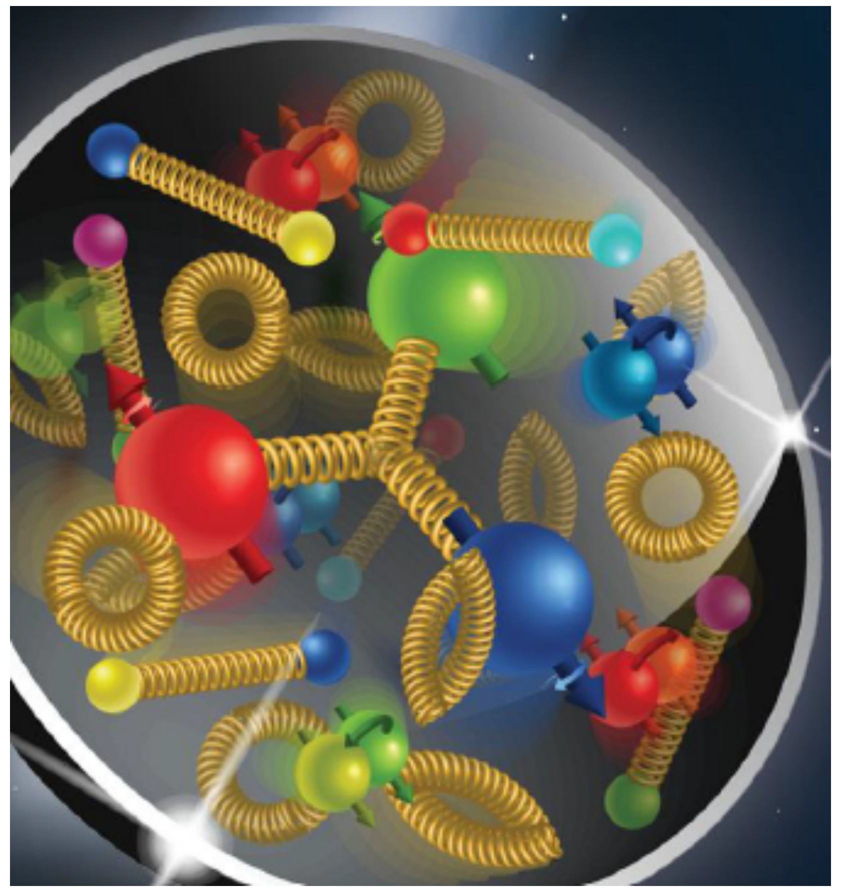
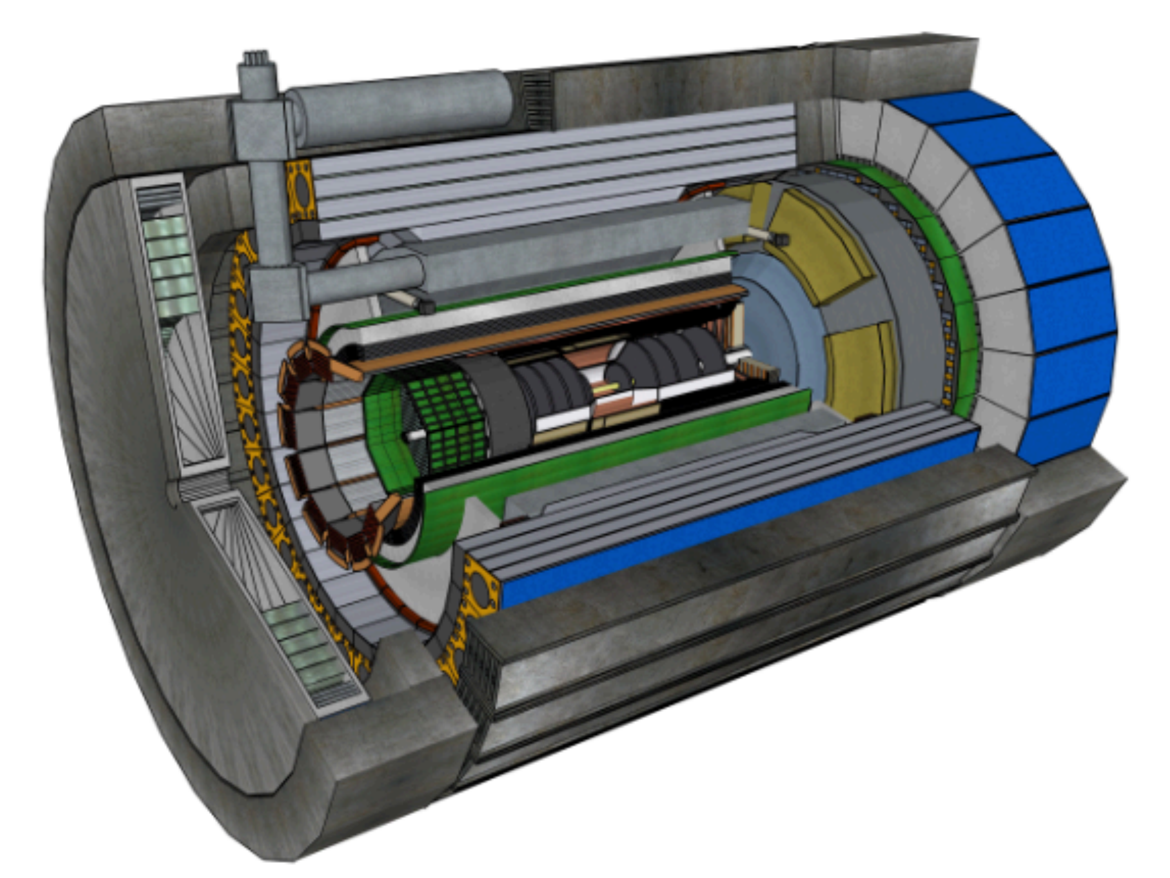
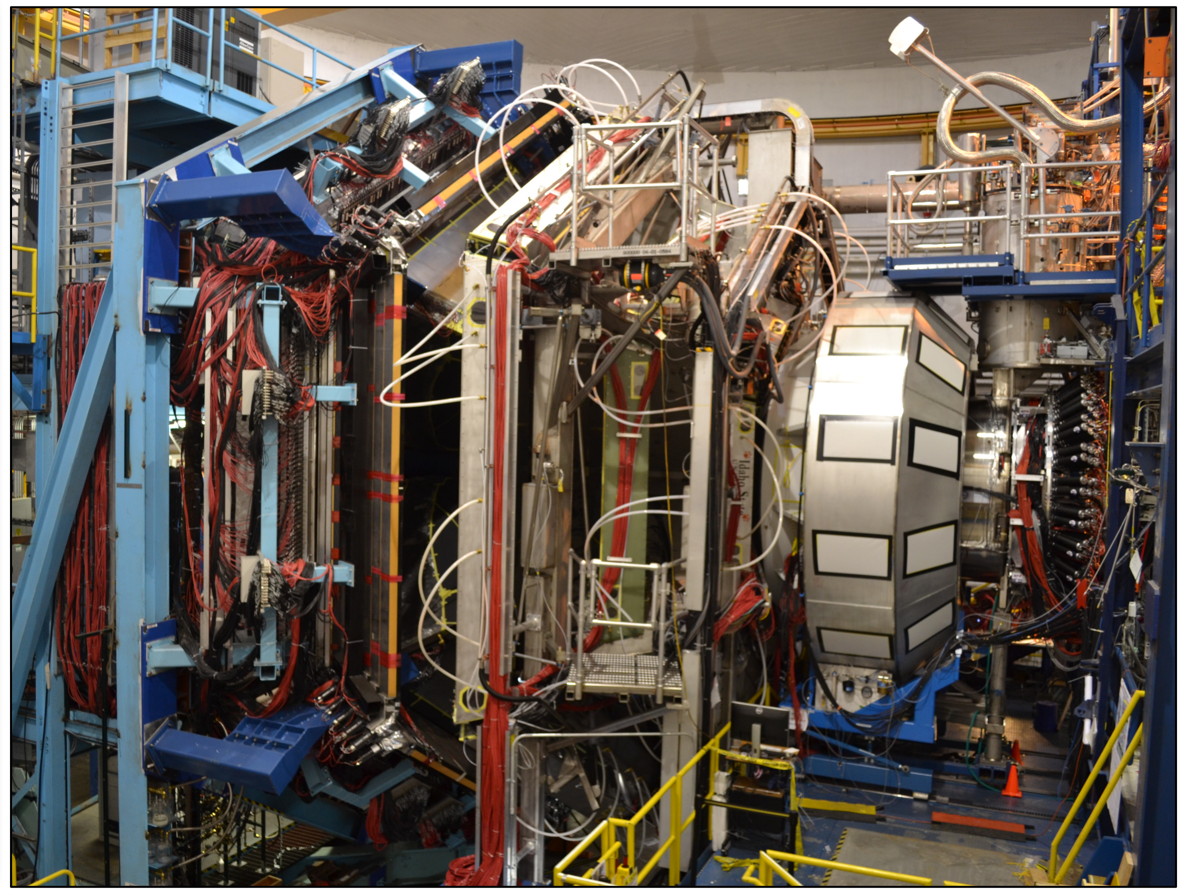



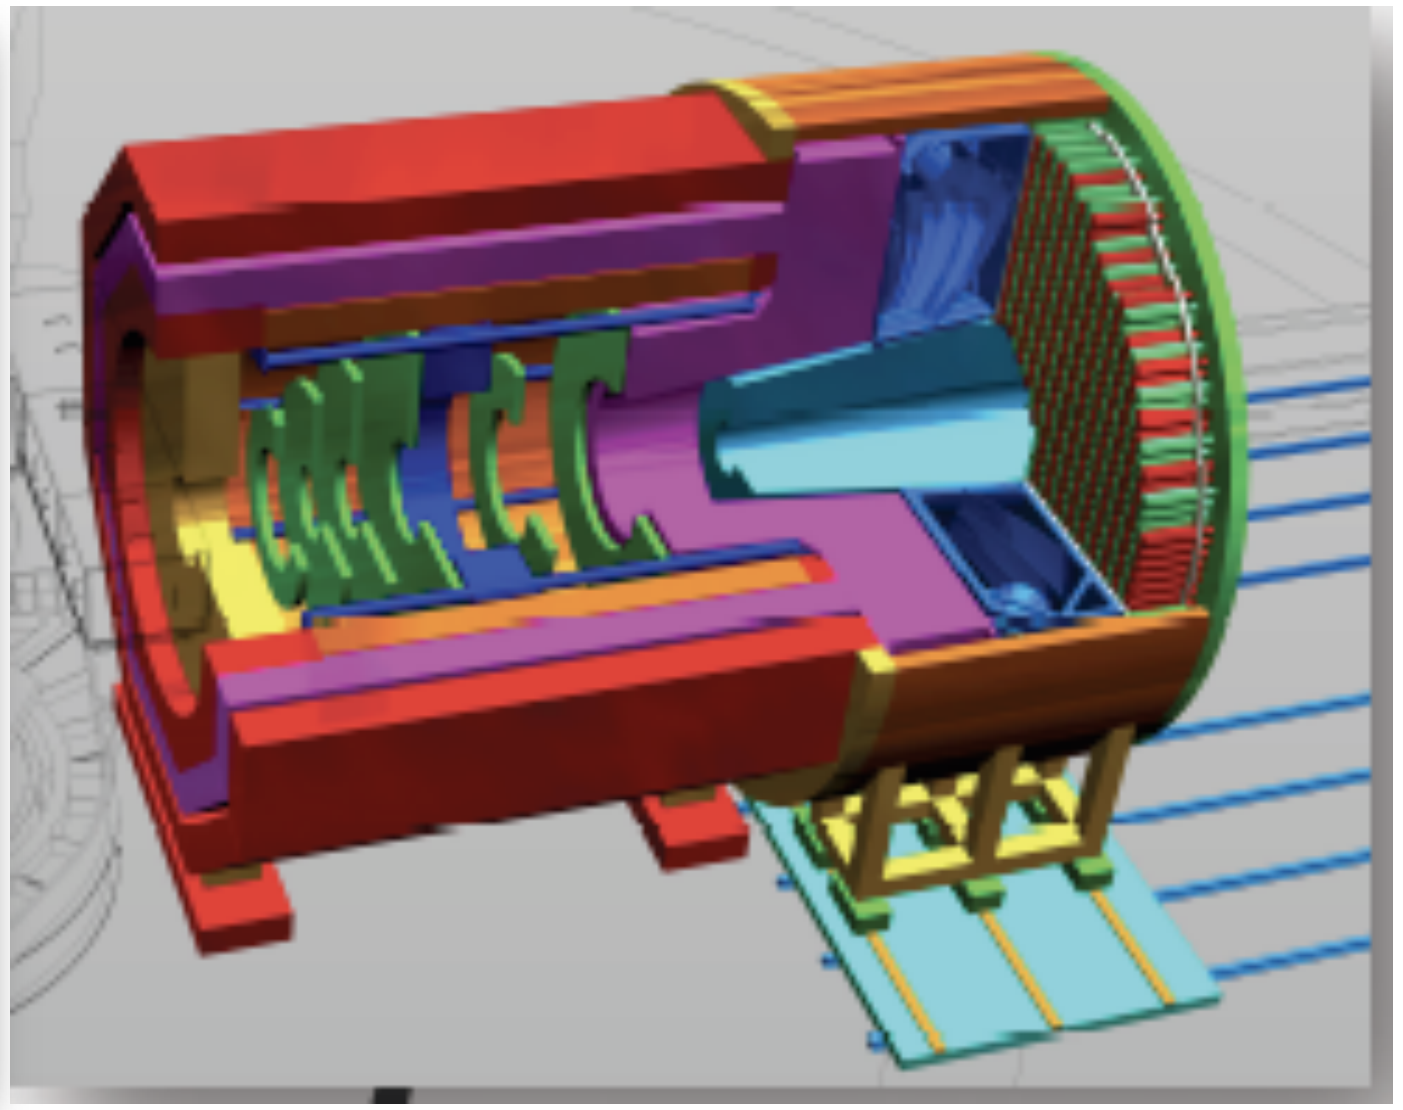

Jefferson Lab streaming readout system

M.Battaglieri (INFN)

From signals to physics

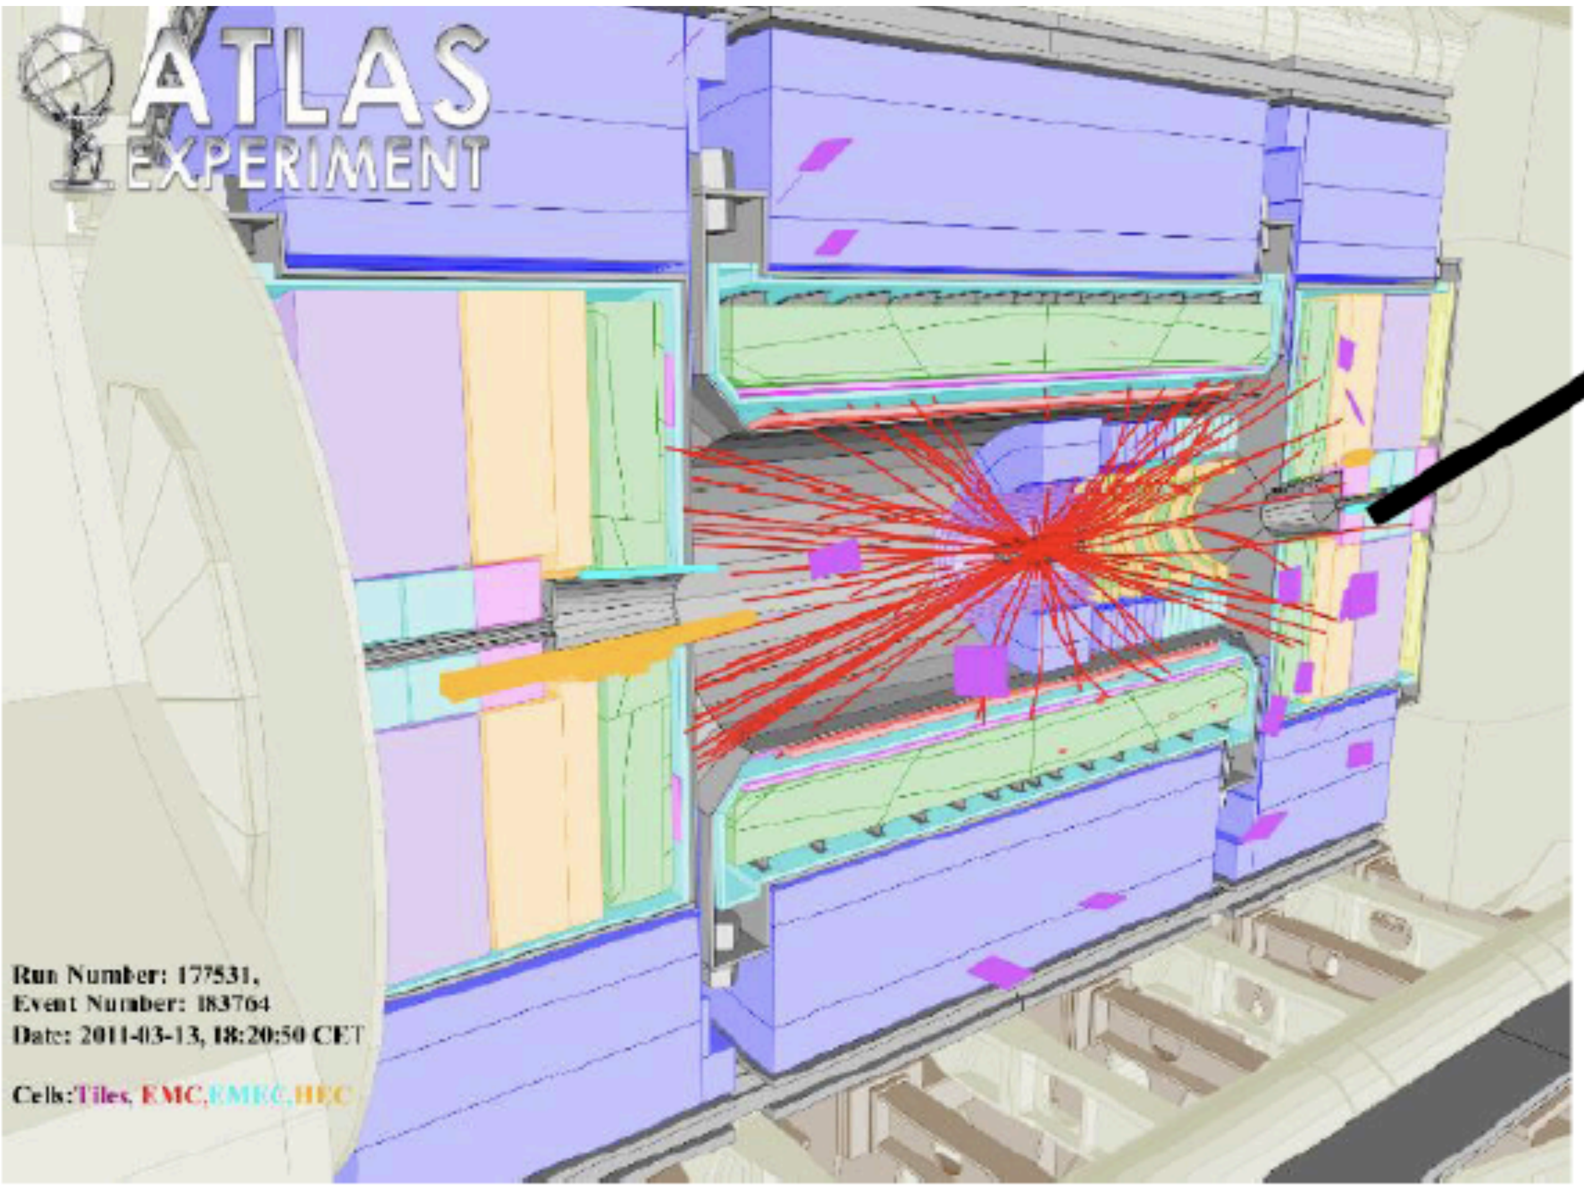

DAQ chain

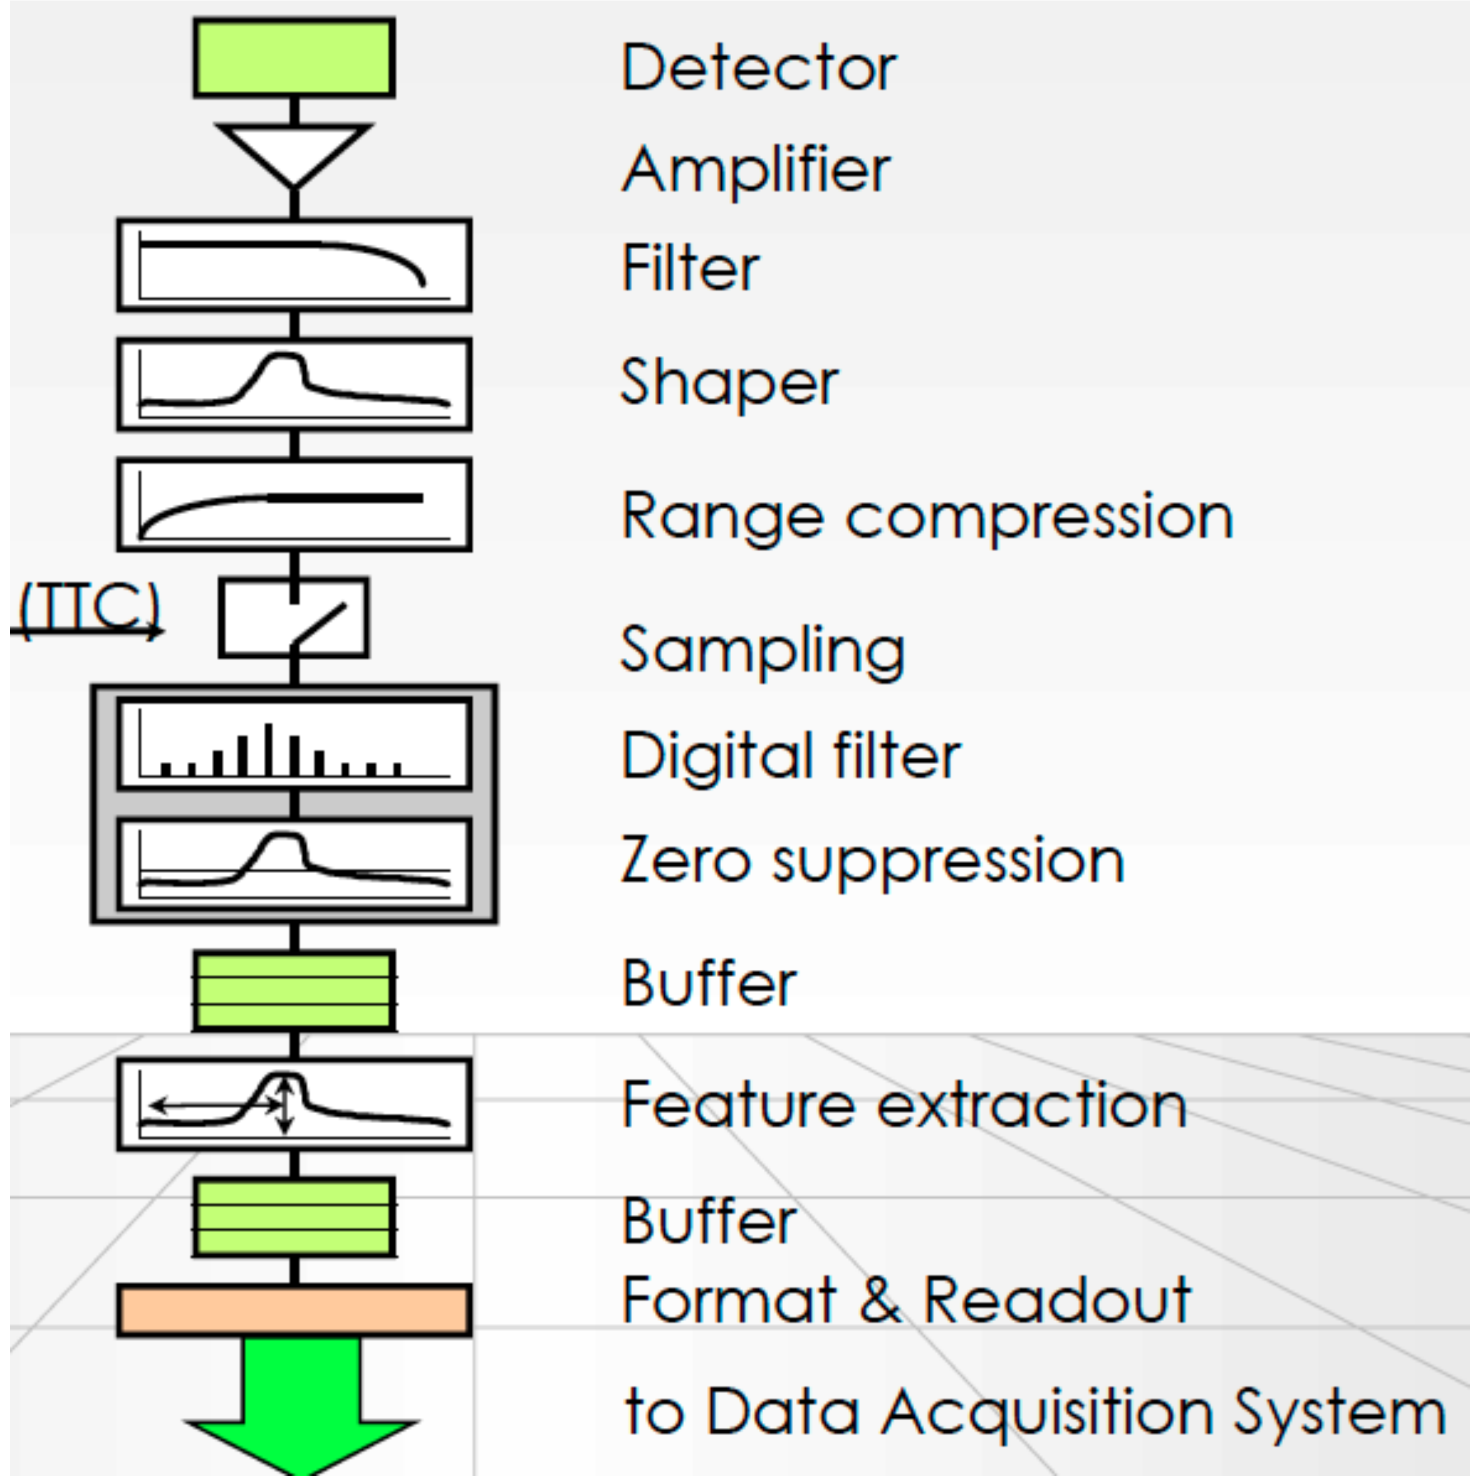

Triggered DAQ

Traditional (triggered) DAQ

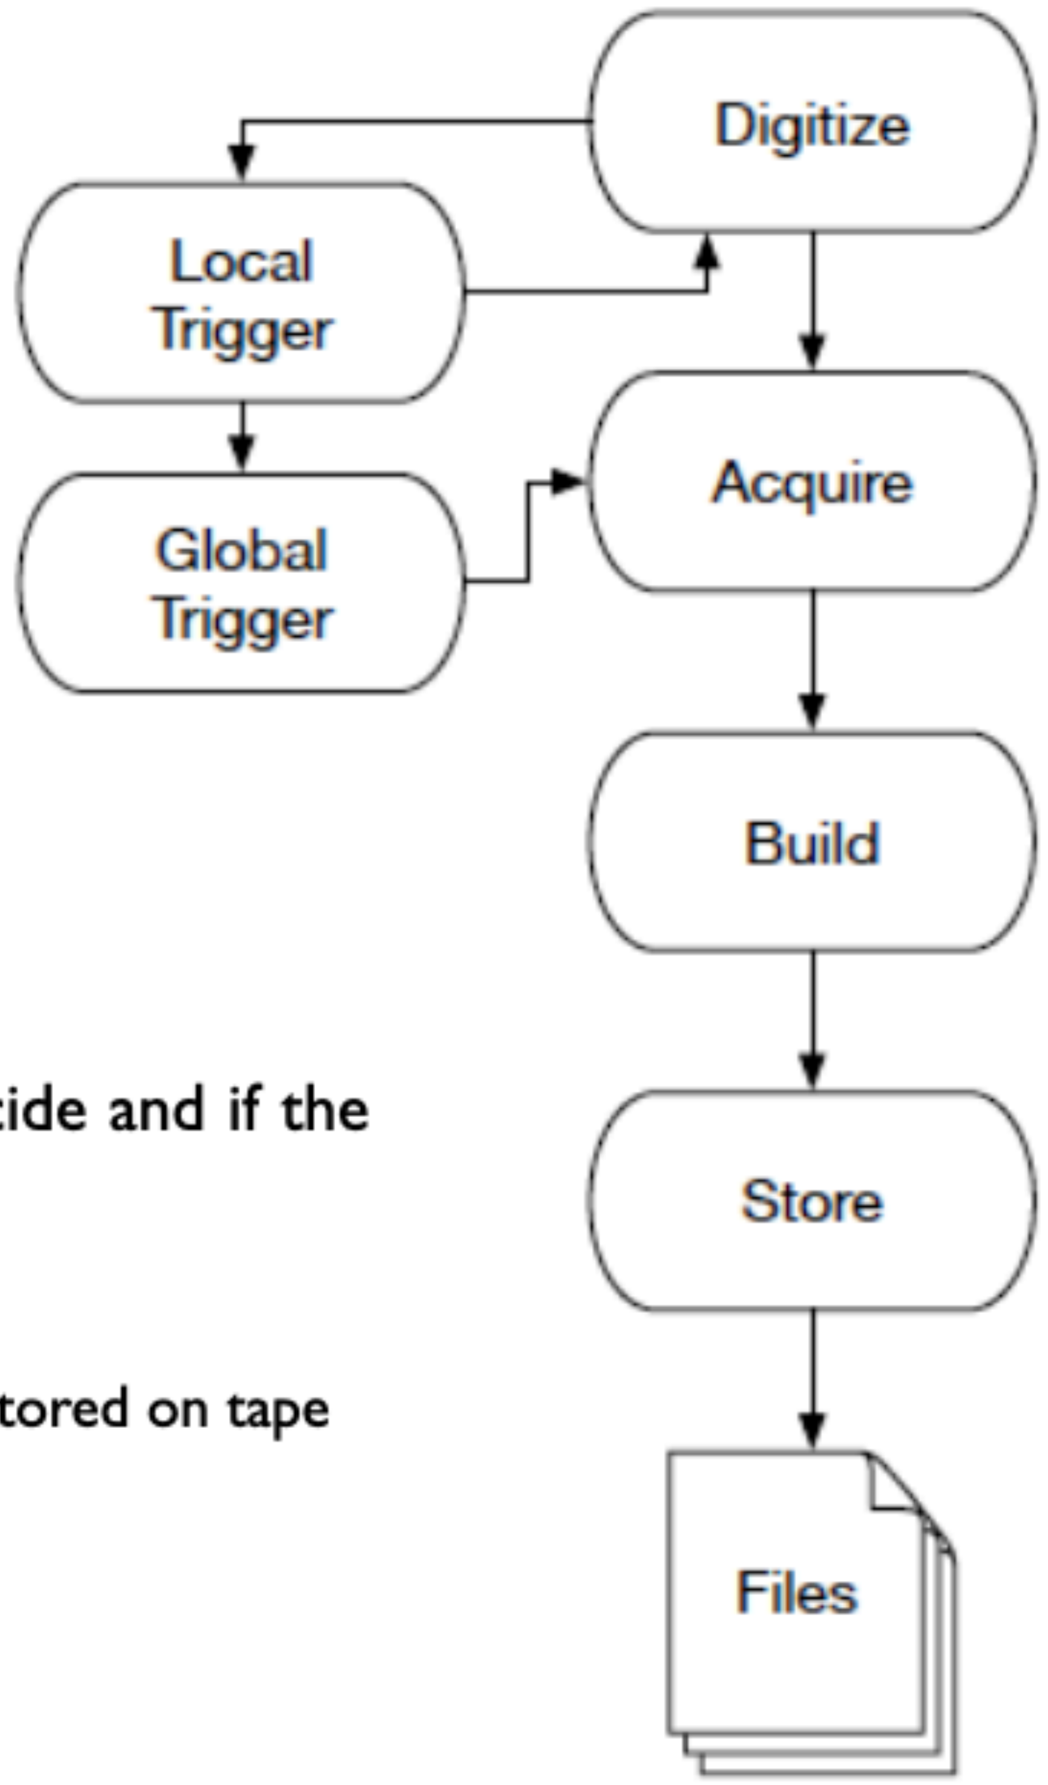

- (few) trigger Channels participating send (partial) information to trigger logic
- All channels continuously measured, hits stored in short term memory
- Trigger logic takes time to decide and if the trigger condition is satisfied:
  - a new 'event' is defined
  - trigger signal back to the FEE
  - data read from memory and stored on tape

