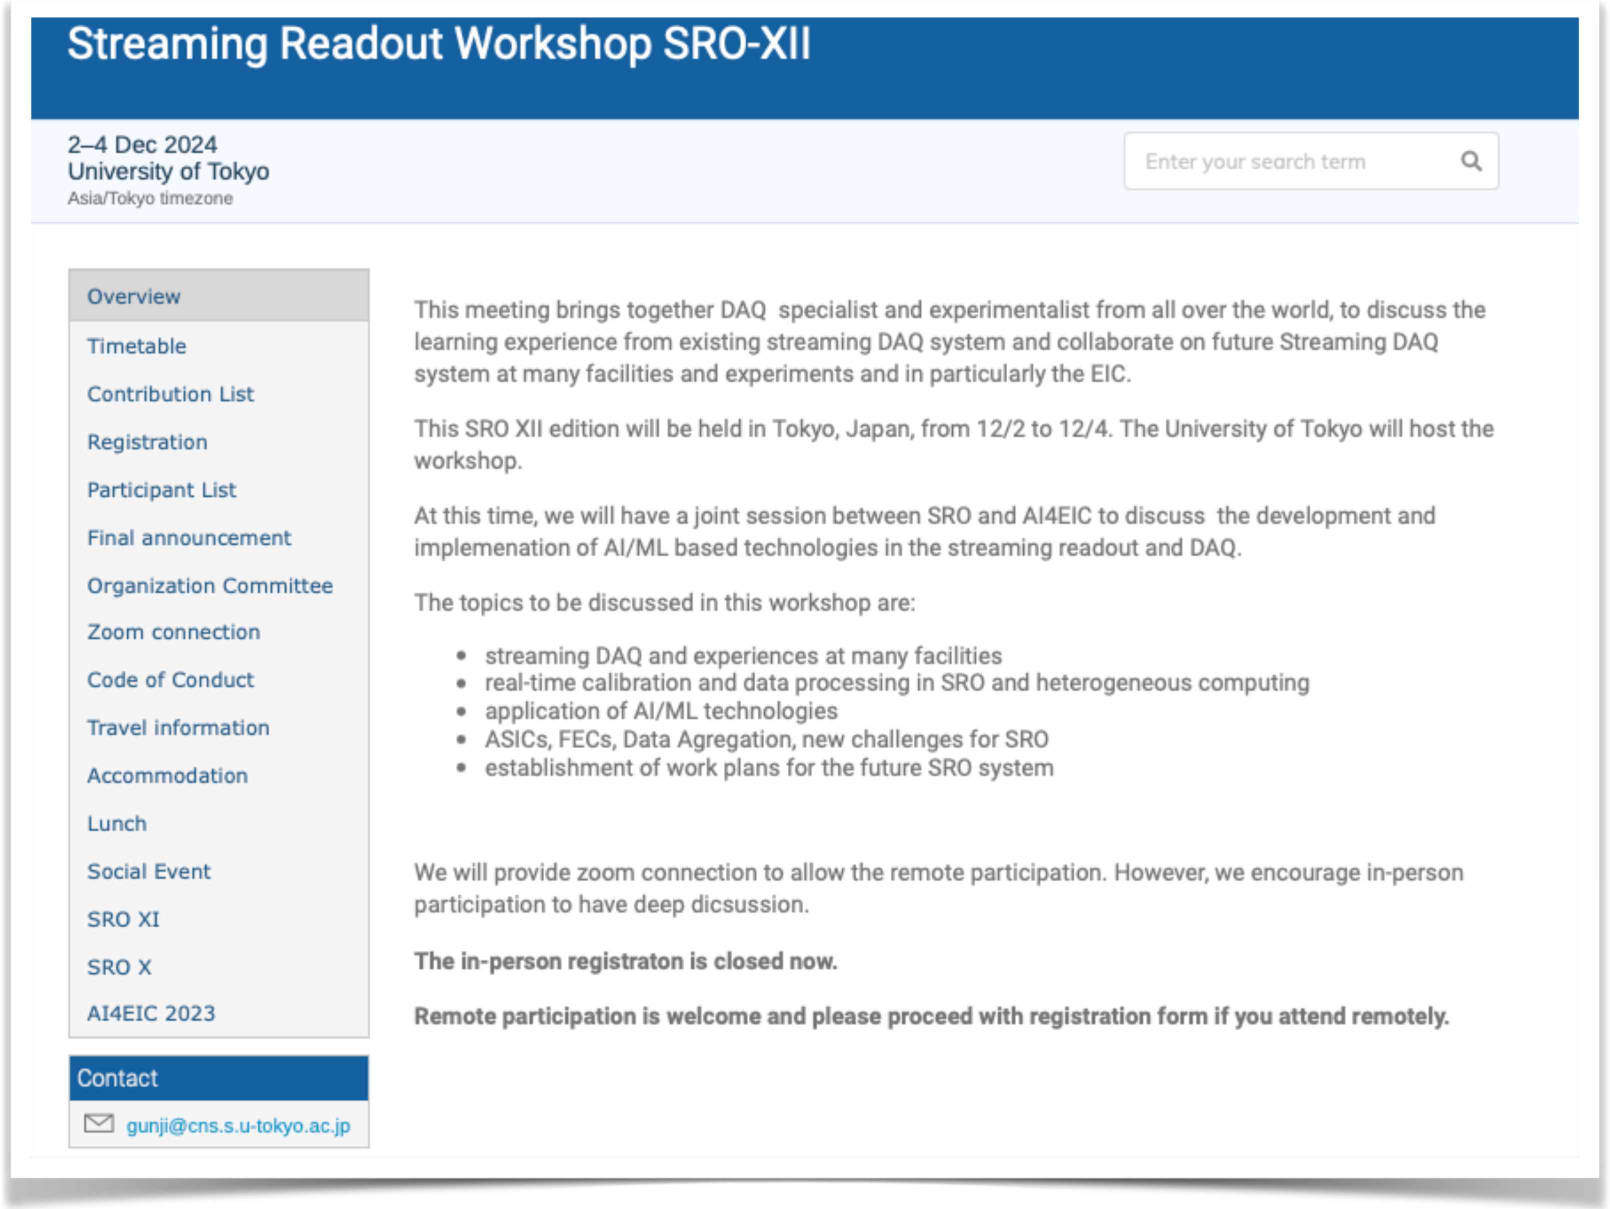

Streaming Readout Workshop SRO-XII

2–4 Dec 2024
University of Tokyo
Asia/Tokyo timezone

Enter your search term

- Overview
- Timetable
- Contribution List
- Registration
- Participant List
- Final announcement
- Organization Committee
- Zoom connection
- Code of Conduct
- Travel information
- Accommodation
- Lunch
- Social Event
- SRO XI
- SRO X
- AI4EIC 2023

Contact
gunji@cns.s.u-tokyo.ac.jp

This meeting brings together DAQ specialist and experimentalist from all over the world, to discuss the learning experience from existing streaming DAQ system and collaborate on future Streaming DAQ system at many facilities and experiments and in particularly the EIC.

This SRO XII edition will be held in Tokyo, Japan, from 12/2 to 12/4. The University of Tokyo will host the workshop.

At this time, we will have a joint session between SRO and AI4EIC to discuss the development and implemenation of AI/ML based technologies in the streaming readout and DAQ.

The topics to be discussed in this workshop are:

- streaming DAQ and experiences at many facilities
- real-time calibration and data processing in SRO and heterogeneous computing
- application of AI/ML technologies
- ASICs, FECs, Data Agregation, new challenges for SRO
- establishment of work plans for the future SRO system

We will provide zoom connection to allow the remote participation. However, we encourage in-person participation to have deep dicsussion.

The in-person registraton is closed now.

Remote participation is welcome and please proceed with registration form if you attend remotely.