Traditional triggered DAQ

- **Pros**
  - we know it works reliably!
- **Drawbacks:**
  - only few information forms the trigger
  - Trigger logic (FPGA) difficult to implement and debug
  - not easy to change and adapt to different conditions

Streaming RO

Streaming read out (SRO)

✴ A HIT MANAGER receives hits from FEE, order them and ship to the software defined trigger

✴ Software defined trigger re-aligns in time the whole detector hits applying a selection algorithm to the time-slice

- the concept of 'event' is lost
- time-stamp is provided by a synchronous common clock distributed to each FEE

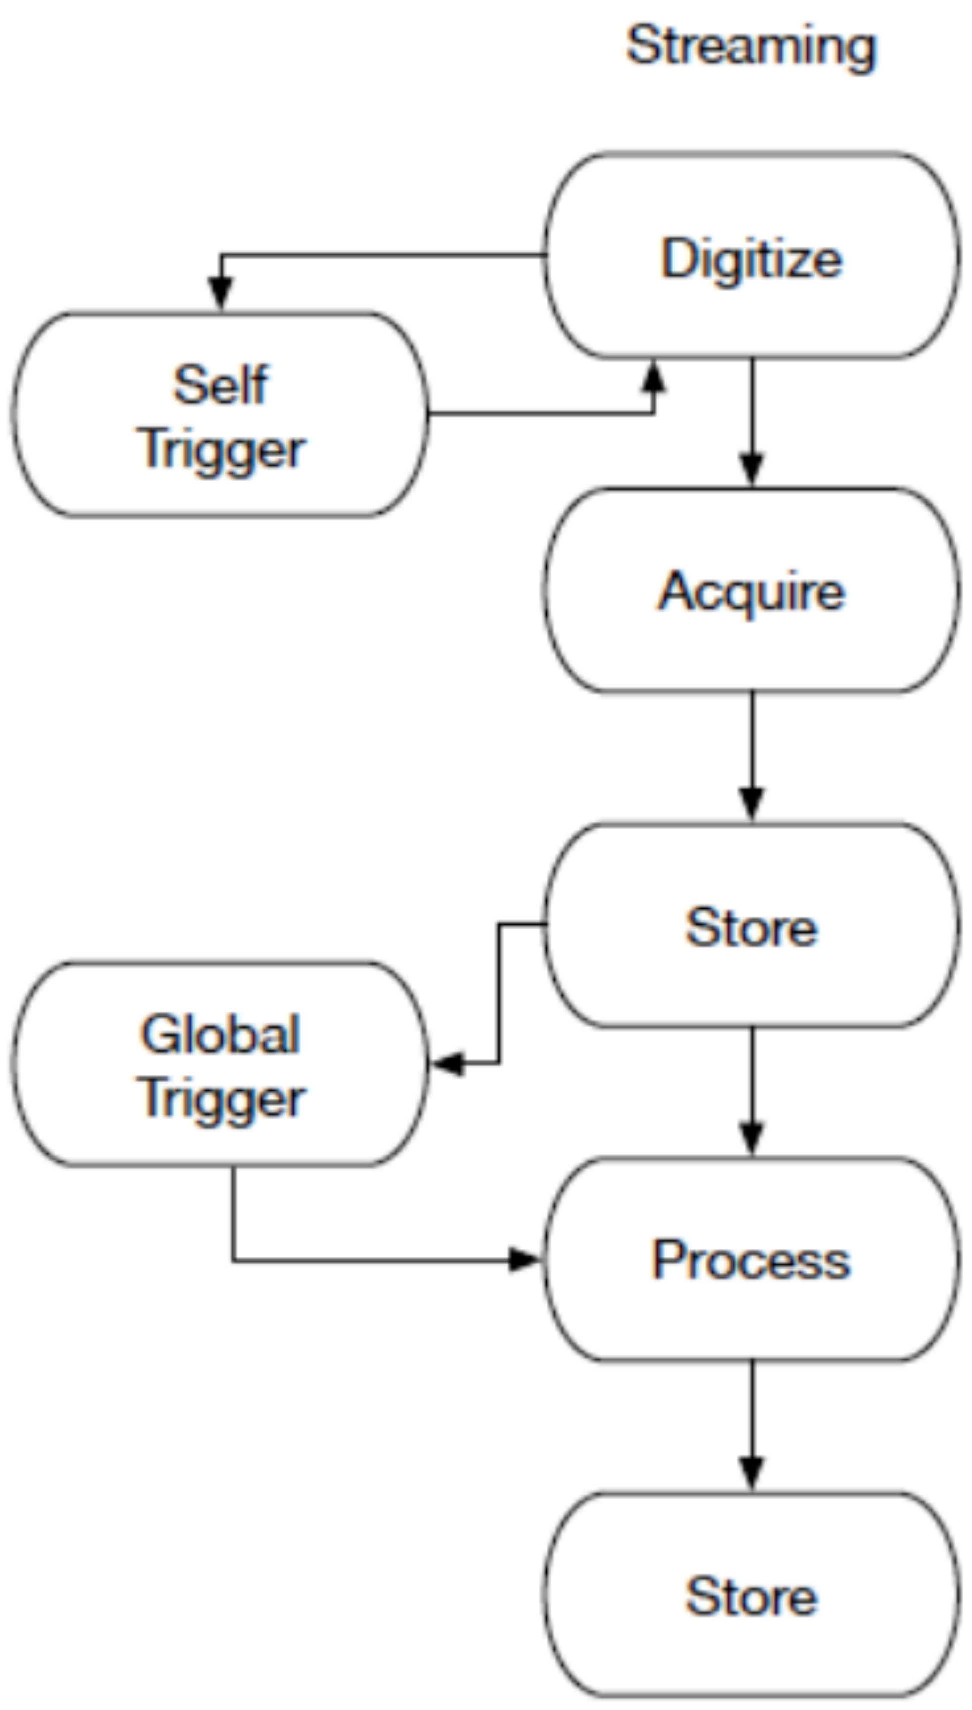

✴ All channels continuously measured and hits streamed to a HIT manager (minimal local processing) with a time-stamp

SRO DAQ

- **Pros**
  - All channels can be part of the trigger
  - Sophisticated tagging/filtering algorithms
  - high-level programming languages
  - scalability
- **Drawbacks:**
  - we do not have the same experience as for TRIGGERED DAQ

Why SRO is so important?

✴ **High luminosity experiments**
- Write out the full DAQ bandwidth
- Reduce stored data size in a smart way (reducing time for off-line processing)

✴ **Shifting data tagging/filtering from the front-end (hw) to the back-end (sw)**
- Optimize real-time rare/exclusive channel selection
- Use of high-level programming languages
- Use of existing/ad-hoc CPU/GPU farms
- Use of available AI/ML tools
- (future) use of quantum-computing

✴ Scaling
- Easier to add new detectors in the DAQ pipeline
- Easier to scale
- Easier to upgrade

Many NP and HEP experiments adopt a SRO DAQ

- CERN: LHCb, ALICE, AMBER
- FAIR: CBM
- DESY: TPEX
- FRIBS: GRETA
- BNL: sPHENIX
- JLAB: SOLID, BDX, CLAS12, …

Jefferson Lab

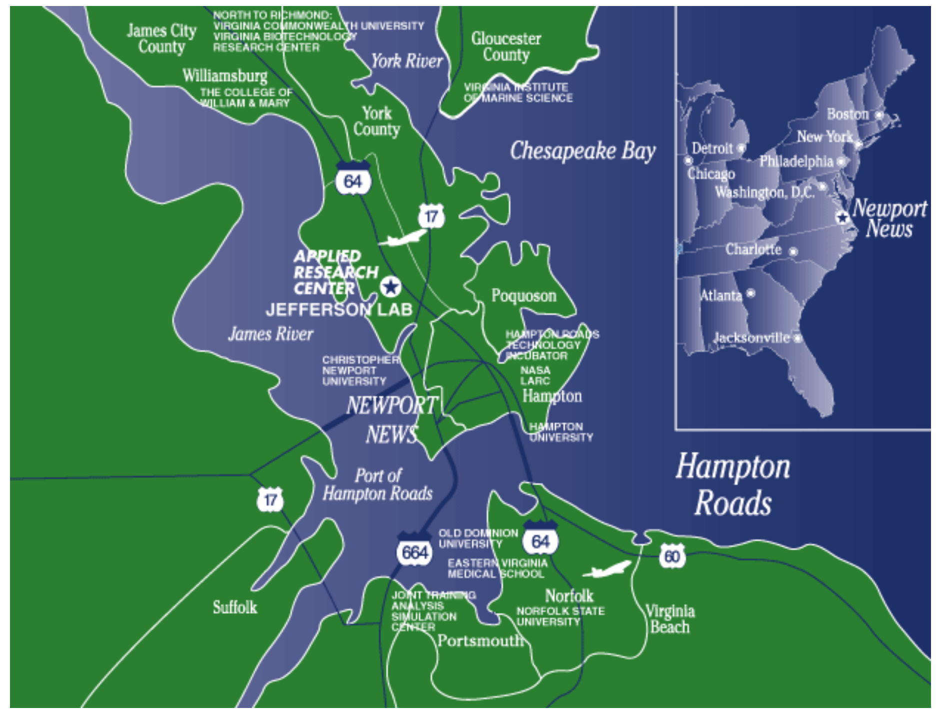


Jefferson Lab's accelerator site

✳Primary Beam: Electrons

✳ Beam Energy: 12 GeV
- 10 > λ > 0.1 fm
- nucleon → quark transition
- baryon and meson excited states

✳100% Duty Factor (cw) Beam
- coincidence experiments
- Four simultaneous beams
- Independent E and I

✳ Polarization
- spin degrees of freedom
- weak neutral currents

Luminosity > 10^7 - 10^8 x SLAC at the time of the original DIS experiments!

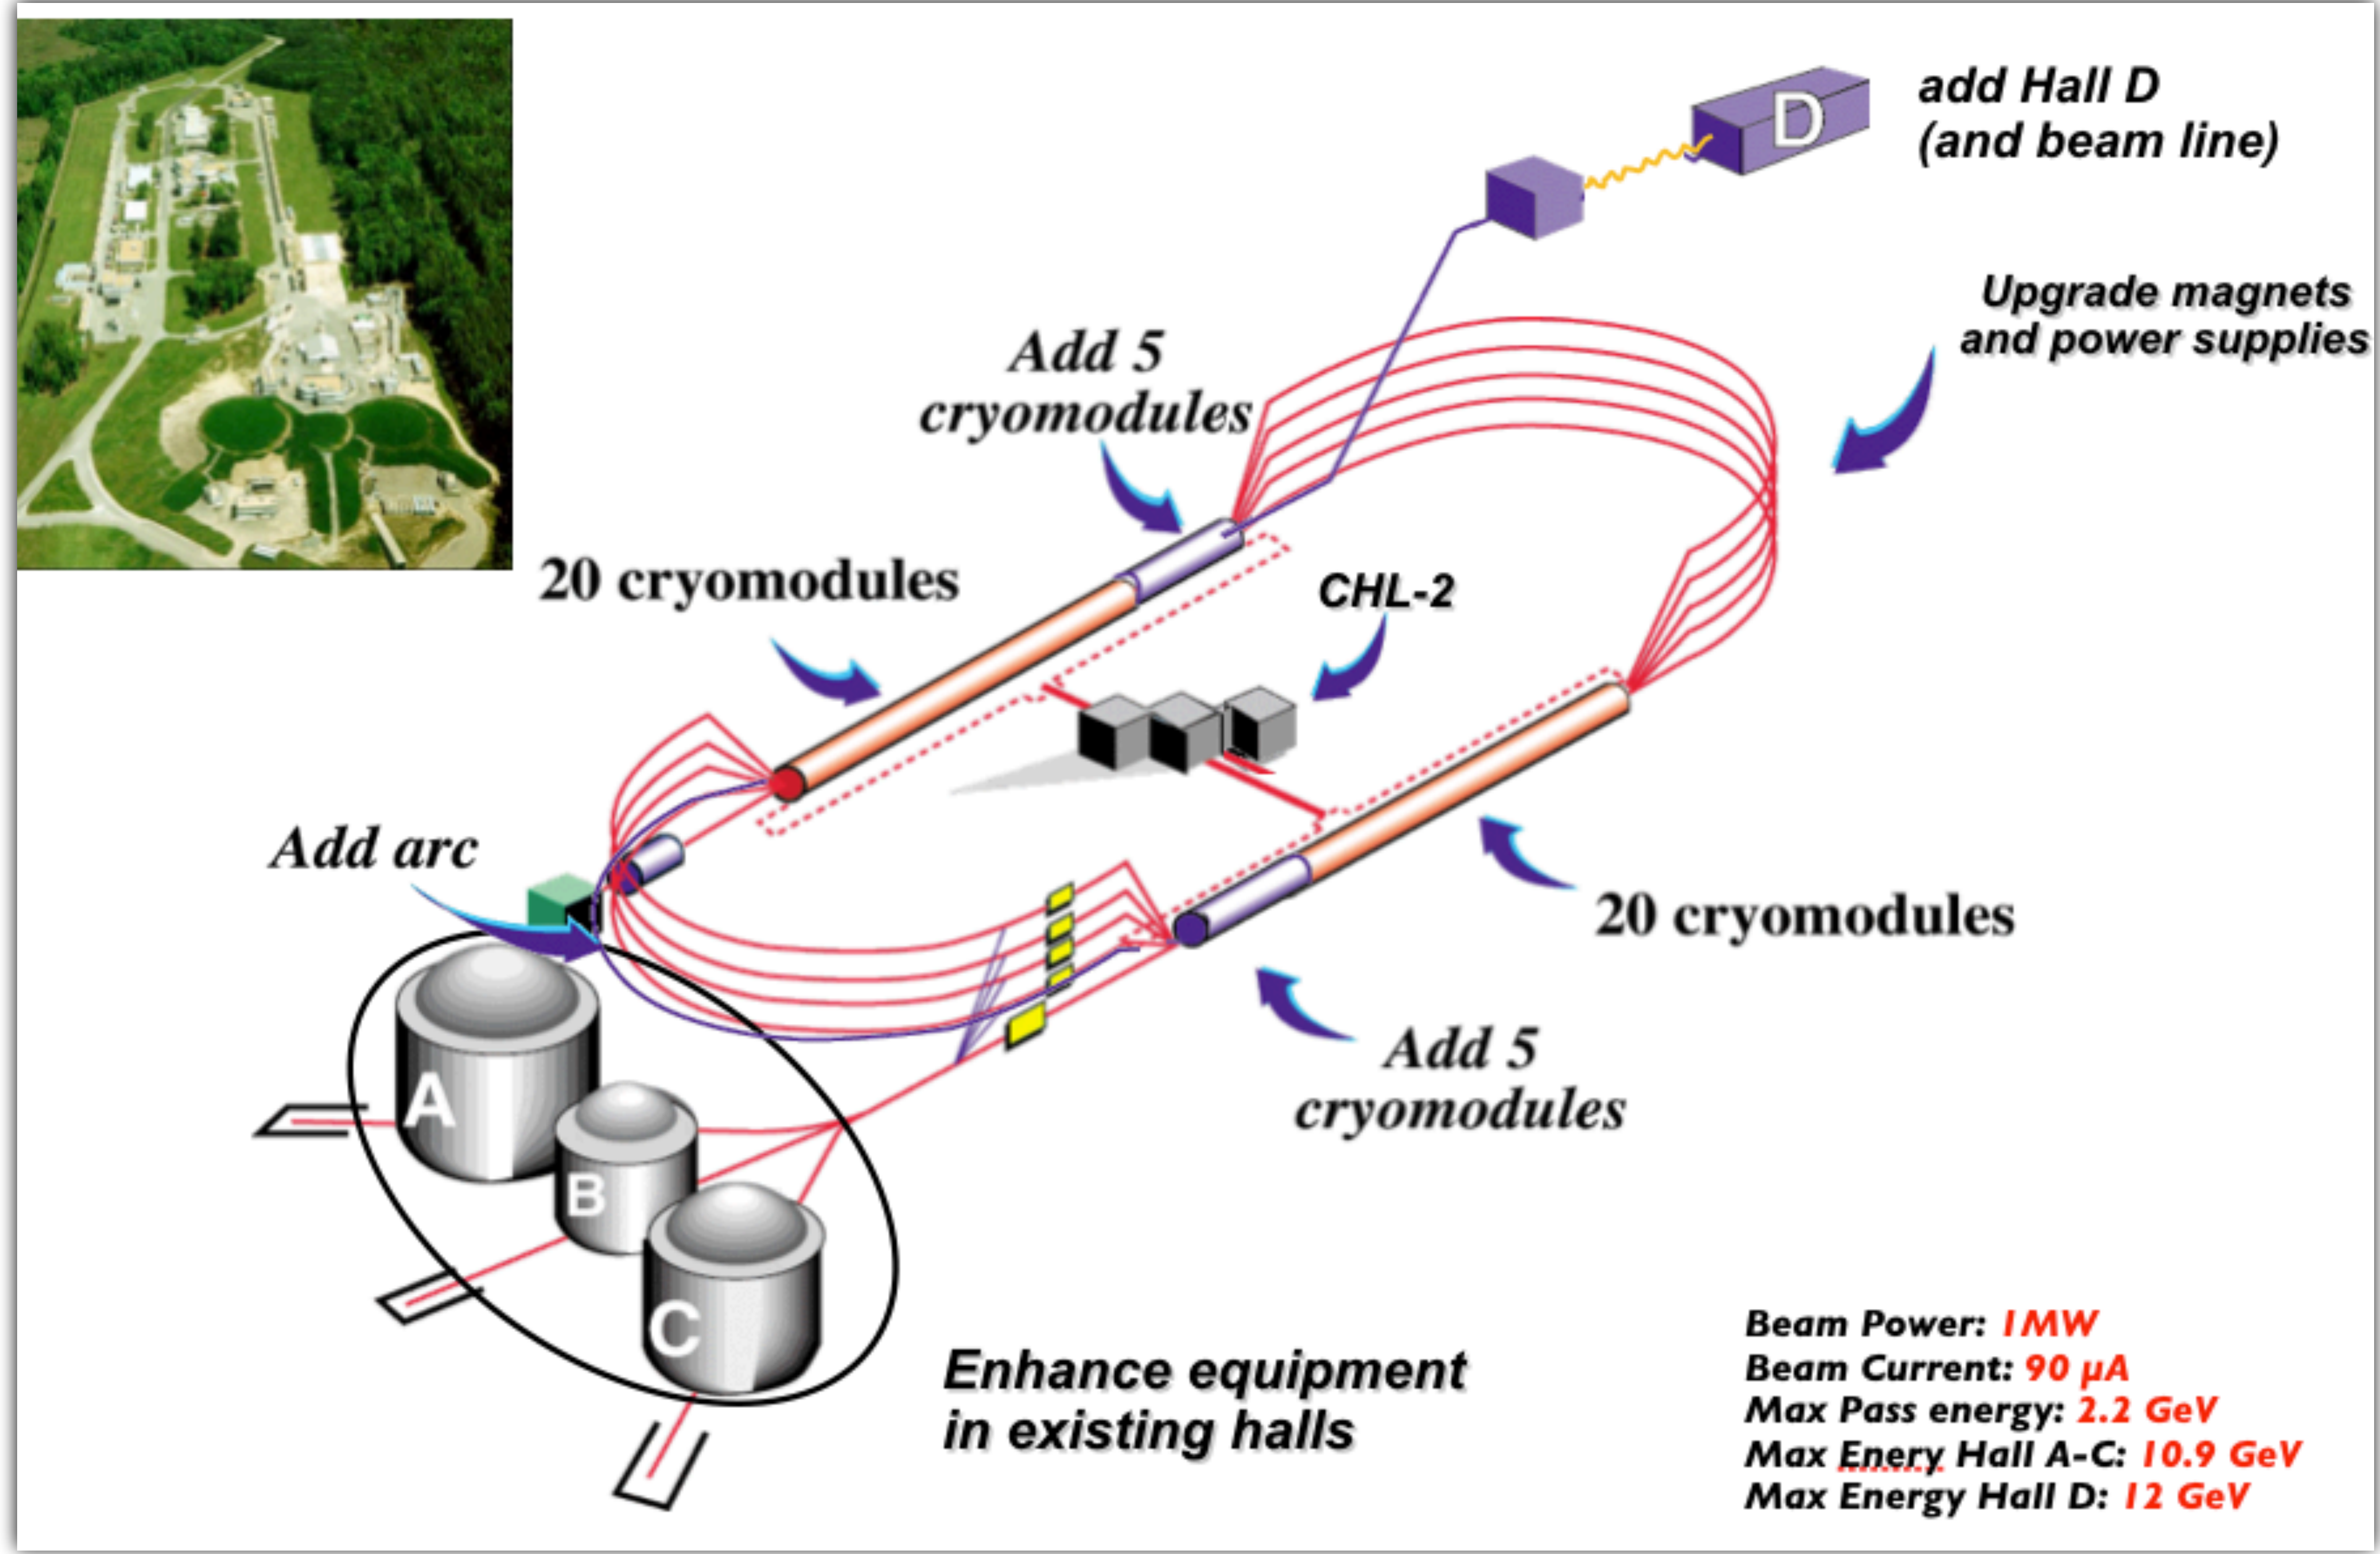

Hall A

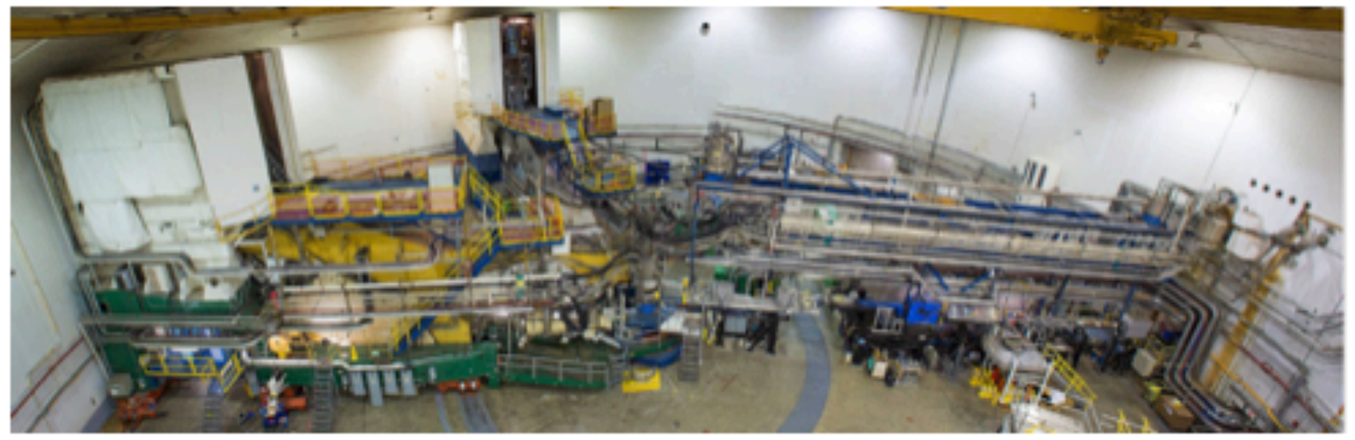

Hall C

Hall B

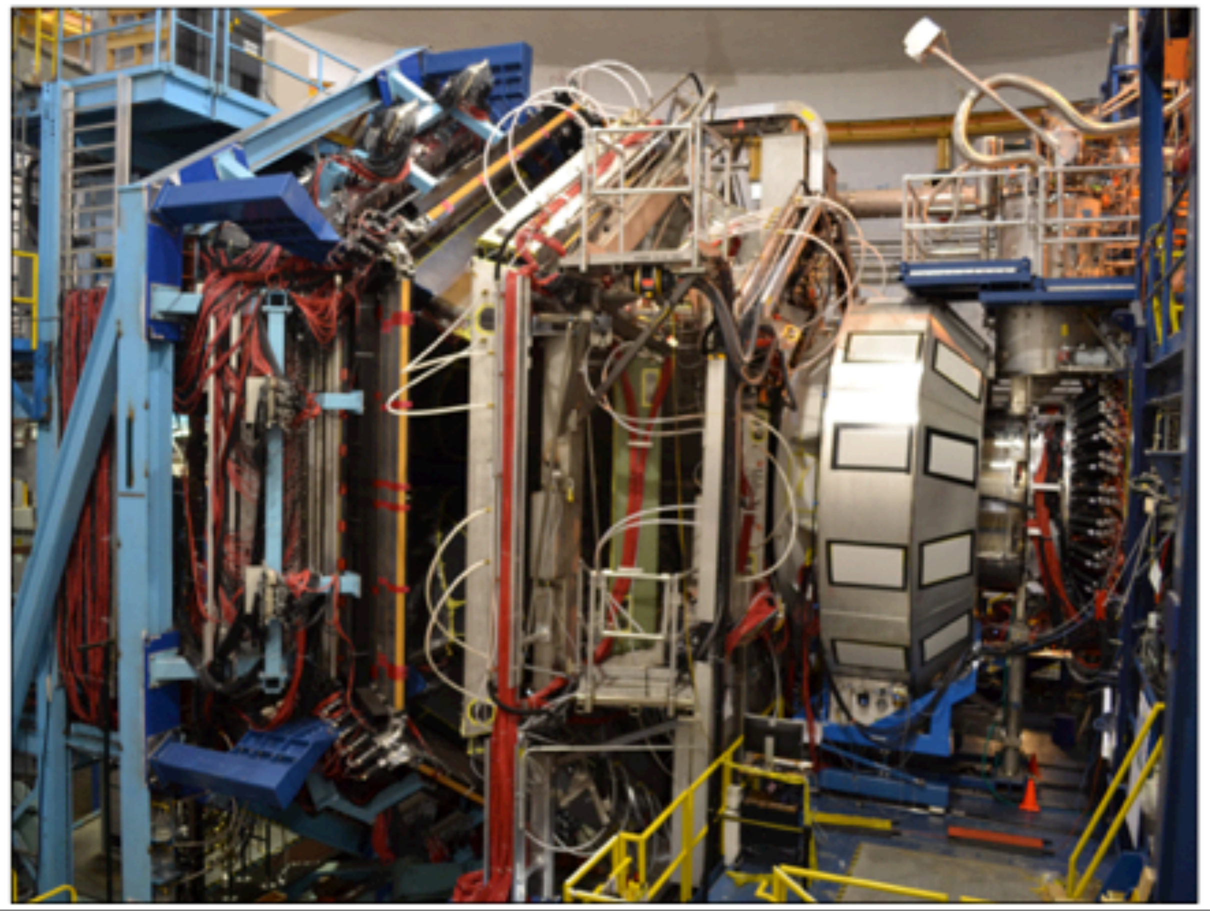

Hall D

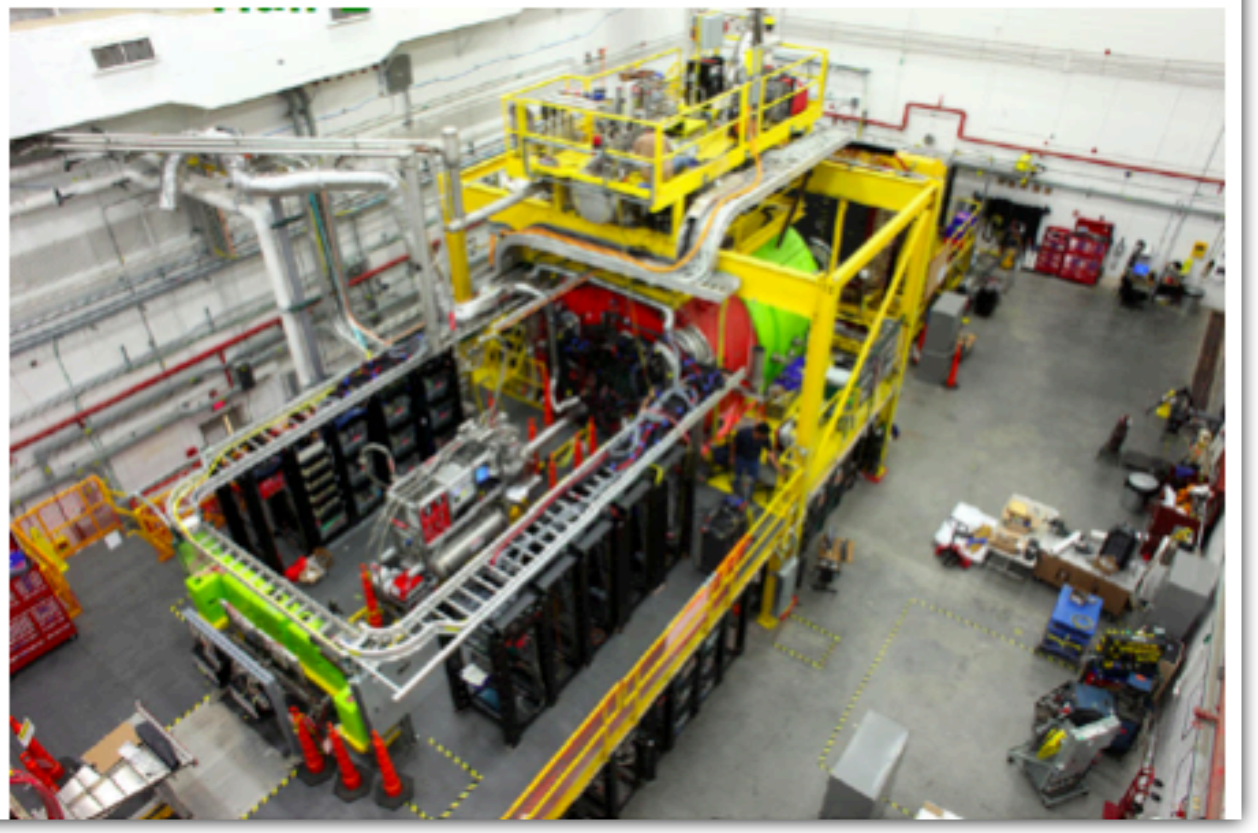

Jefferson Lab

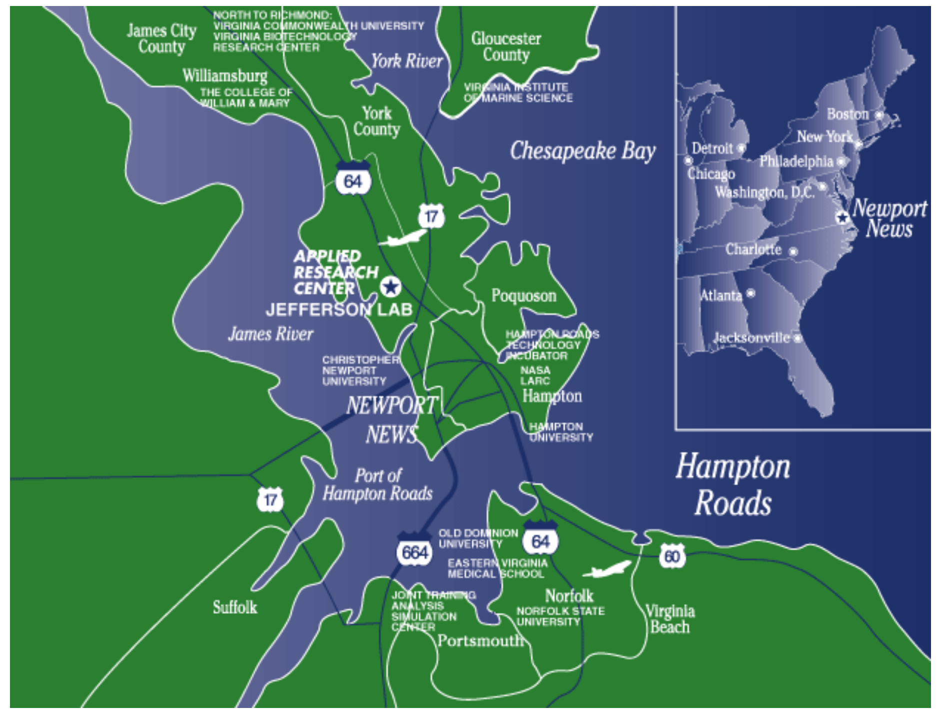


Jefferson Lab's accelerator site

✳Primary Beam: Electrons

✳ Beam Energy: 12 GeV
- 10 > λ > 0.1 fm
- nucleon → quark transition
- baryon and meson excited states

✳100% Duty Factor (cw) Beam
- coincidence experiments
- Four simultaneous beams
- Independent E and I

✳ Polarization
- spin degrees of freedom
- weak neutral currents

Luminosity > 10^7 -10^8 x SLAC at the time of the original DIS experiments!

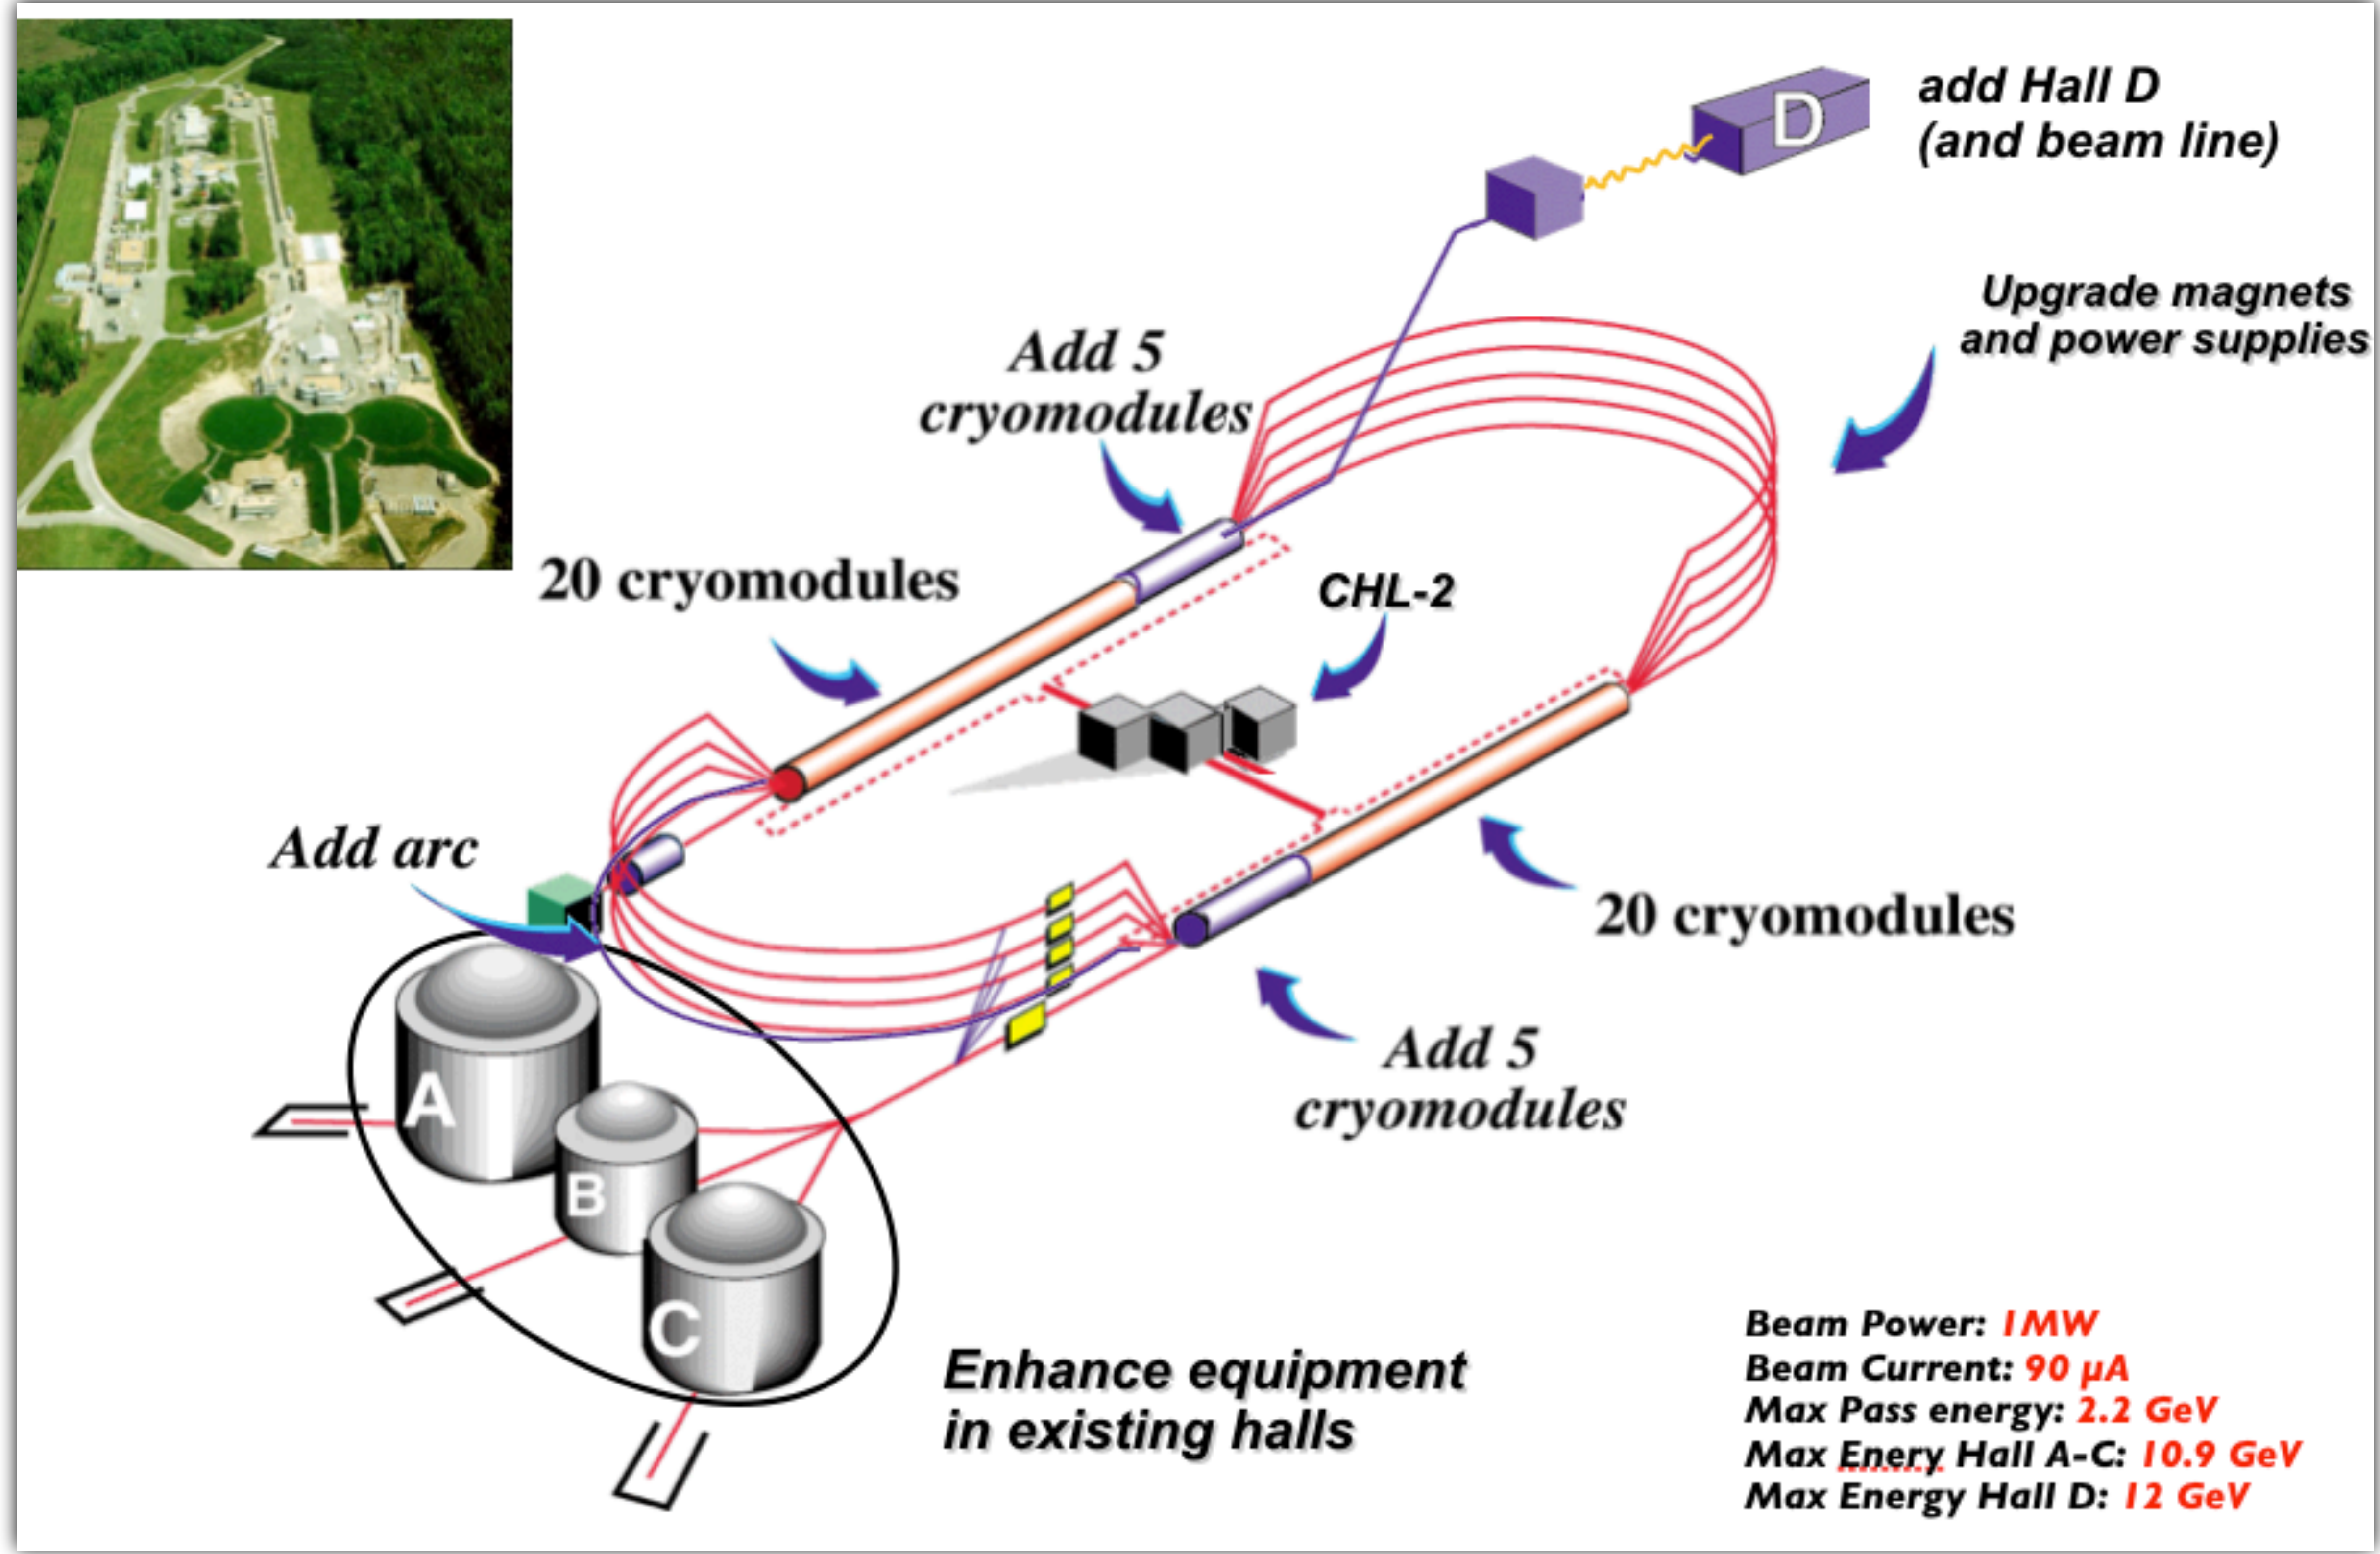

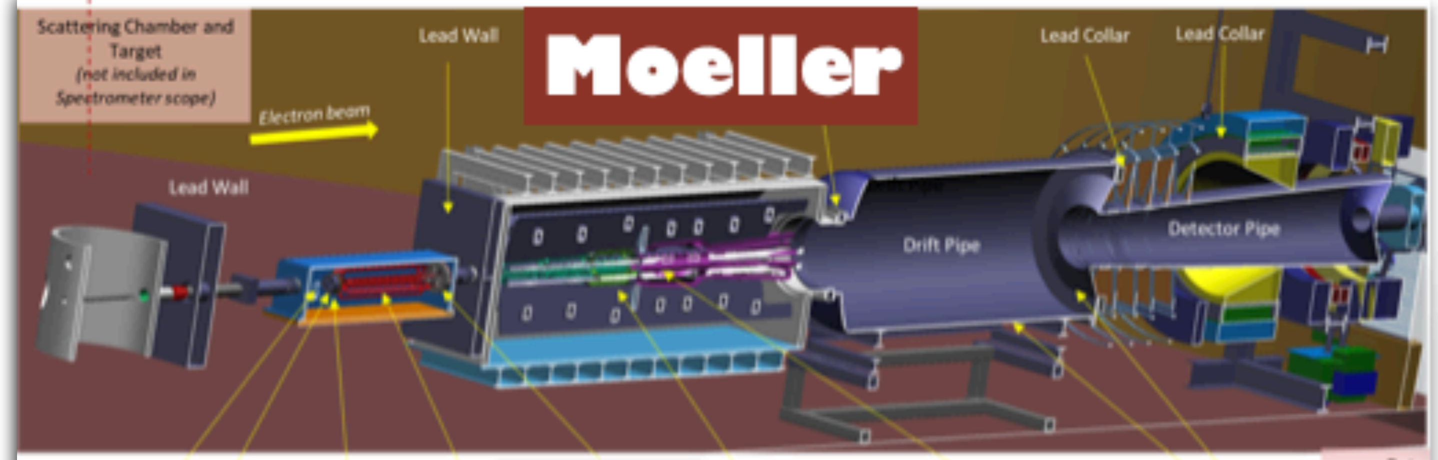

- Unique discovery space for new physics up to 38 TeV mass scale, with a purely leptonic probe
- CD-1 approved Dec 2020
- Expected to operate in FY26

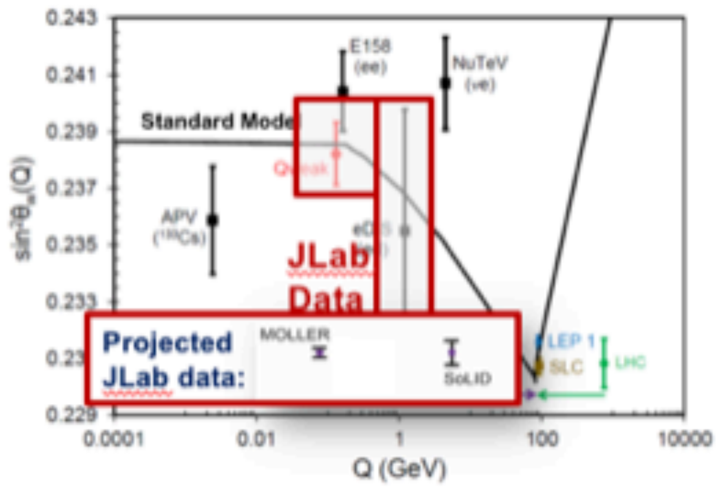

- **SO**lenoidal **L**arge **I**ntensity **D**evice – new multipurpose detector facility optimized for high luminosity (10^{37-39} cm^{-2} s^{-1}) and large acceptance

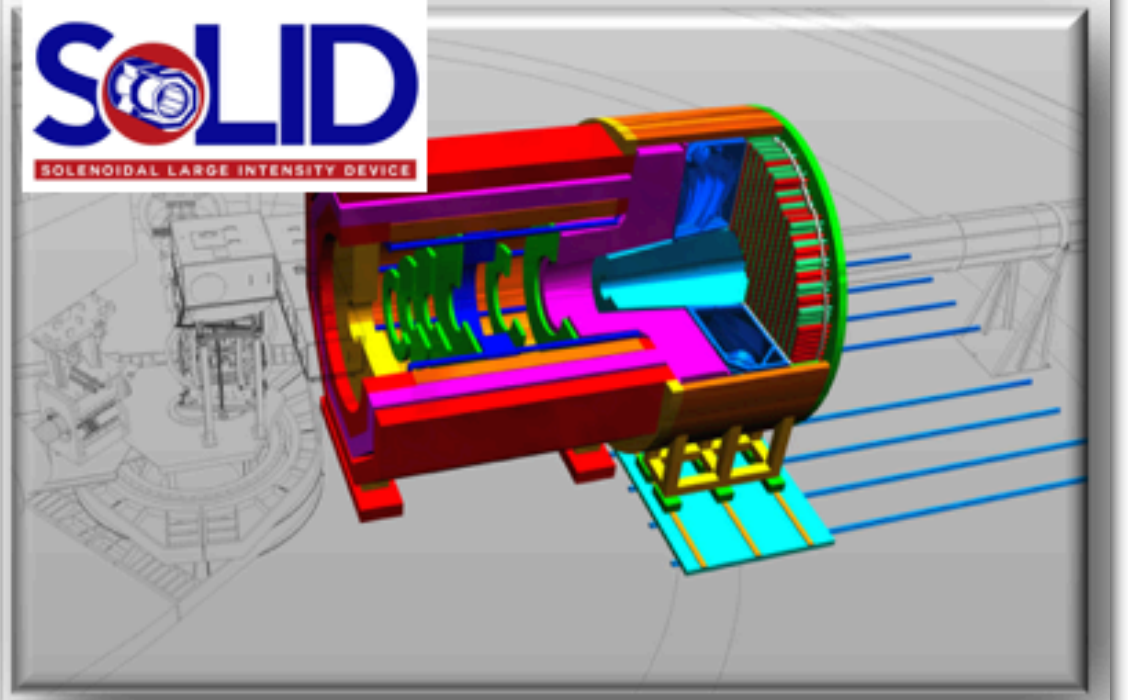

EIC

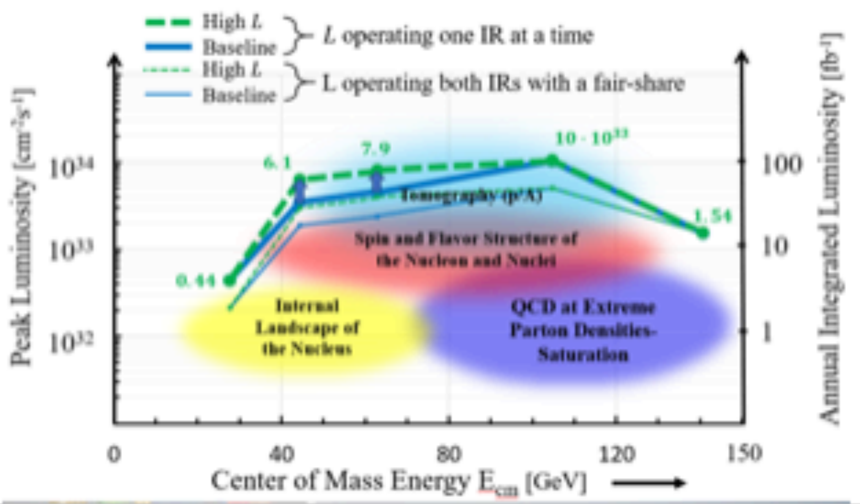

- Luminosity 100-1000 times that of HERA
- Polarized protons and light nuclear beams
- Nuclear beams of all A (p→U)
- Center mass variability with minimal loss of luminosity

- Large acceptance
- Frwrd/Bckw angles

- Precise vertexing
- HRes Tracking
- Excellent PID

DAQ in Hall-B: CLAS and CLAS12

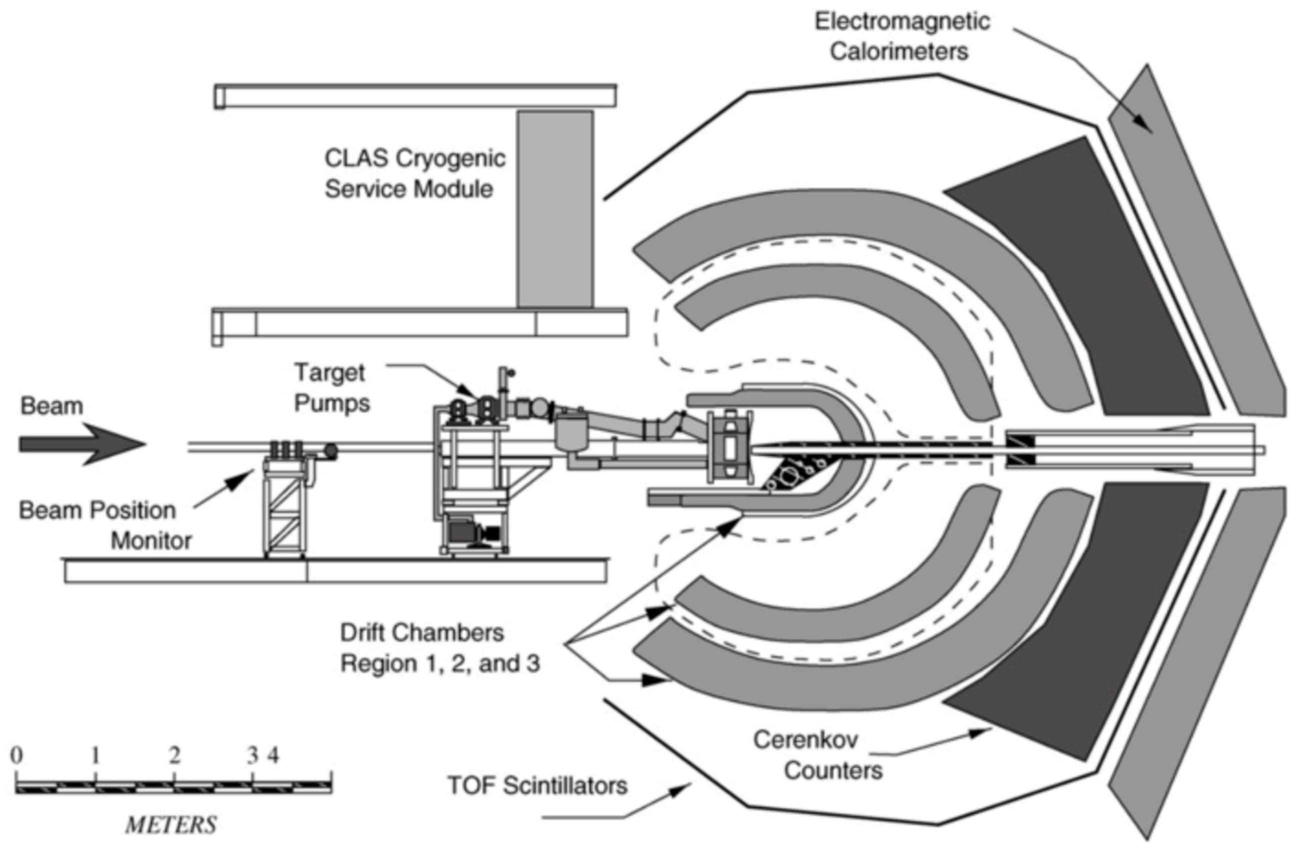

CLAS (1997-2012)