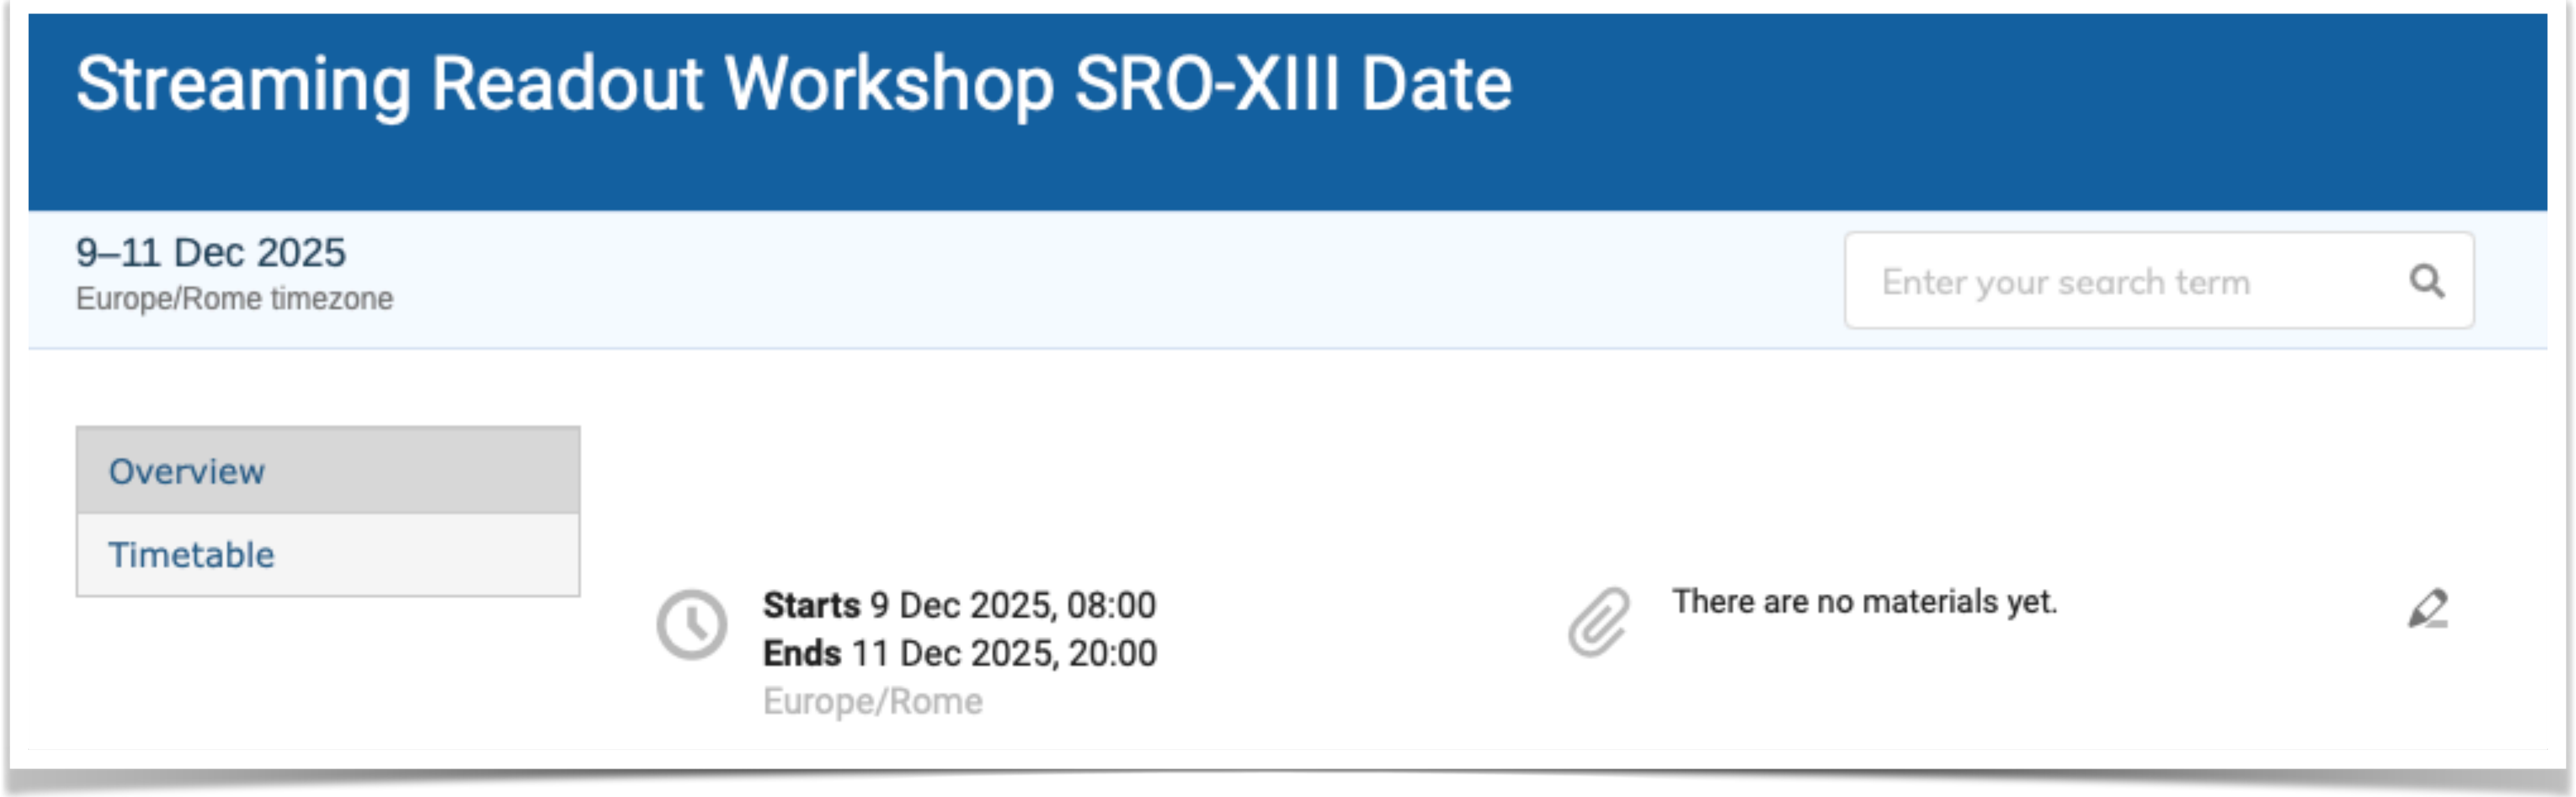

Streaming Readout Workshop SRO-XIII Date

9–11 Dec 2025
Europe/Rome timezone

Enter your search term

- Overview
- Timetable

Starts 9 Dec 2025, 08:00
Ends 11 Dec 2025, 20:00
Europe/Rome

There are no materials yet.

✳ **Next Streaming ReadOut workshop (SRO XIII) will be held from Dec 9 to 11 2025 at Auditorium ex Chiesa della Purita' in Catania (IT)**
- LOC: M.Battaglieri, J.Bernauer, M.Bondî, A.Camsonne, M.Diefenthaler, T.Gunji, M.Spreafico, C.Tuvé, S.Vallarino
- Topics: SRO testbeds (NESTDAQ, CODA, RCDAQ, JUNO, KM3NET), …, software (RUCIO, ALLEN, …), real-time calibration (AI solutions), RT data processing (Leonardo, object data storage, data transfer, …)
- Demonstration session

Summary and outlook

- ePIC DAQ will be streaming
- SRO is expected to facilitate (and extend!) the ePIC science and shorten the time between data taking and physics output
- SRO provides new opportunities but poses challenges (e.g. new detector-computing relationship)
- Strong link with DAQ, SW, and Physic WGs to design a sound framework (+ trigger/analysis algorithms)
- Key components of SRO workflows are under definitions
  - Data concentrator
  - Data reduction
  - Orchestrator
  - Time Frame builder and router
  - Data management
- AI-supported algorithms will play a significant role in ePICS SRO computing model
- Alignment/calibration procedures are under definitions with feedback from sub-detector groups
- Inputs from physics analysis/simulations is necessary to shape the framework
- On-field validation needs to be pursued in parallel
- Strong commitment from ePIC-Itay with leading roles in management and contributions

See you in Catania at SRO-XIII in December!