TRIGGERED system limited by:

- Available technology
- Low latency (~100ns) defined by FASTBUS ADCs
- Level-1 trigger: calorimeter/TOF/Cherenkov/StartCounter/Tagger (time coincidence of 'fast' detectors discriminators)
- Level-2 trigger: drift chamber (no tracking, just multiplicity)
- Hardware: - since 1997: JLAB-made; - since 2008: CAEN v1495's (Cyclone I - first experience with FPGAs)

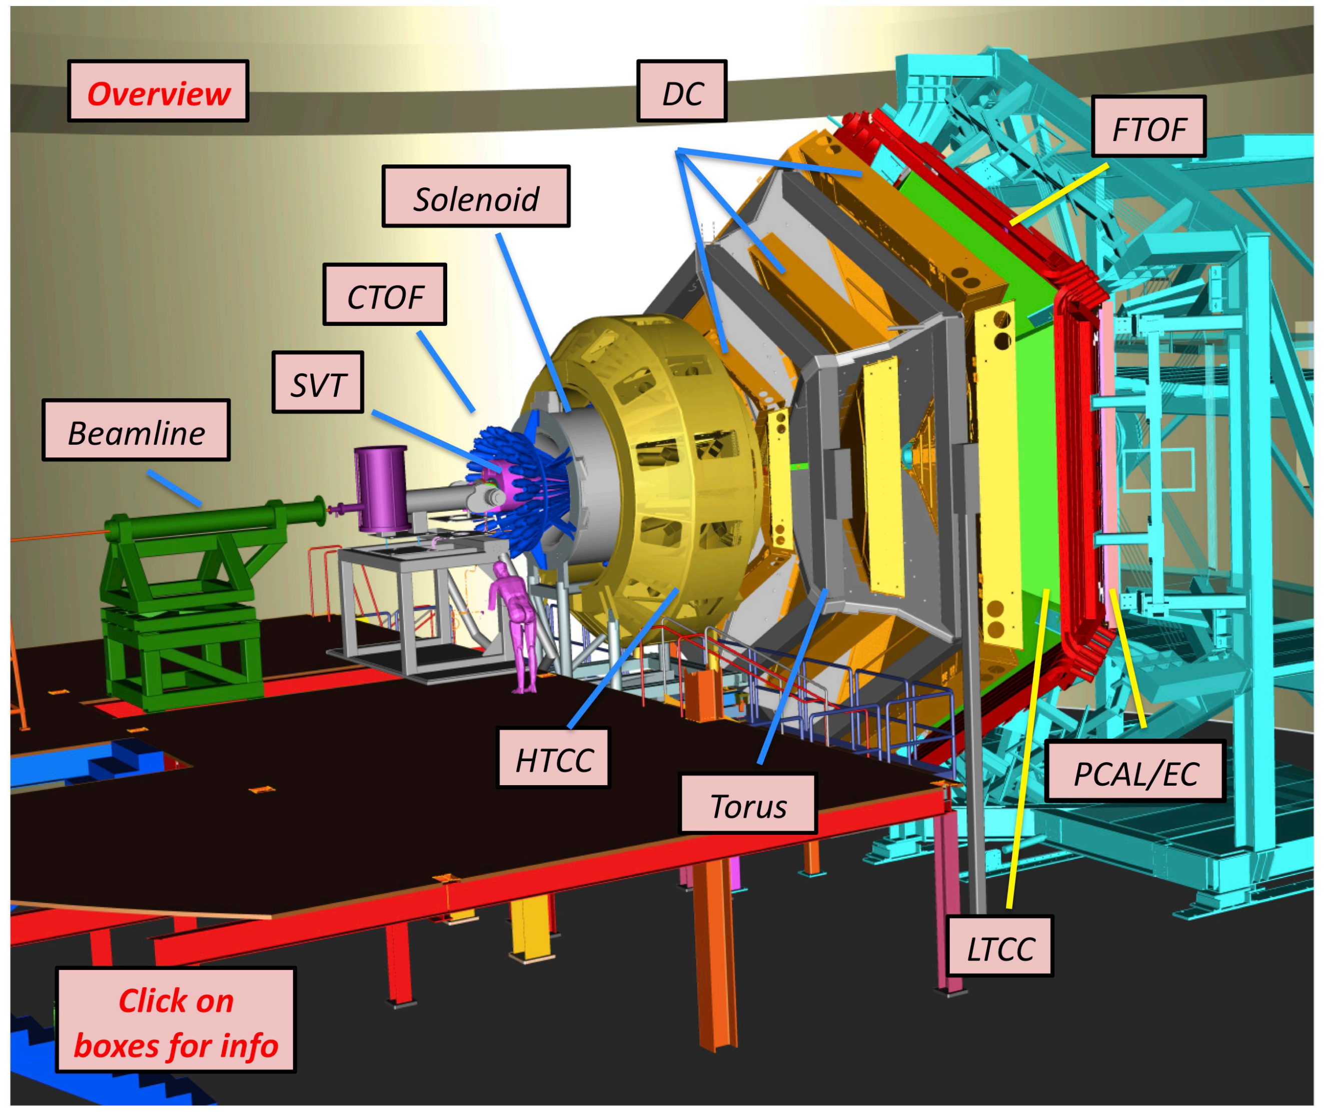

CLAS12 (2017 - present)

TRIGGERED system improved by:

- Advanced FPGA technology
- high latency (~8 us) defined by fADCs
- Level1 trigger: clusters in 3 calorimeters and hodoscope (FADC pulse integrals) + hits in 2 TOFs and 2 Cherenkovs (FADC pulse integrals) + tracks in Drift Chamber (discriminators/TDCs)
- Hardware: JLAB-made VTP (VXS Trigger Processor) boards (Virtex7 FPGA)
- Trigger algorithms: energy, position, and timing of the clusters in calorimeters; energy, position, and timing of TOFs and Cherenkov hits + drift chamber track finding (dictionary-based) + timing and position matching between clusters, hits and track

Credit: Sergey Boiarinov

Streaming RO in Hall-B

Credit: Sergey Boiarinov

CLAS12 (Future)

Streaming ReadOut

- Alternative use of the VXS Trigger Processor (VTP) board
- Up to 4 x 10 Gbit/s
- More than 50 crates (VME/VXS) can be reused

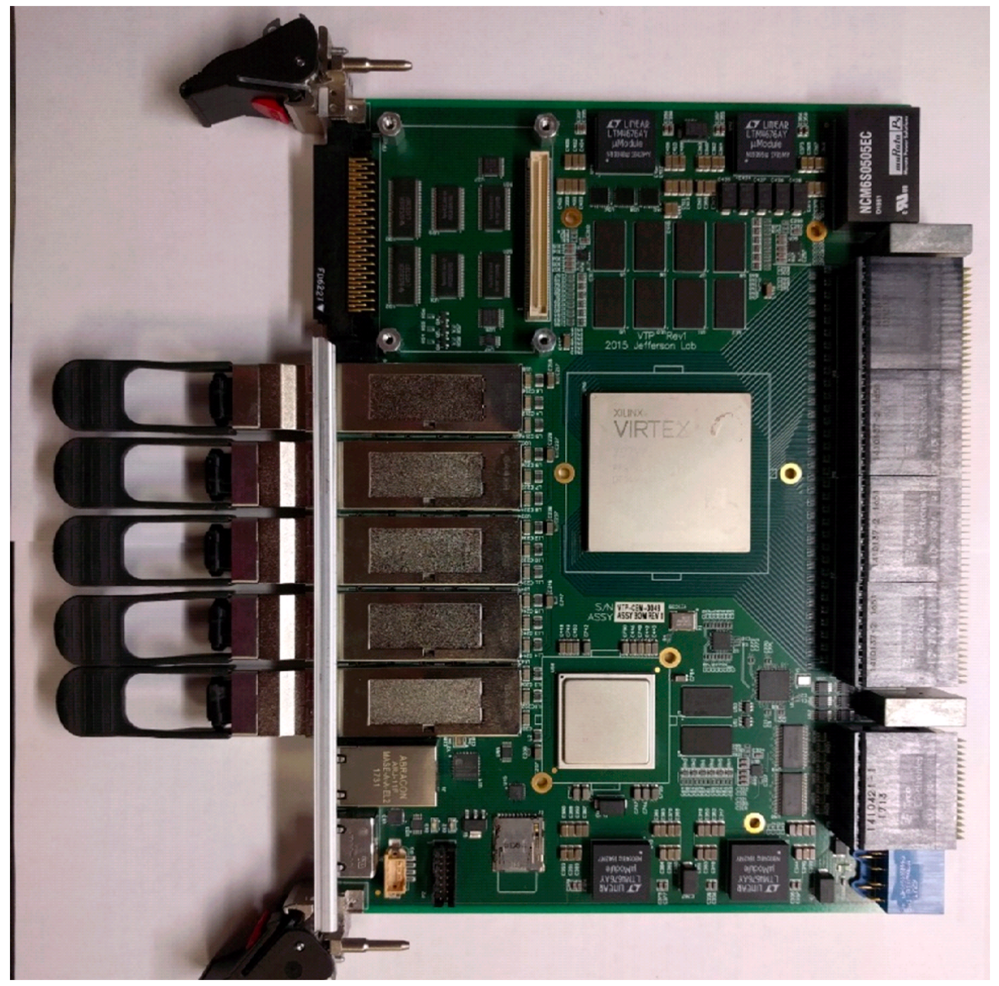

Status of SRO in Hall-B

- VTP firmware updates
- Beam tests with local resources (Hall-B counting house)
- Tests with remote resources (Computer Center) using Arista smart switch
- Planned test to stream ~half CLAS12 data to CC
- Final goal (~5y): x2-3 CLAS12 Luminosity (~50GB/s)

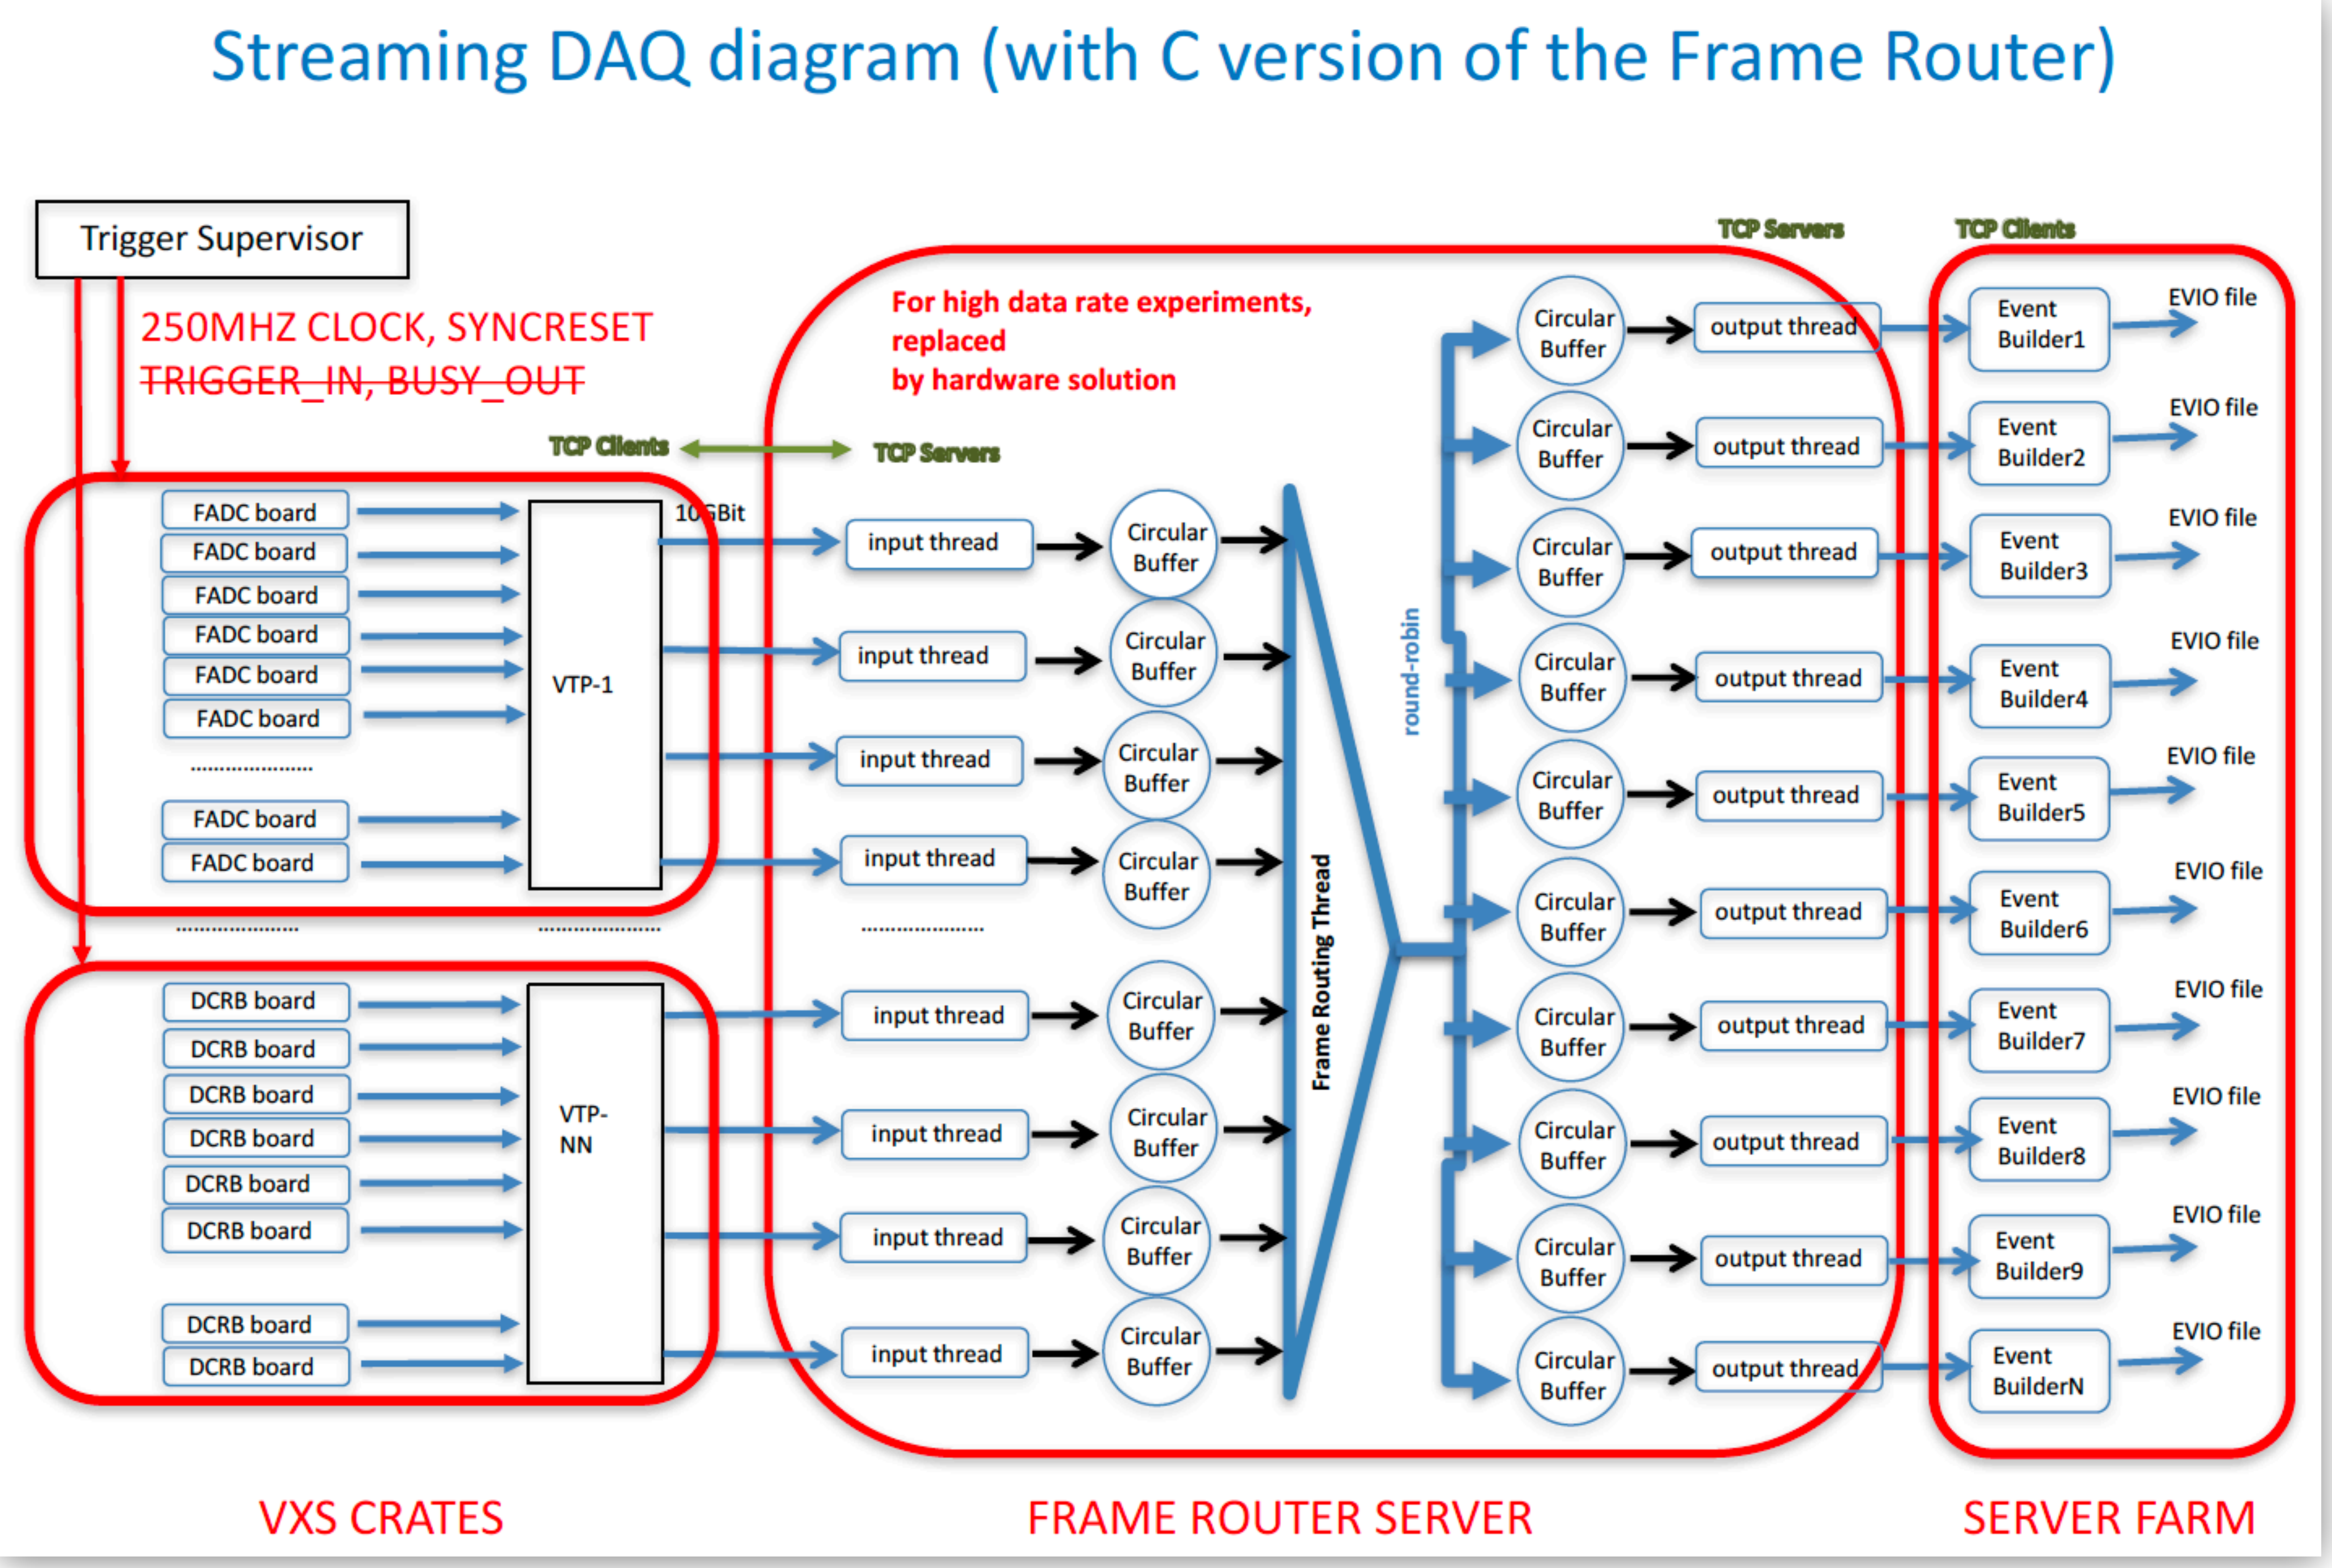

Streaming RO at Jlab: the proof of principle

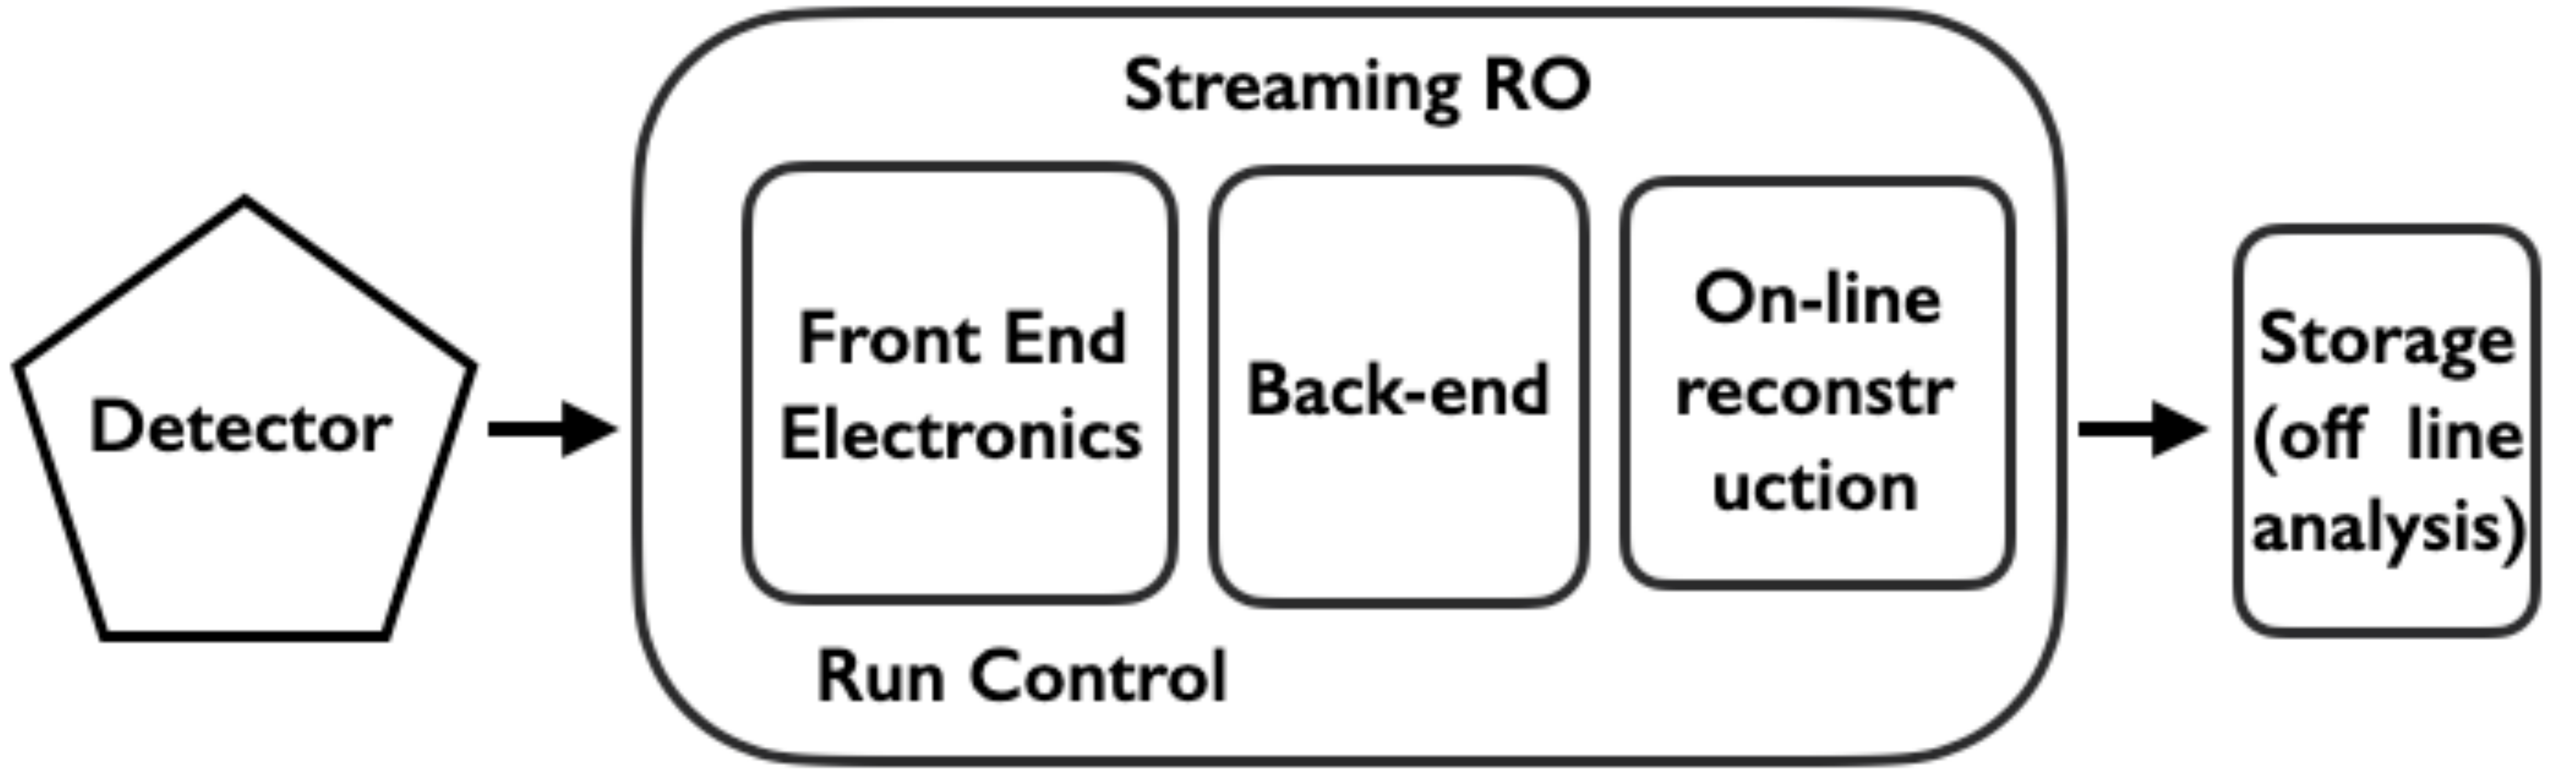

SRO concept validation
1) Assemble SRO components
2) Test SRO DAQ in lab
3) Test SRO DAQ on-beam

FrontEnd
D.Abbott, F.Ameli, C.Cuevas, P. Musico, B.Raydo

✳ **JLab fADC250 + VTP bord**
- JLab 250 MHz flash ADC digitizer currently used in many experiments
- Overcome VXS limitations (<24 Gb/s) using JLab VTP board (<40 Gb/s)
- Not optimised but reuse of existing boards: ready-to-go solution while waiting for fADC250.v2

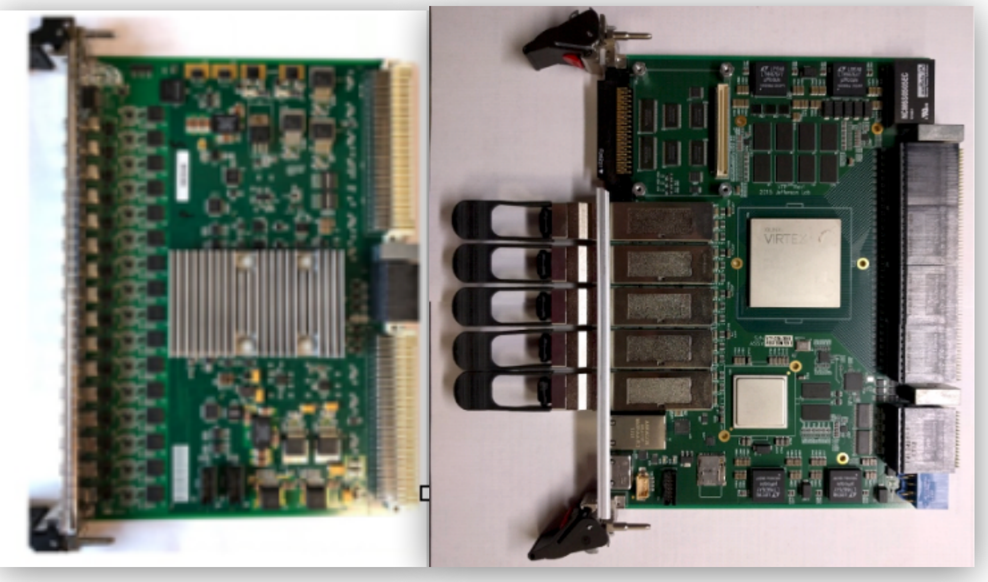

Jana2 + reconstruction
N.Brei, D.Lawrence,
M.Bondi', A.Celentano, C.Fanelli, S.Vallarino

✳ **JANA2 + TriDAS**
- Integration between On-line and off-line
- Real-time tagging/filtering data
- Offline algorithm development immediately available for use in Software Trigger
- Level 1 "minimum-bias": at least one crystal with E> 2 GeV
- Level 2 plugins (*tagging* and *filtering*)
  - "standard" FT-CAL clustering ($N_{cluster}\geq 1, 2, 3$)
  - cosmic tracking
  - AI clustering algorithm: at least two cluster in the FT-CAL

BackEnd
L.Cappelli, T.Chiarusi, F.Giacomini, C.Pellegrino

✳ **TRIggerless Data Acquisition System (TriDAS)**
- Developed for KM_3NET
- Installed on Hall-B DAQ cluster
- Multi CPUs, rate up to 20-30 MHz

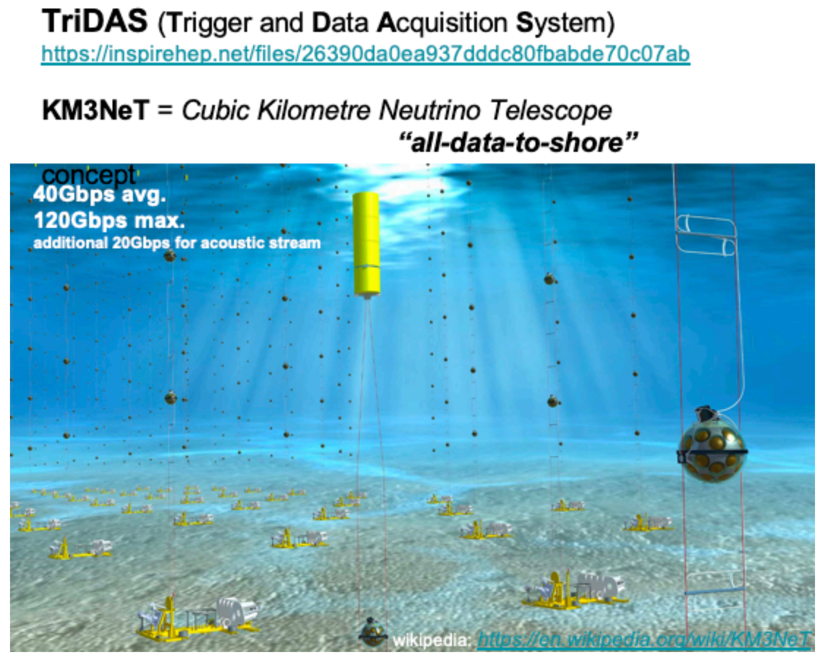

Cebaf Online Data Acquisition (CODA)
D.Abbott, S.Boyarinov, B.Raydo, G.Heyes

- Originally designed for trigger-based readout systems
- Controllers (ROCs) and VXS Trigger Boards (VTPs)
- The Trigger Supervisor (TS) synchronizes components using clock, sync, trigger and busy signals.-time tagging/filtering data
- CODA adapeted to the SRO
  - Replaced EB to use timestamp)
  - ROC communication via VTP (not VXS bus)

Streaming RO at Jlab: the proof of principle

Streaming Readout for Next Generation Electron Scattering Experiments
https://link.springer.com/article/10.1140/epjp/s13360-022-03146-z

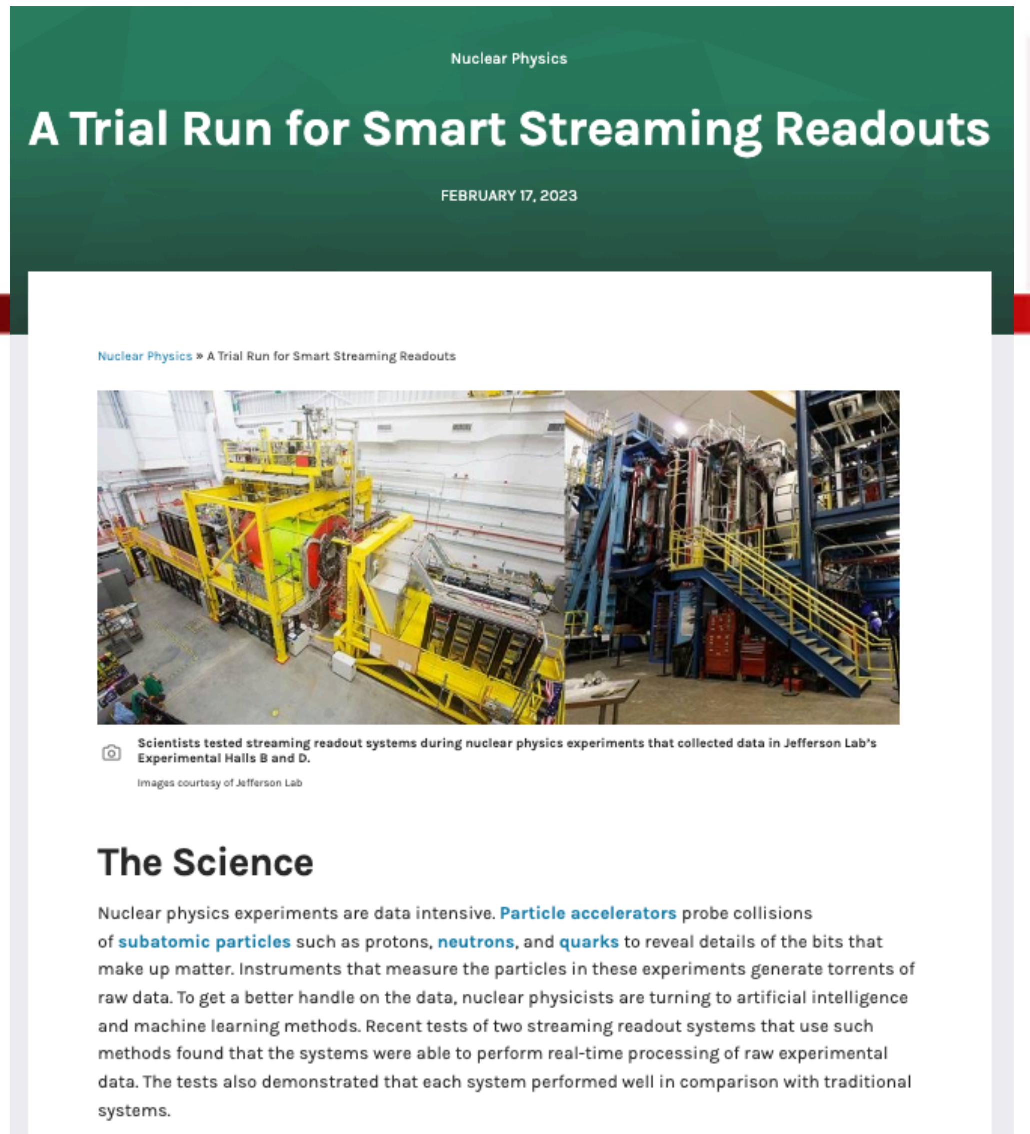
Nuclear Physics

A Trial Run for Smart Streaming Readouts

FEBRUARY 17, 2023

Nuclear Physics » A Trial Run for Smart Streaming Readouts

Scientists tested streaming readout systems during nuclear physics experiments that collected data in Jefferson Lab's Experimental Halls B and D.

Images courtesy of Jefferson Lab

The Science

Nuclear physics experiments are data intensive. Particle accelerators probe collisions of subatomic particles such as protons, neutrons, and quarks to reveal details of the bits that make up matter. Instruments that measure the particles in these experiments generate torrents of raw data. To get a better handle on the data, nuclear physicists are turning to artificial intelligence and machine learning methods. Recent tests of two streaming readout systems that use such methods found that the systems were able to perform real-time processing of raw experimental data. The tests also demonstrated that each system performed well in comparison with traditional systems.

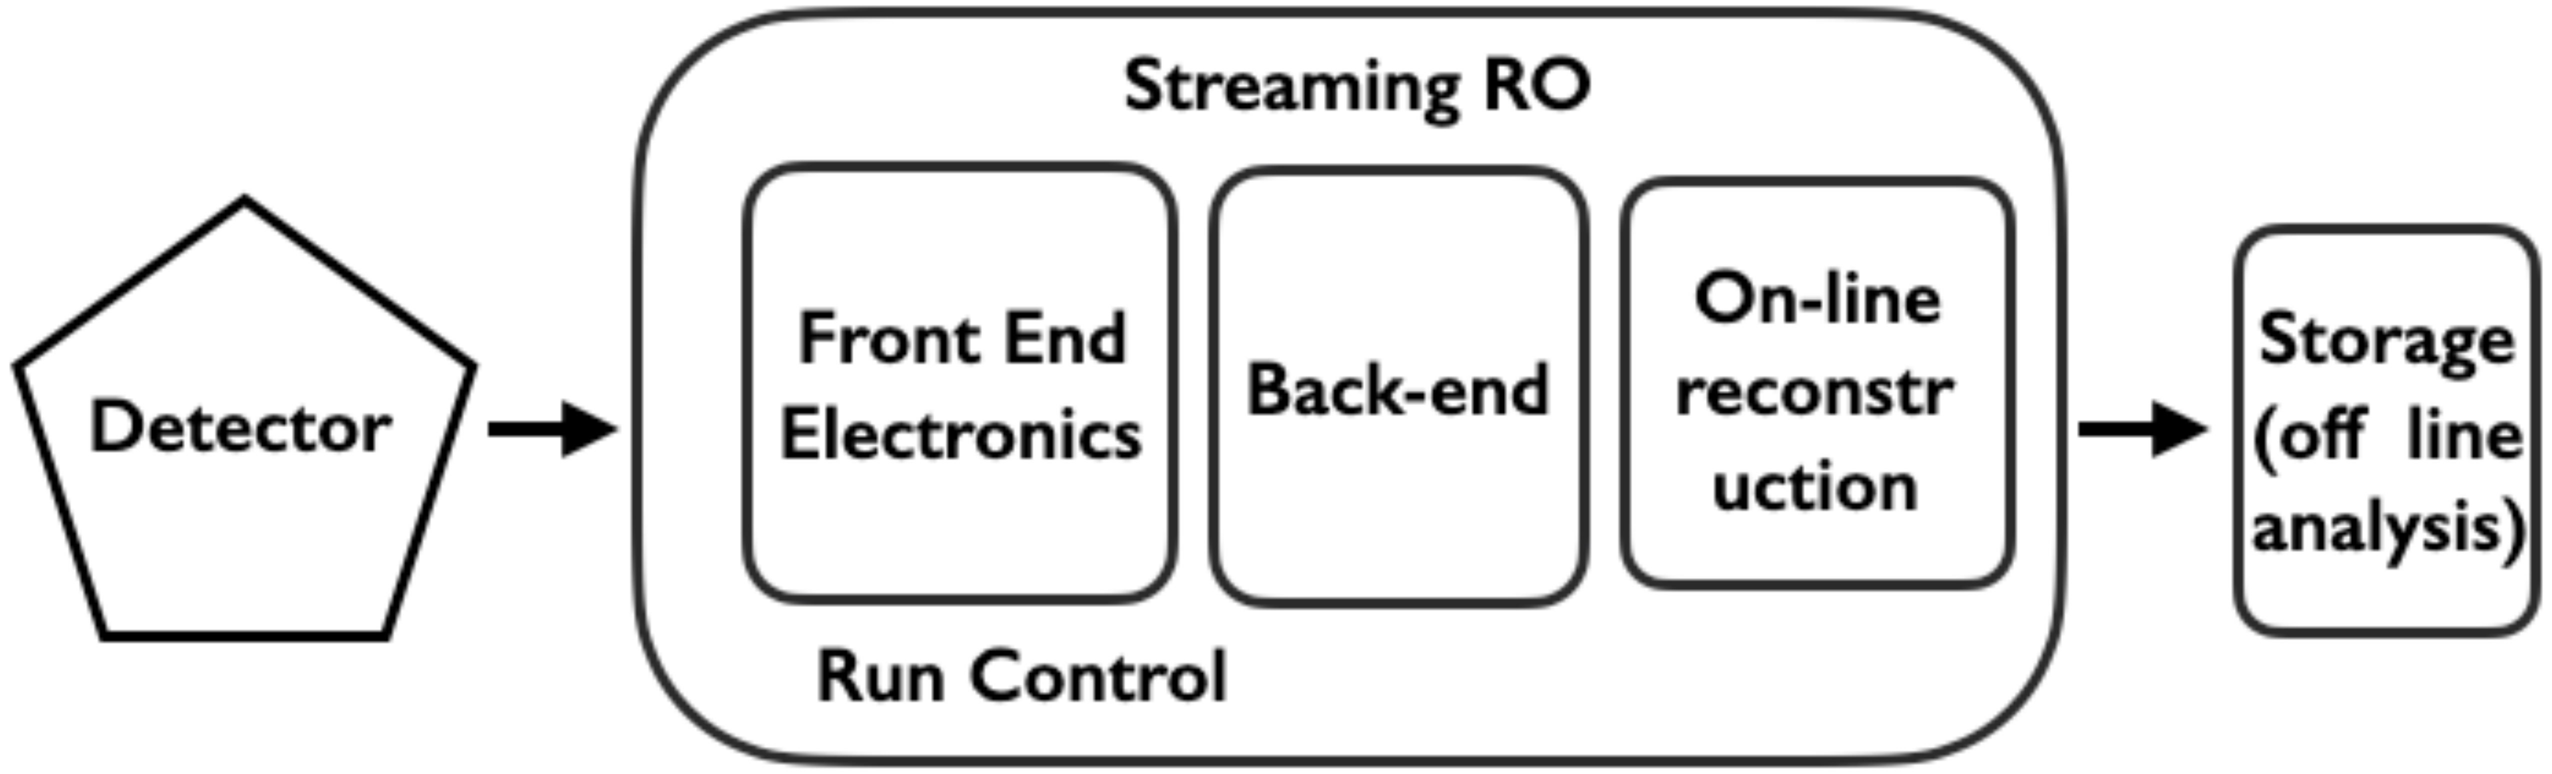

SRO concept validation
1) Assemble SRO components
2) Test SRO DAQ in lab
3) Test SRO DAQ on-beam

JLab SRO validation

✳ CLAS12 Forward Tagger
- Complete system that include calorimetry, PiD, Traking in a simpler (than CLAS12) set up
- FT-ECAL: 332 PbWO crystals, APD readout
- FT-HODO: 224 plastic scintillator tiles, SiPM readout
- FT-TRK: ~3000 channels, MicroMegas
- fADC250 digitizers + DREAMs for MM

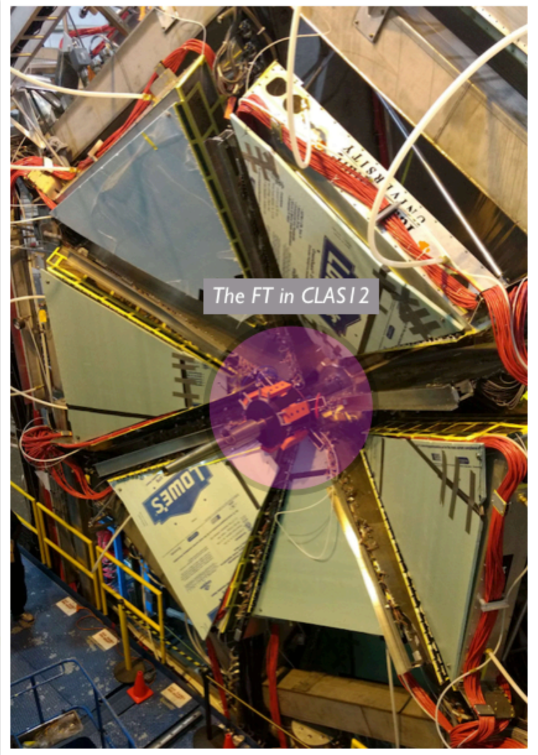

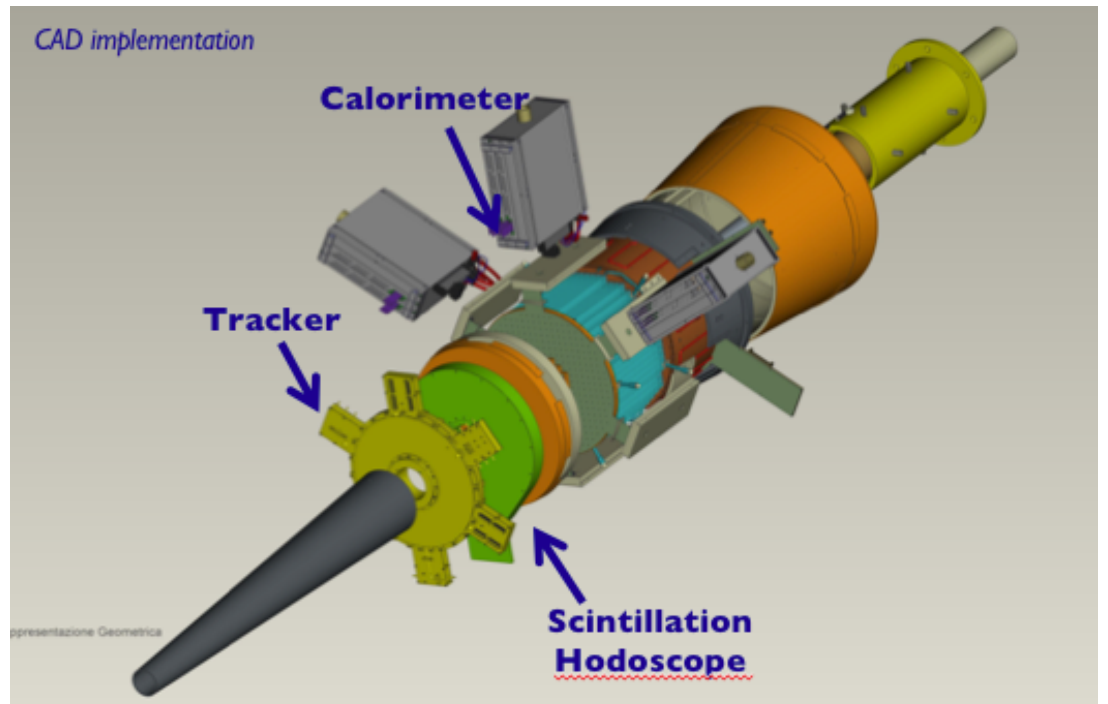

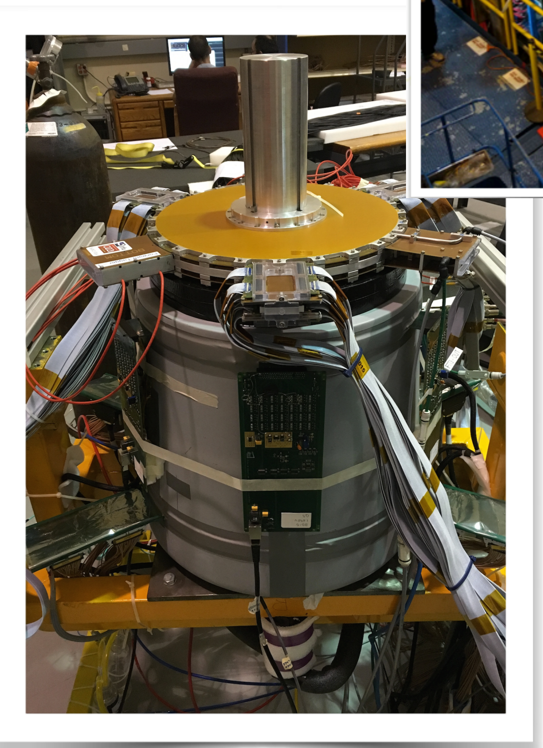

✳ CLAS12 Forward Tagger
- Inclusive pi0 electroproduction
- Two gammas detected into FT-CAL
- EM clusters identification, anti coincidence with FT-Hc
- Self-calibration reaction (pi0 mass)

- **On-beam tests:**
  - 10.4 GeV e- beam on thin Pb/Al target
  - Inclusive pi0 production
    - e + Pb/Al -> X$e\pi^0$ -> (X)e$\gamma\gamma$
  - Two gammas detected in FT-CAL

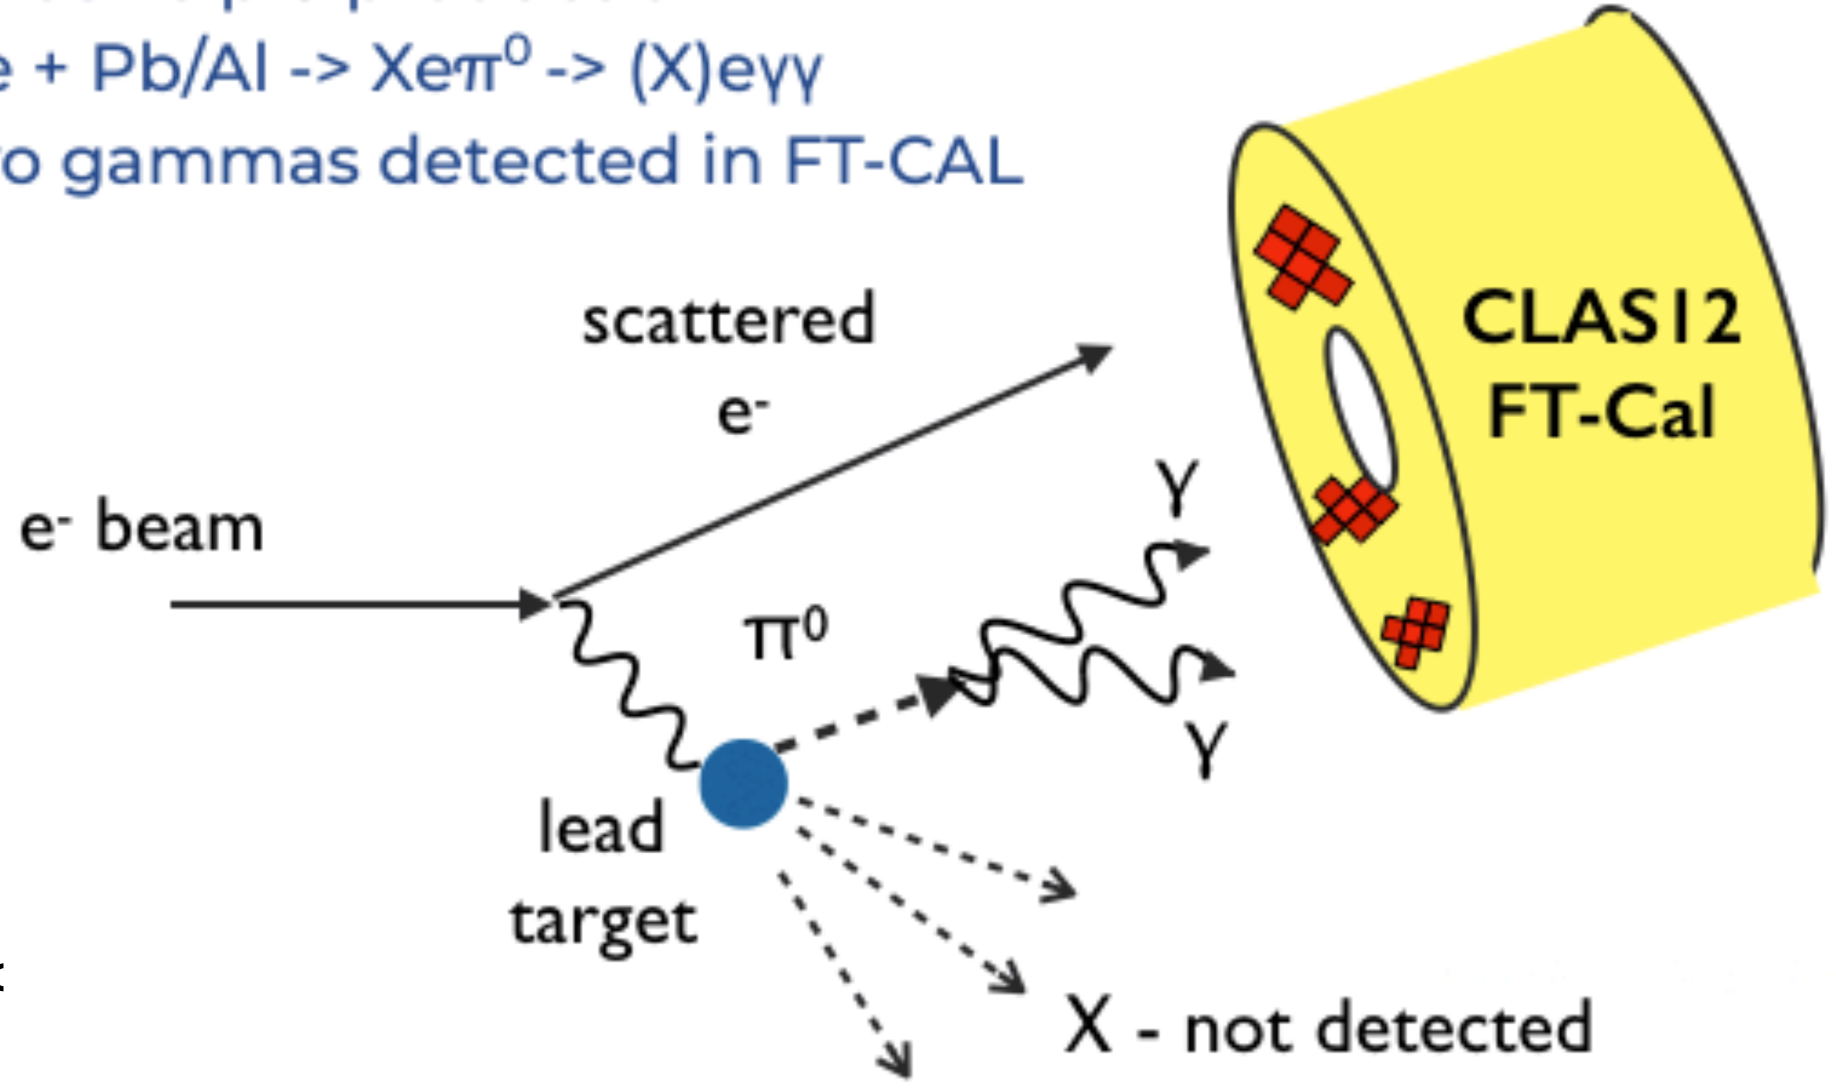

Streaming RO at Jlab: the proof of principle

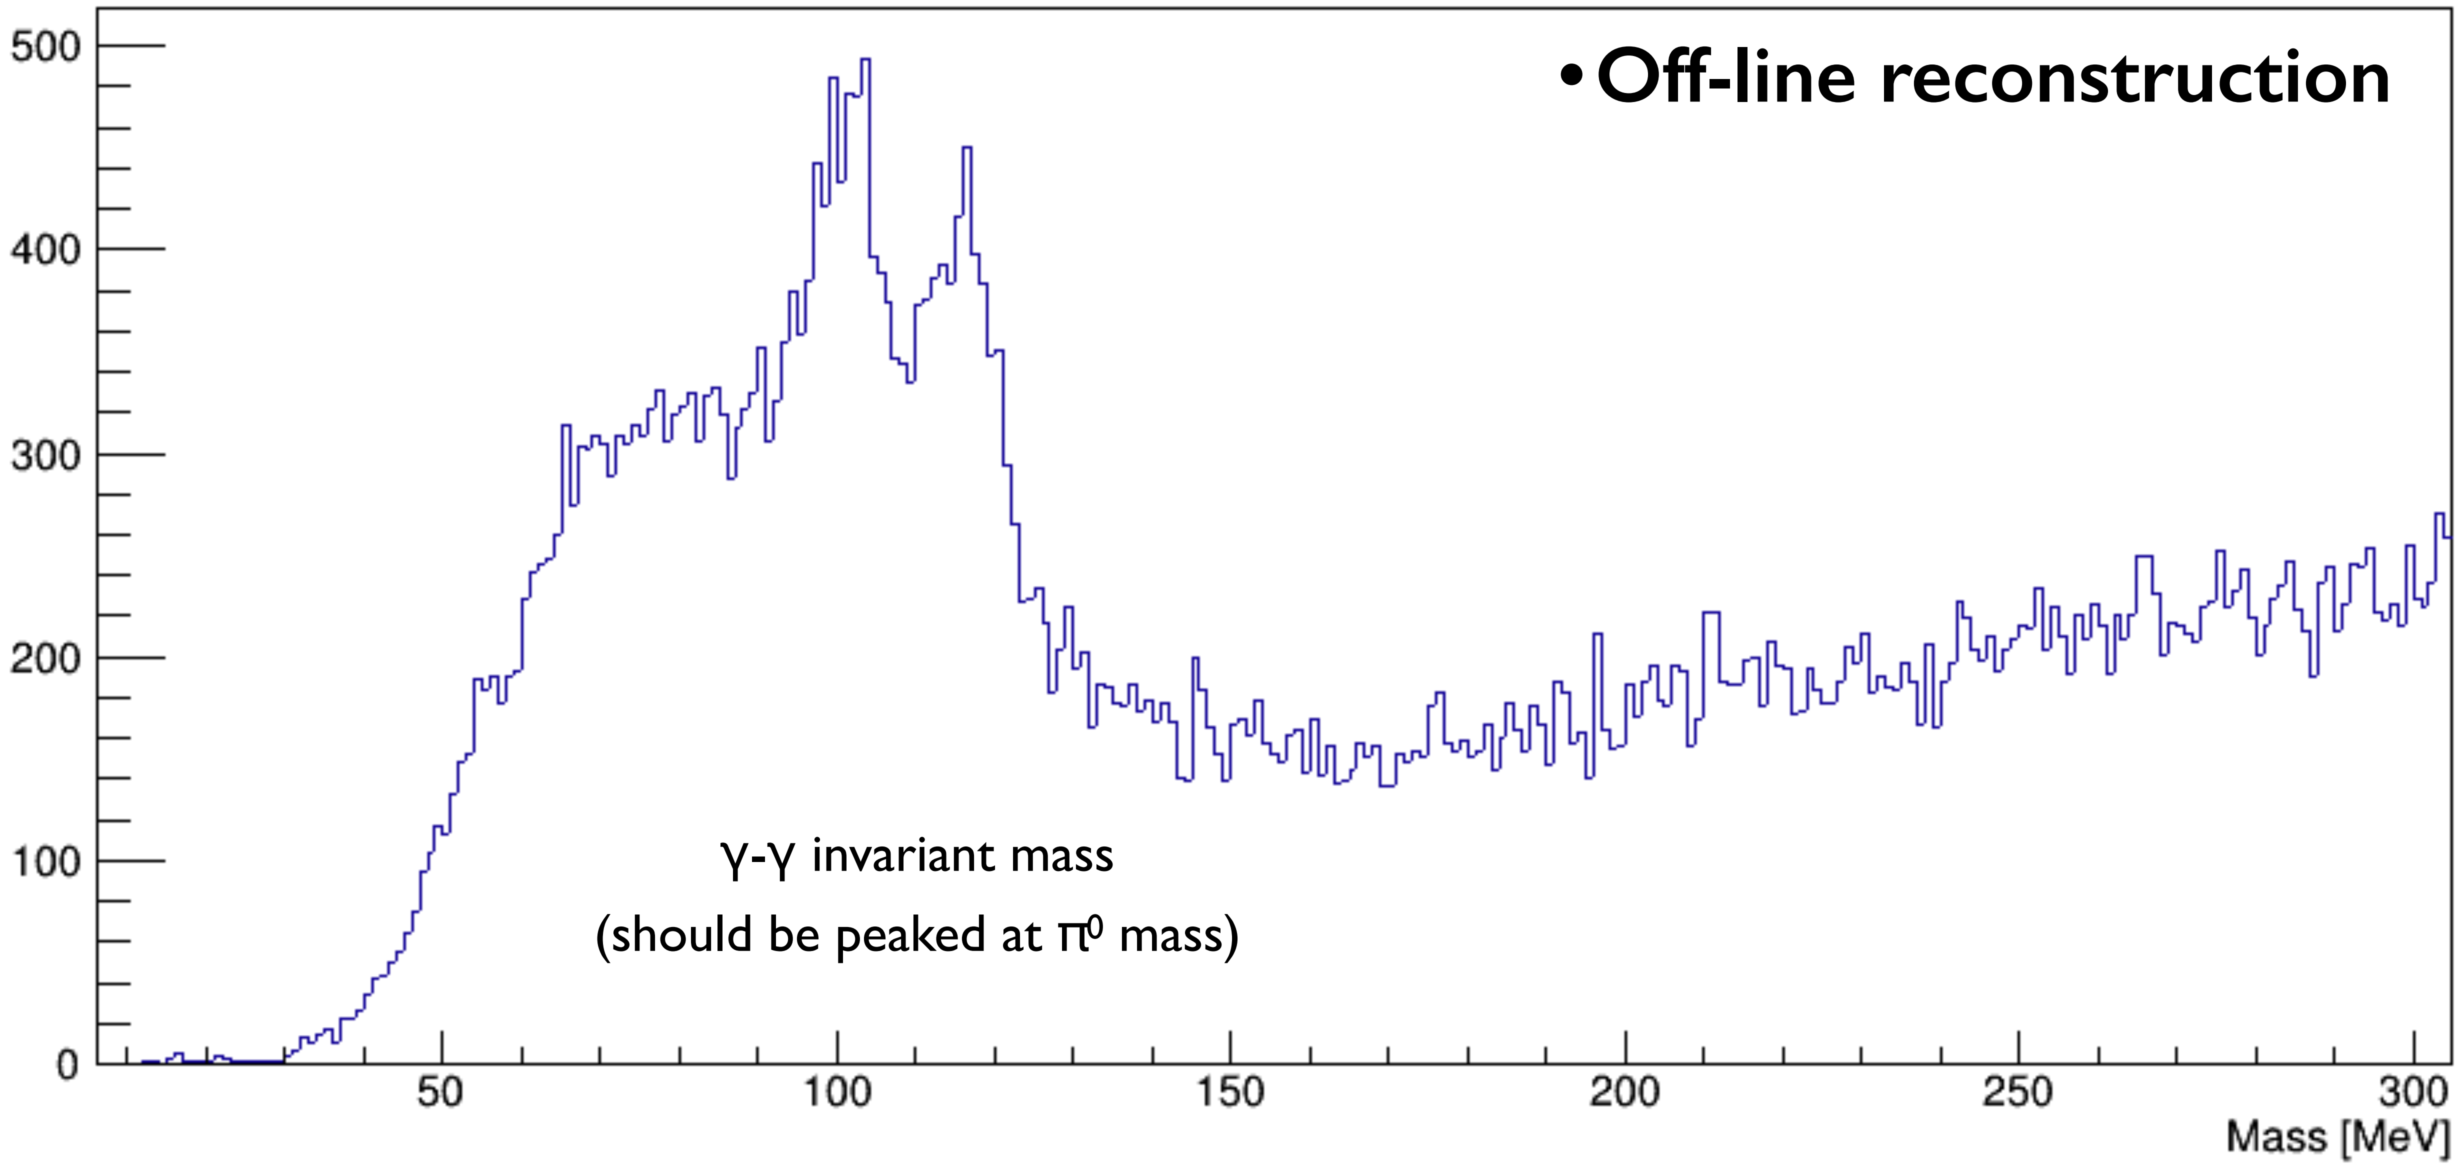

• Two pi0 peaks corresponding to two vertices (and a wrong assumption on the vertex position)

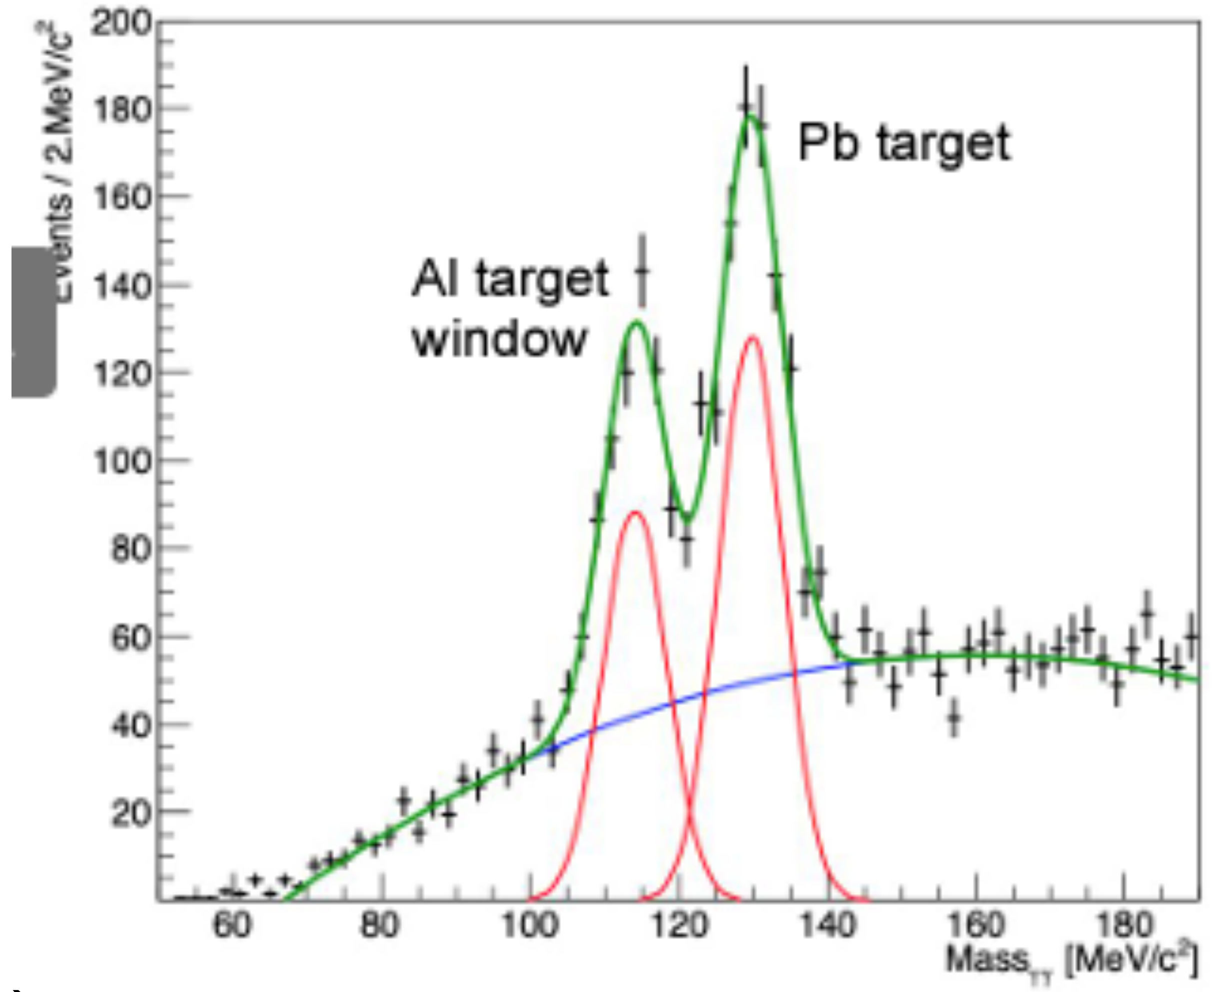

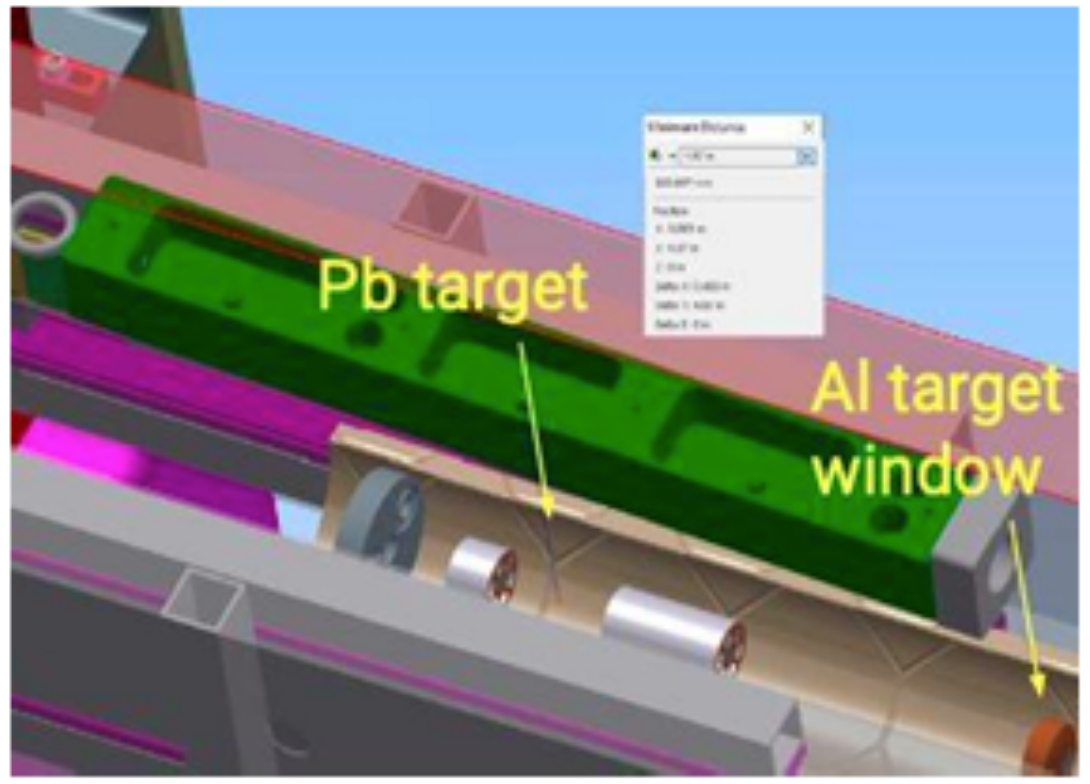

AI to analyze data in real time, and extract features (e.g. number of peaks and position)

- AI clustering inspired by *Hierarchical Density-Based Spatial Clustering of Applications with Noise* (HDBSCAN)
  - It is not cut-based
  - it is able to cope with a large number of hits
- Compared γγ-invariant mass spectrum obtained utilizing both the standard and the HDBSCAN clustering algorithm
  - AI significantly improves signal-to-background ratio in the π0 region
  - A longer runtime of ~30% relative to the standard clustering algorithm
- AI clustering approach promising alternative to traditional cut-based approaches

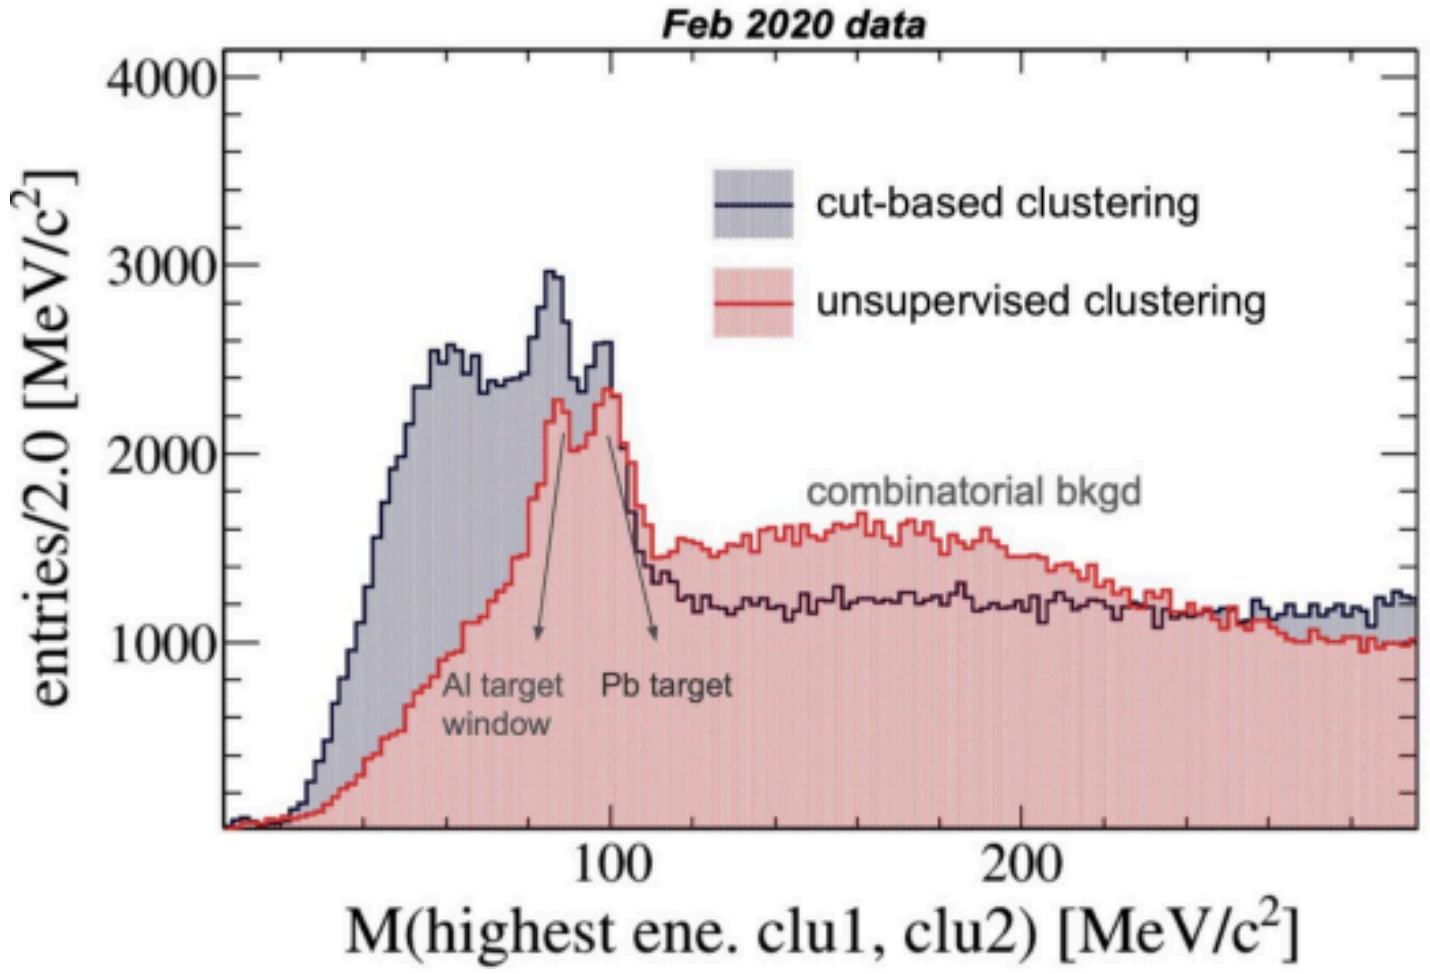

SRO at Jefferson Lab

Streaming Readout Real-Time Development and Testing Platform

PI: David Lawrence (JLab CST)
Jefferson Lab Lab-Directed R&D project

Contributors:	Vardan Gyurjyan, CST
	Xinxin "Cissie" Mei, CST
Advisors/consultants:	Marco Battaglieri, INFN
	Markus Diefenthaler, ENP
	Sergey Furletov, ENP
	Sergey Boyarinov, ENP

1. Develop software platform capable of configuring and launching various existing software and hardware **SRO** components as a **complete chain**
2. Develop **monitoring** system capable of monitoring performance of all components specifically for identifying **bandwidth** and **compute bottlenecks**
3. Develop proxy components that can effectively **simulate** performance of specific **hardware** or **software components** that do not currently exist
4. Develop **multi-stream software source** that can take existing experimental or simulated data and broadcast it into the system with **time structure** and **stream count** that mimics a running experiment.
5. Configure a system comparable in size and bandwidth to a future JLab experiment (e.g. SoLID) which includes a **400Gbps** transfer requirement from the counting house to the Computer Center, at least one **FPGA** component and one **GPU** component for at-scale testing.
6. Identify potential issues relevant to a future **HPDF** in receiving and processing SRO data, including from **remote**, non-JLab **experiments**.

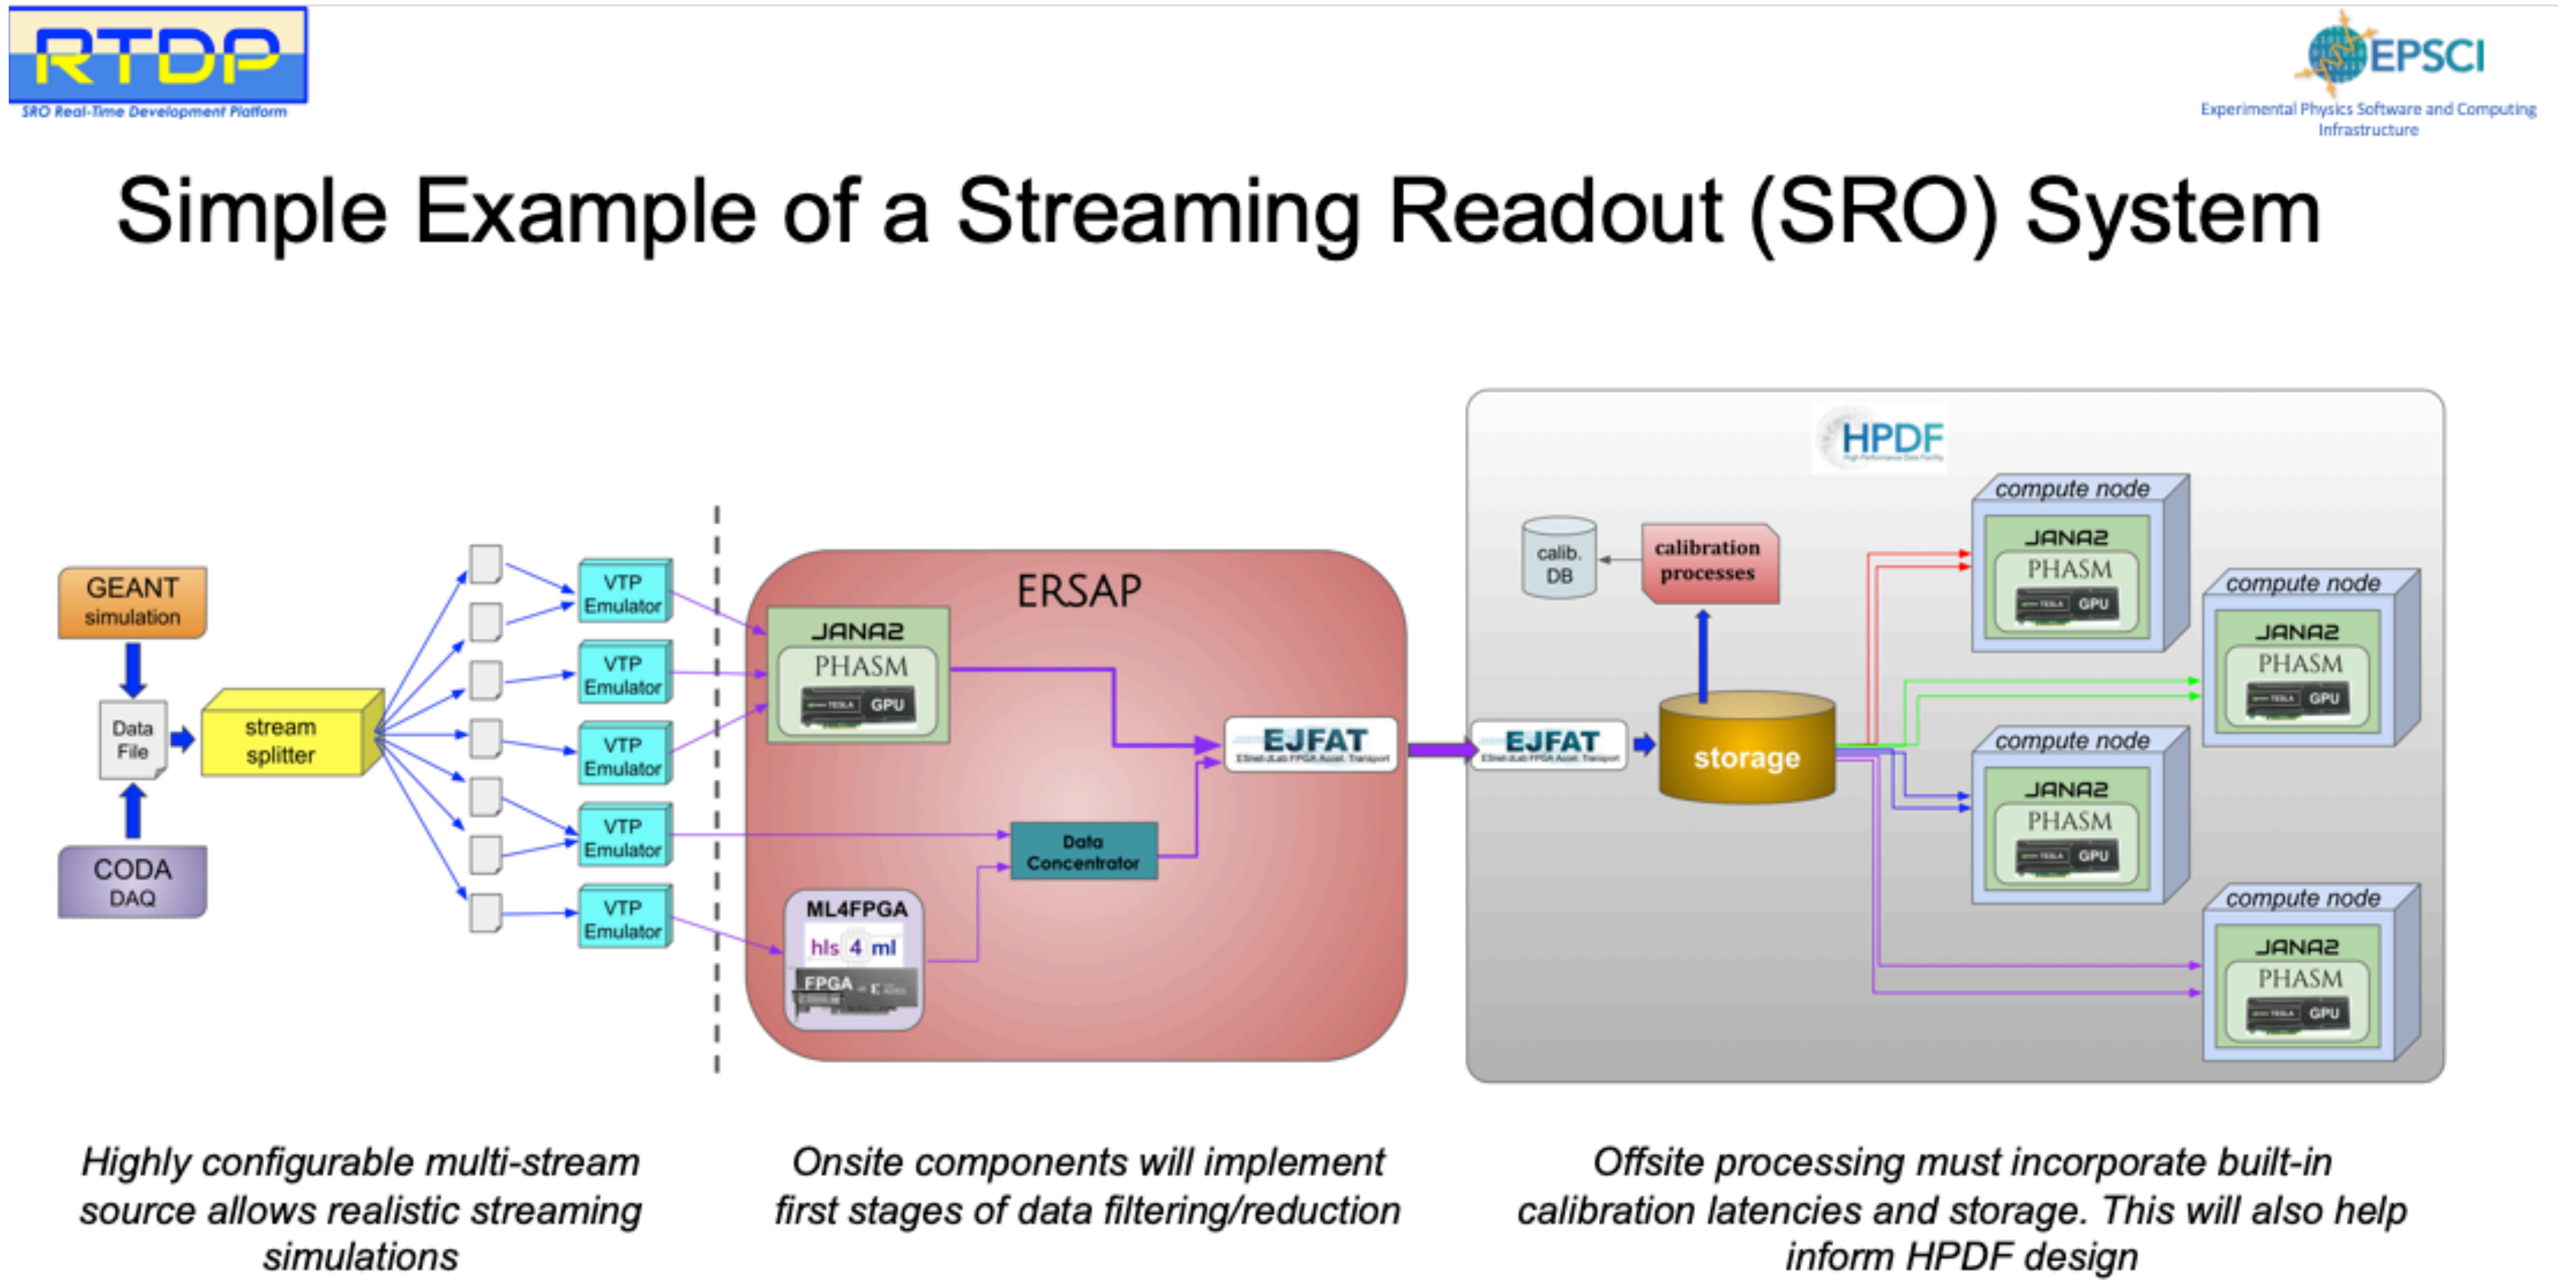

Credit: David Lawrence

Streaming RO at Jlab: the current effort

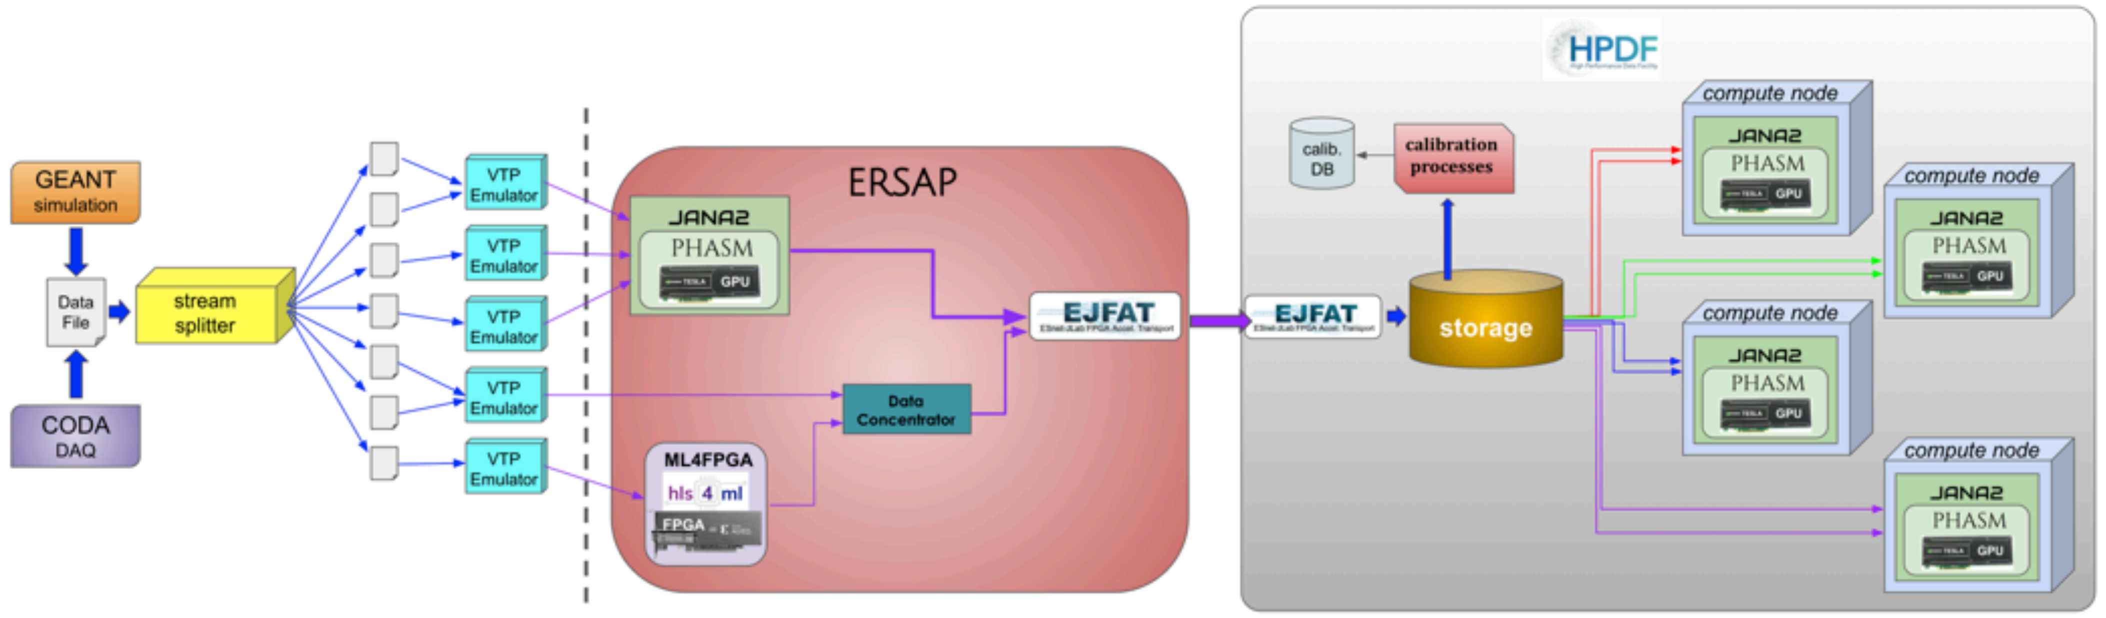

Highly configurable multi-stream source allows realistic streaming simulations

Onsite components will implement first stages of data filtering/reduction

Offsite processing must incorporate built-in calibration latencies and storage. This will also help inform HPDF design

- **Integration of different components in an optimised SRO framework**

ERSAP

- Reactive, event-driven data-stream processing framework that implements micro-services architecture
- Provides basic stream handling services (stream aggregators, stream splitters, etc.)
- Adopts design choices and lessons learned from TRIDAS, JANA, CODA and CLARA

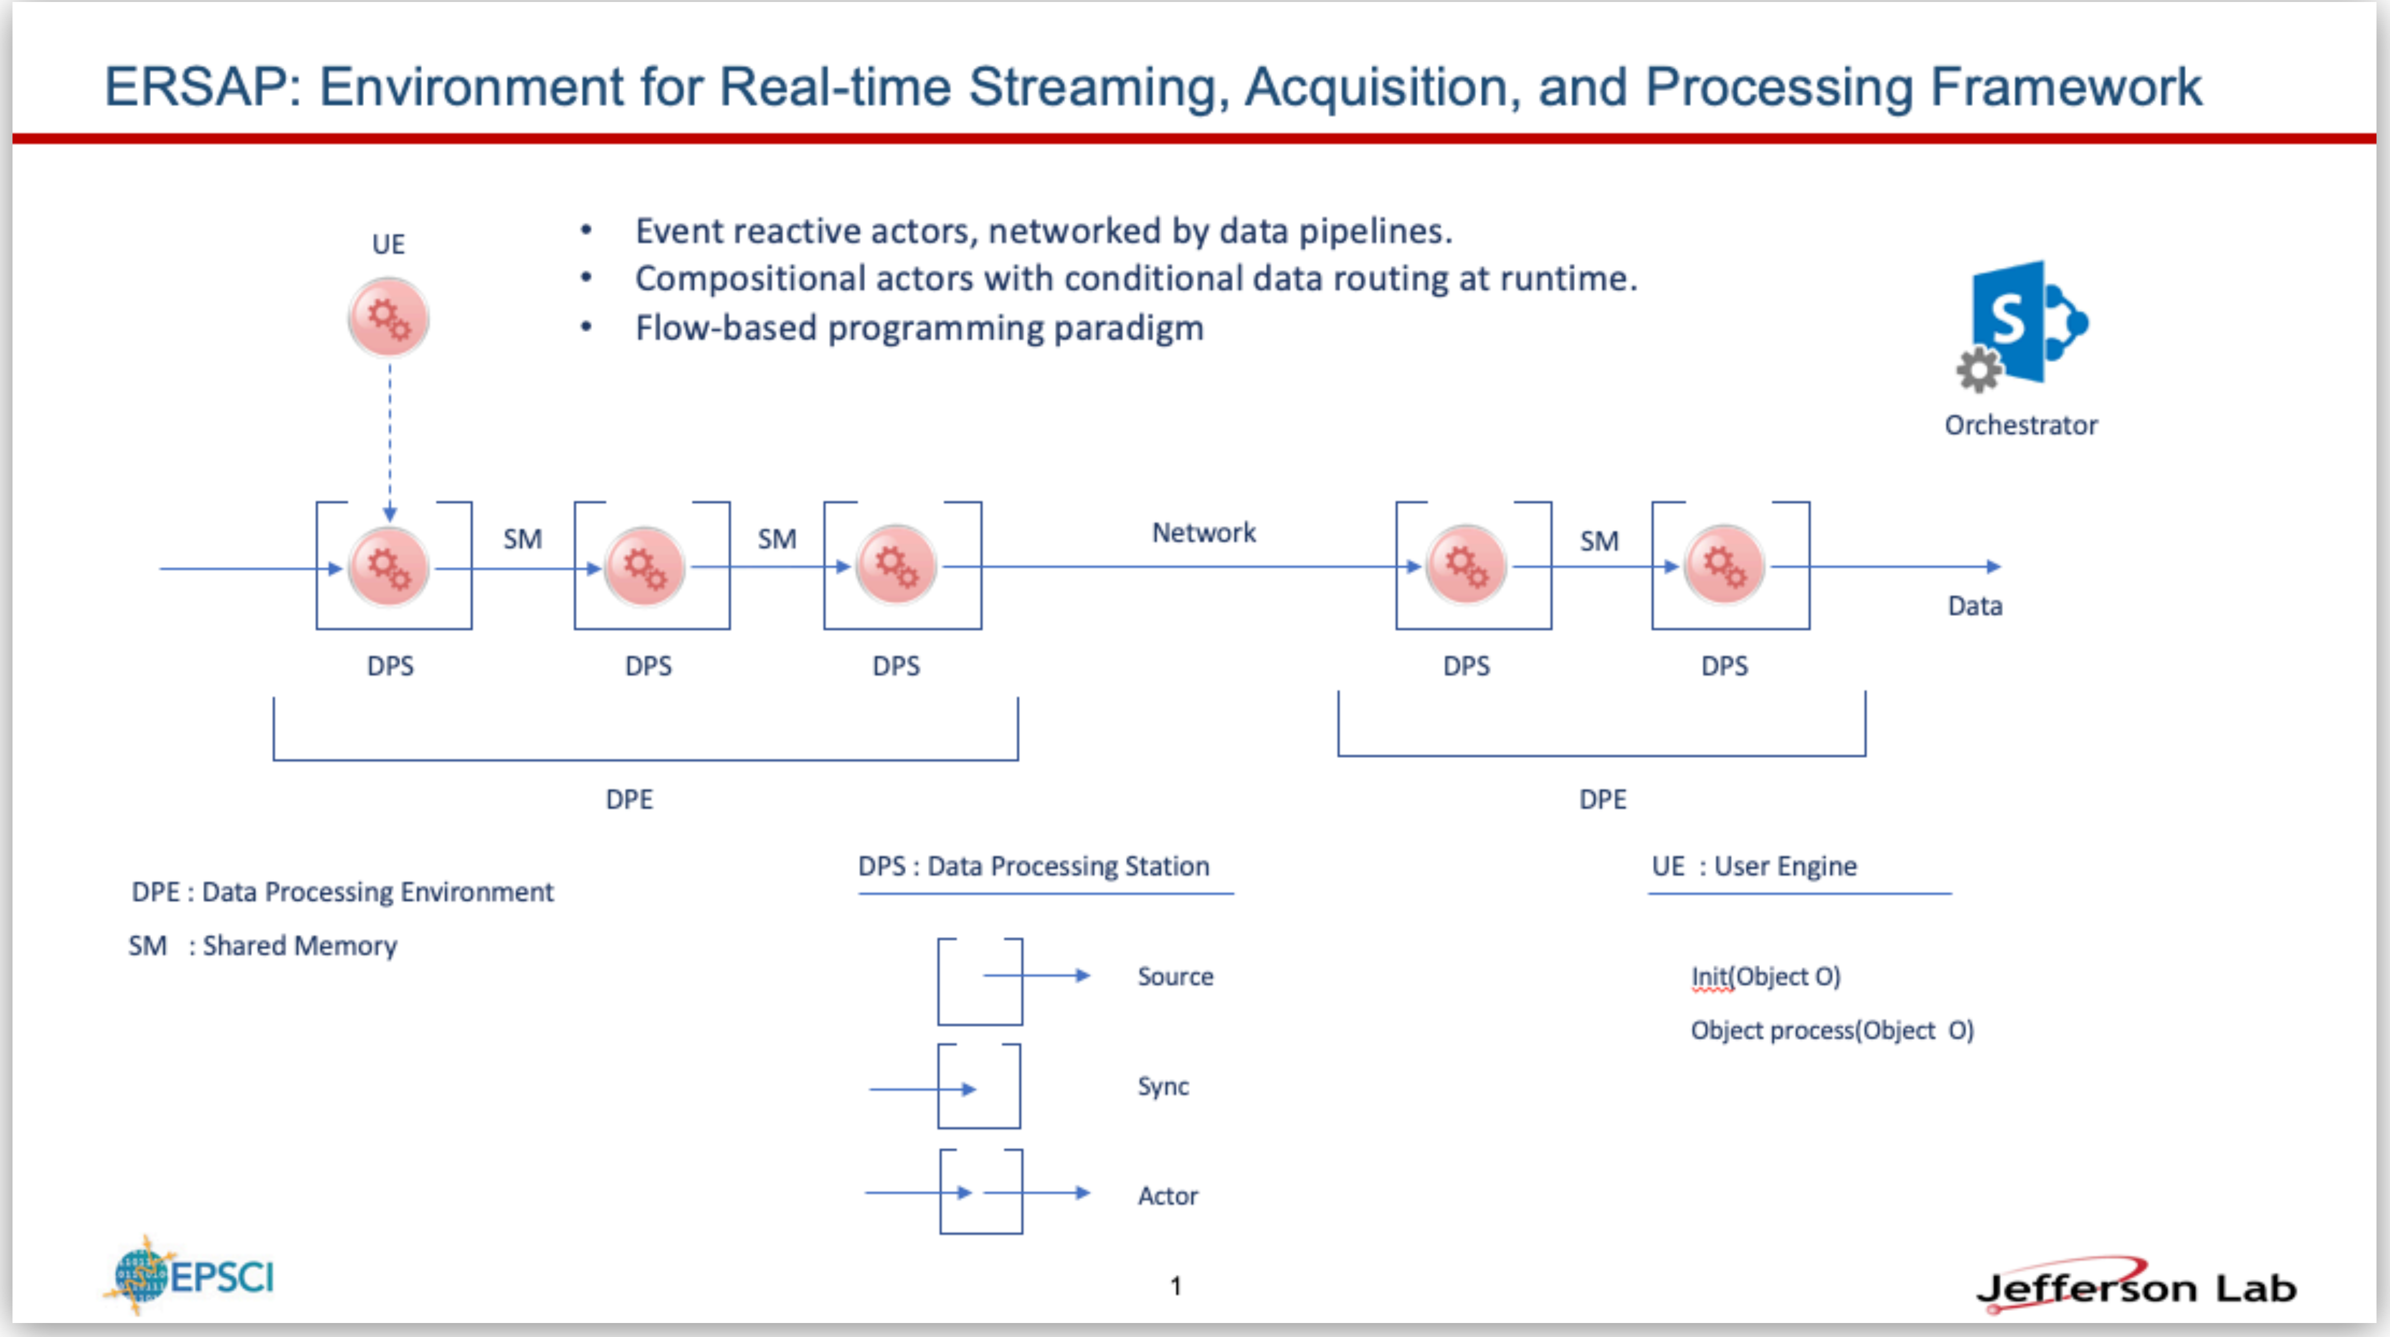

- **Development of SRO DAQ pipeline tools**

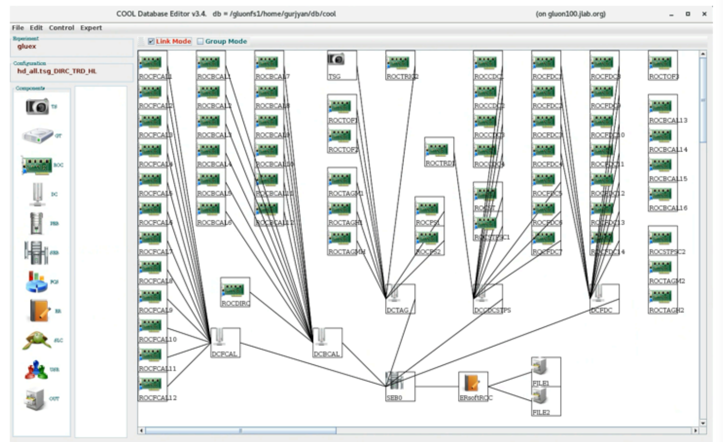

- **Development of JLAB CLAS12 Remote data-stream processing over a distance**

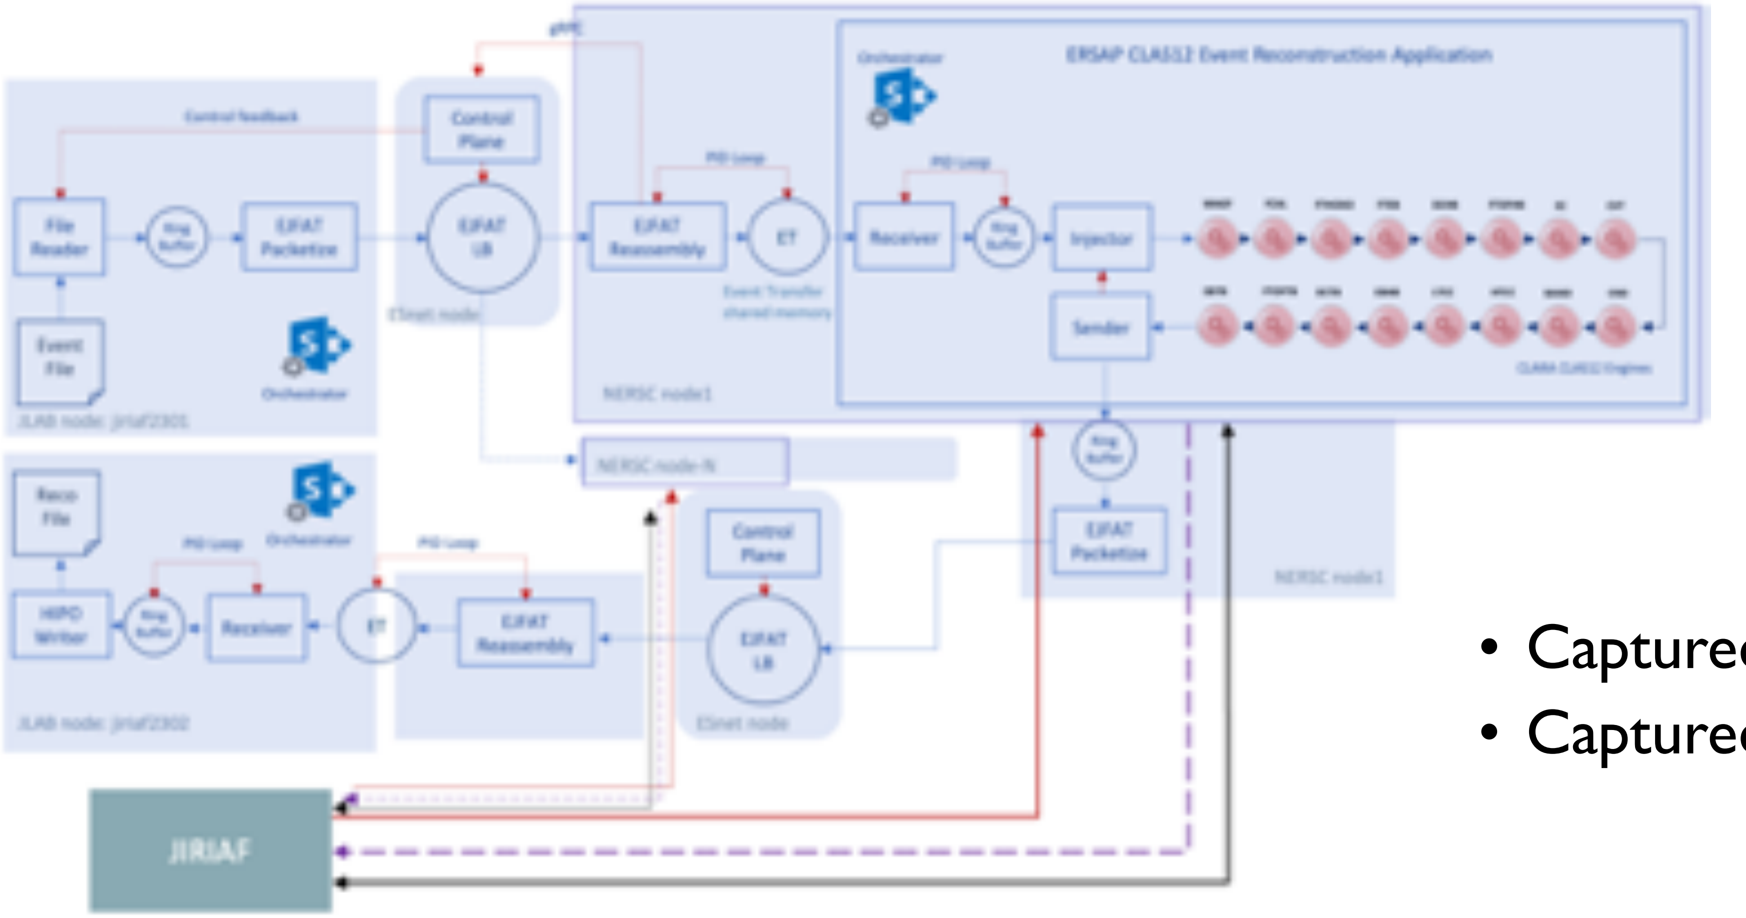

Inaugural demonstration of remote data-stream workflow deployment and processing over a distance

- Captured CLAS12 data, streamed across the Jlab campus using a 100Gbps high-speed NIC featuring hardware timestamps
- Captured data using synchronized streams from multiple network sources

Credit: Vardan Gyurjyan & Ayan Roy

JLab SRO R&D: networking and data transfer

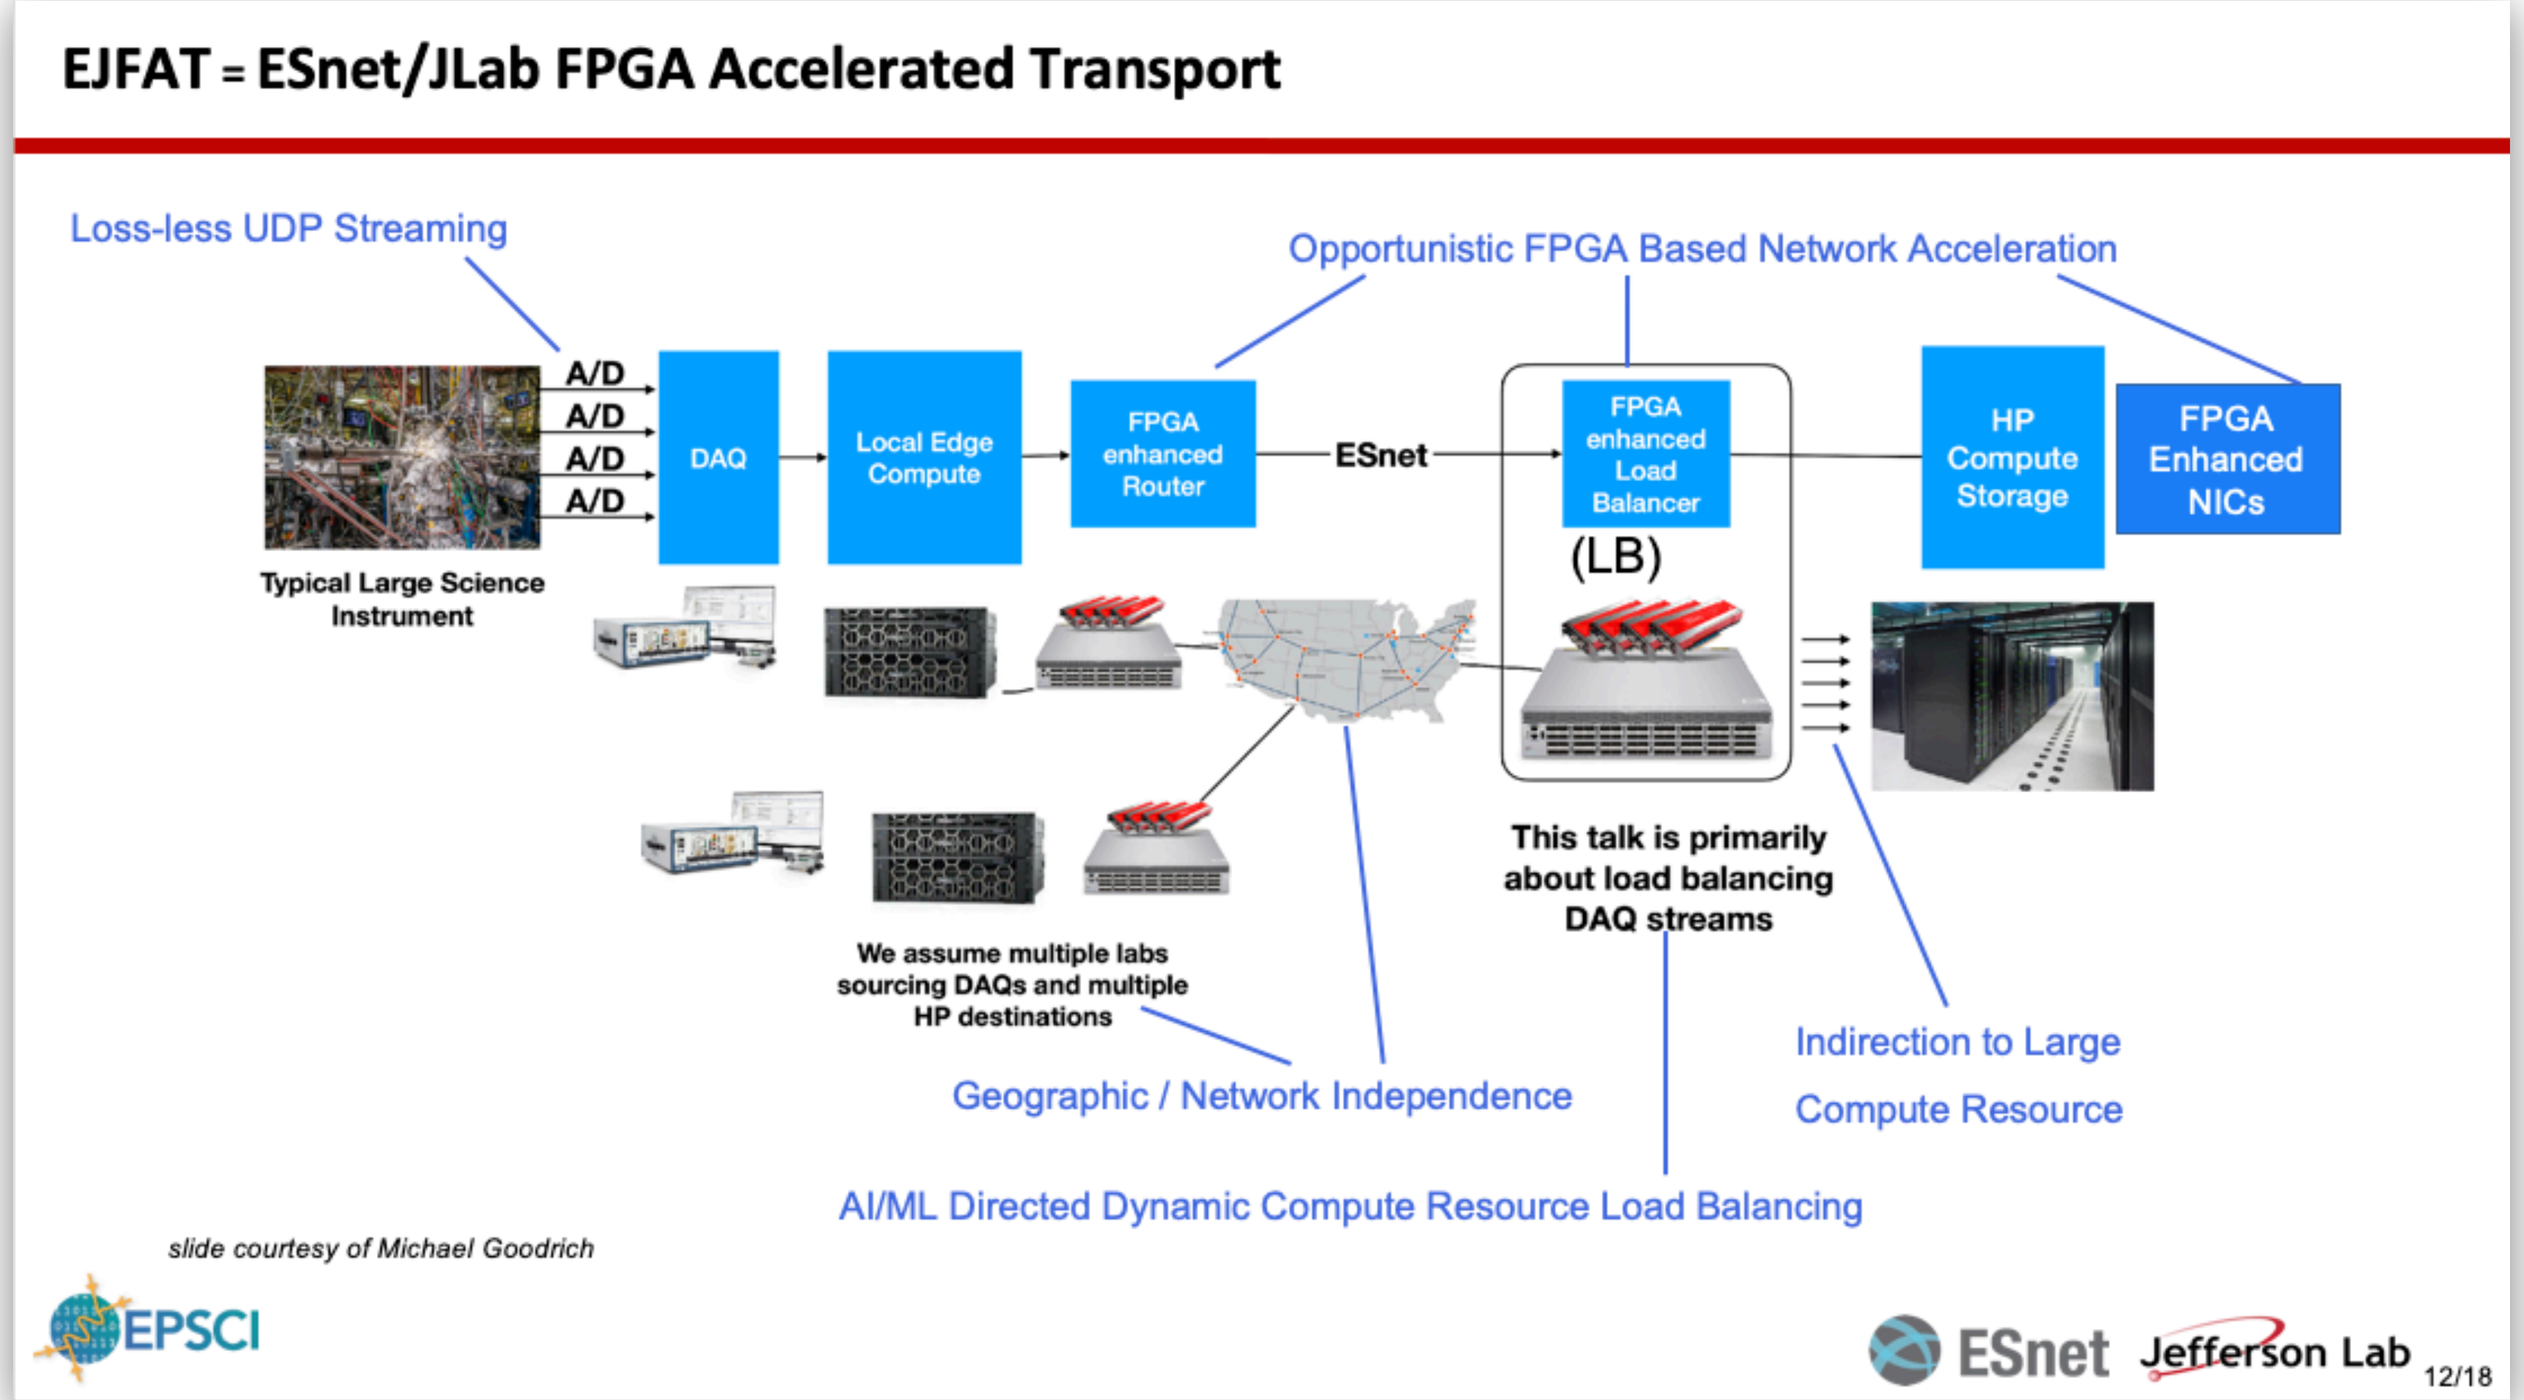

- Re-utilize the existing data transfer infrastructures developed at Jefferson Lab

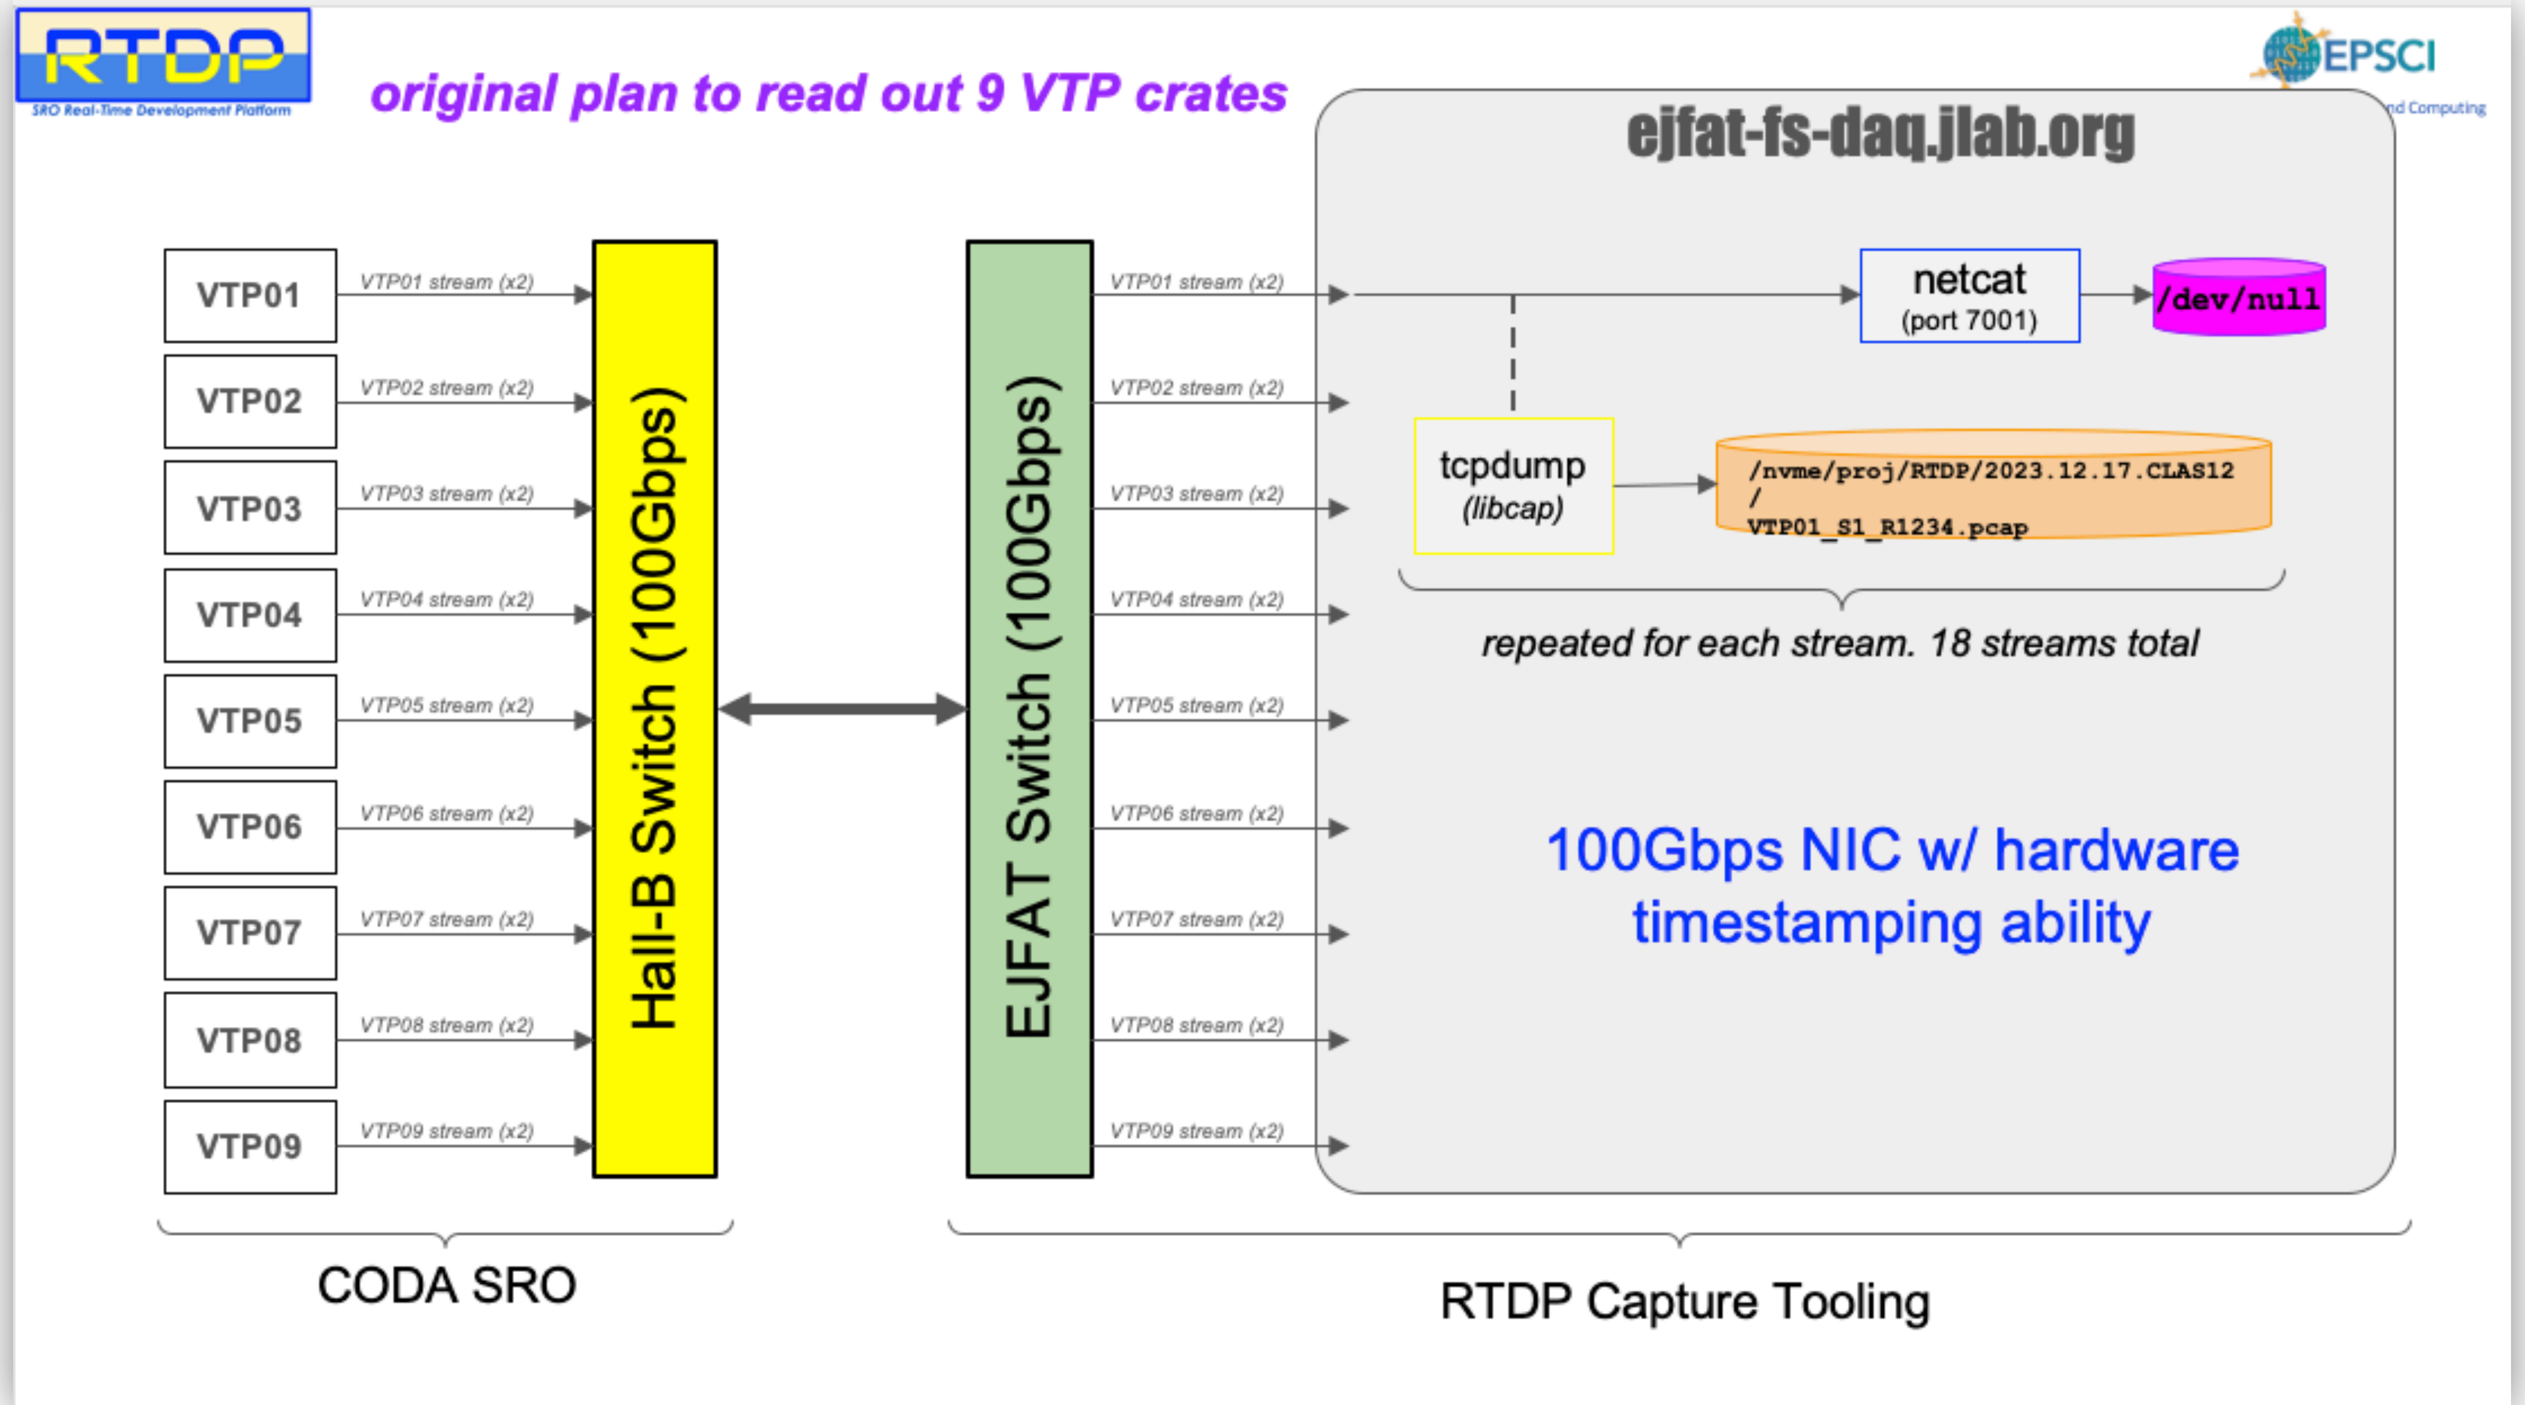

- Networking test for JLab SRO framework just started

JLab SRO R&D: data reduction

Credit: Fabio Rossi

AUTOENCODER

Fully connected auto encoder with dense layer

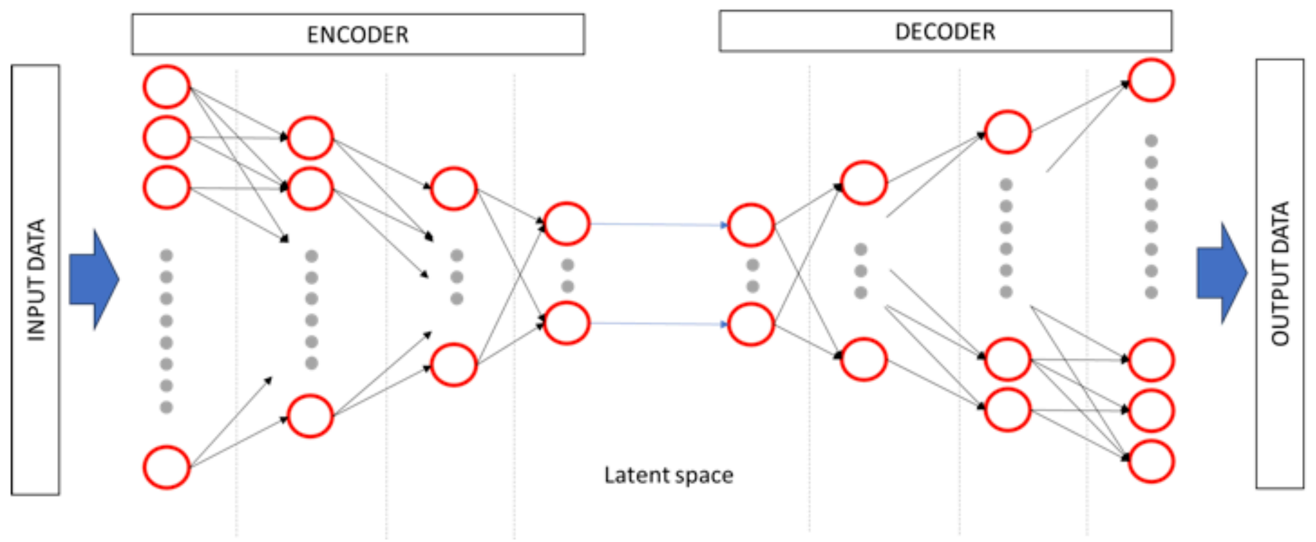

Autoencoder Training: Test results with 96-12 configuration

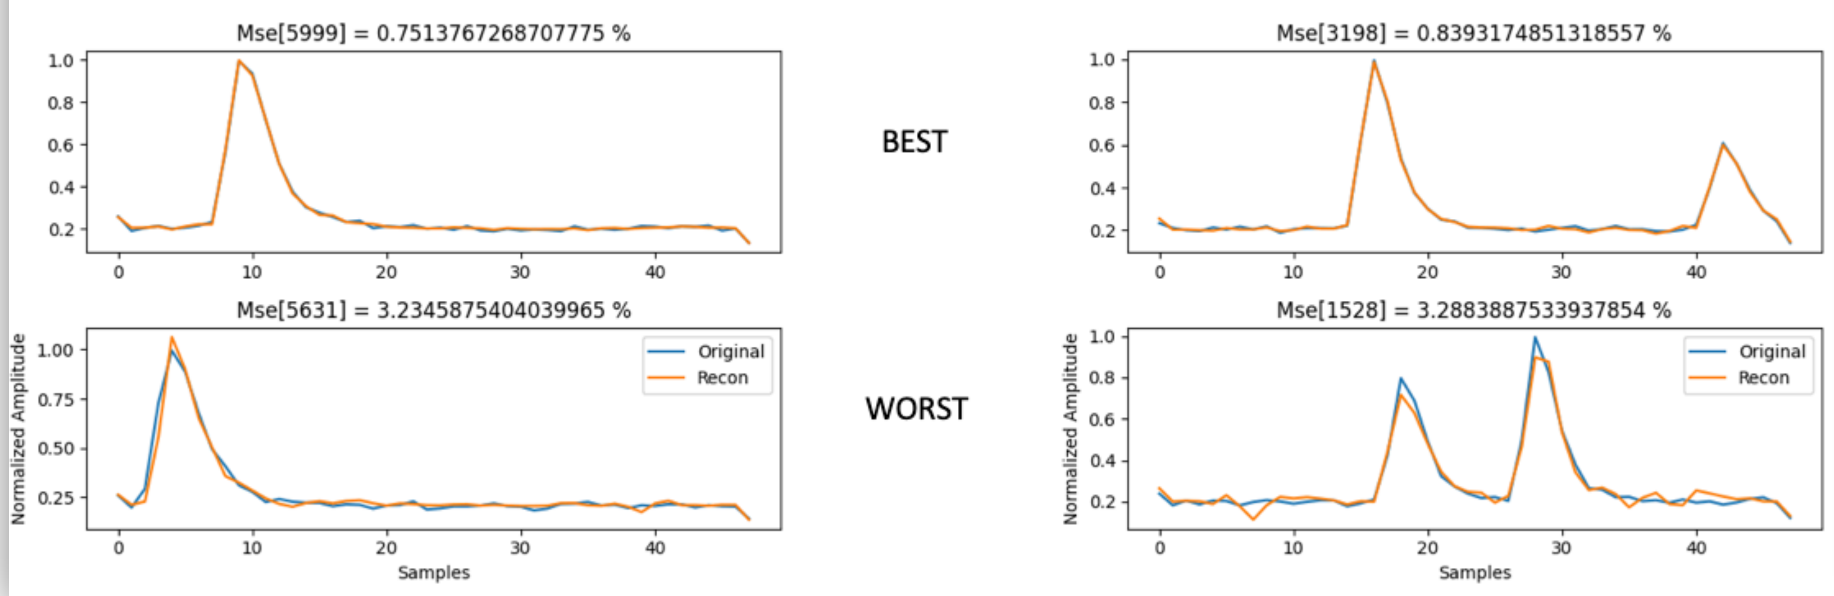

Autoencoder Training: Test results with 96-12 configuration

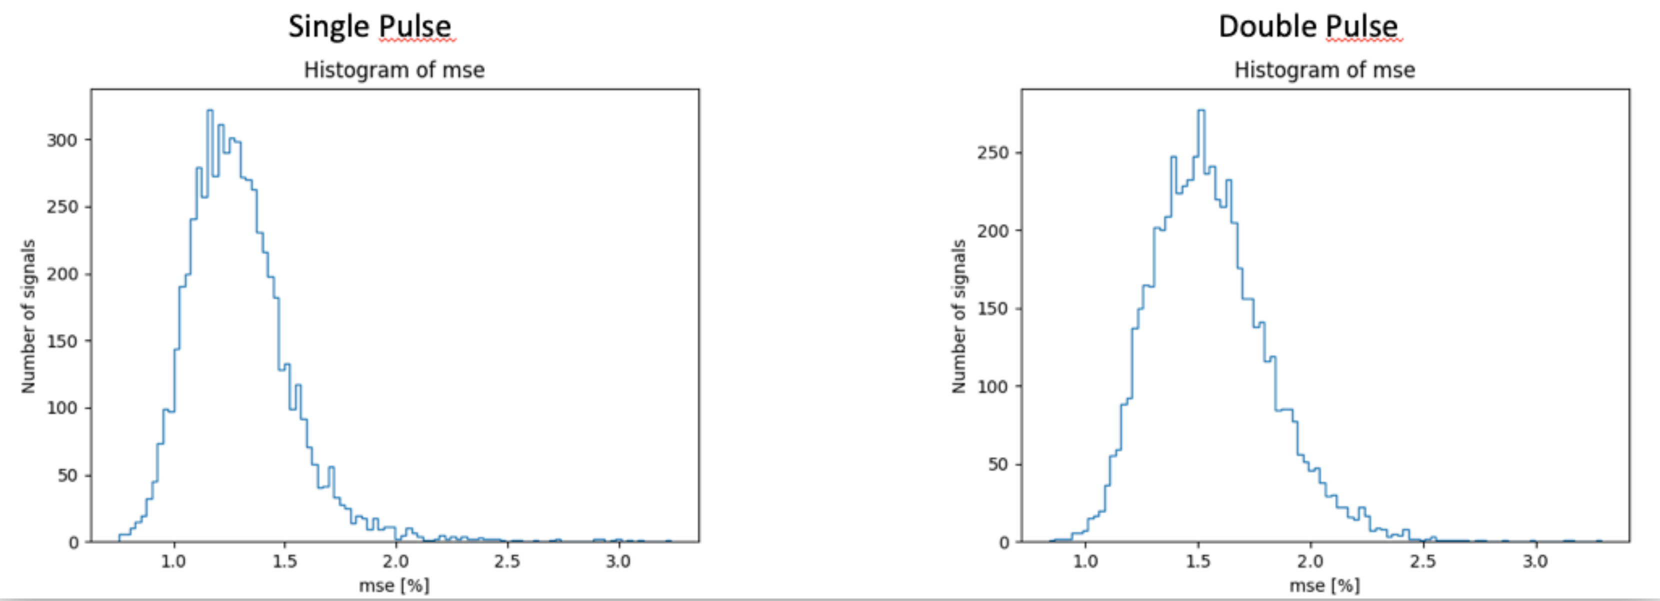

Compression Comparison

Autoencoder perform better on few data

1,52 times better on compression

gzip compress faster than other algoritms

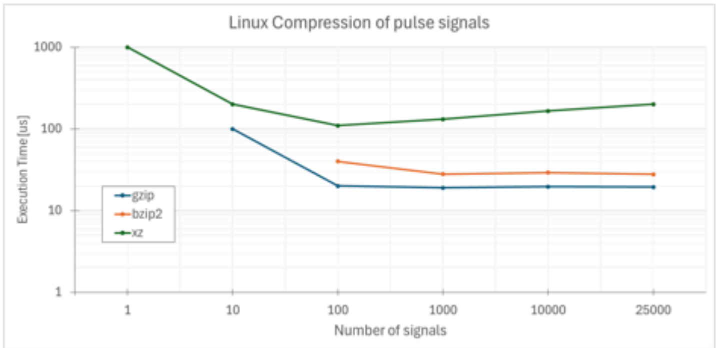

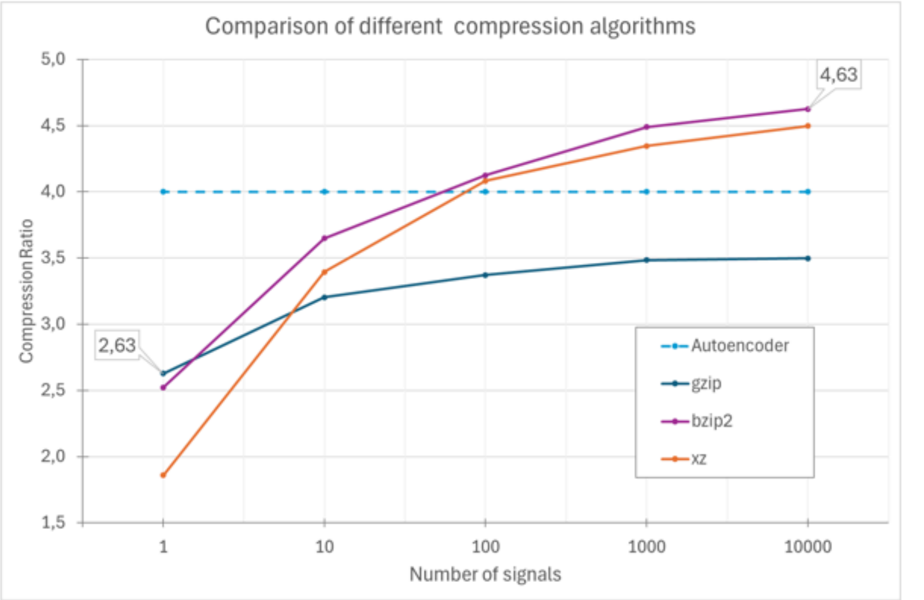

Timing of compression with different hardware

GPU

```
NVIDIA-SMI 535.104.12                 Driver
GPU  Name                  Persistence-M
Fan  Temp   Perf           Pwr:Usage/Cap

  0  Tesla V100-SXM2-32GB            Off
N/A   38C    P0              60W / 300W

  1  Tesla V100-SXM2-32GB            Off
N/A   38C    P0              58W / 300W

  2  Tesla V100-SXM2-32GB            Off
N/A   38C    P0              59W / 300W

  3  Tesla V100-SXM2-32GB            Off
N/A   39C    P0              58W / 300W
```

CPU

AMD Ryzen 5 5600U

6 core, 12 Logic processor

RASPBERRY PI4 (C code)

Quad core Cortex-A72 (ARM v8)

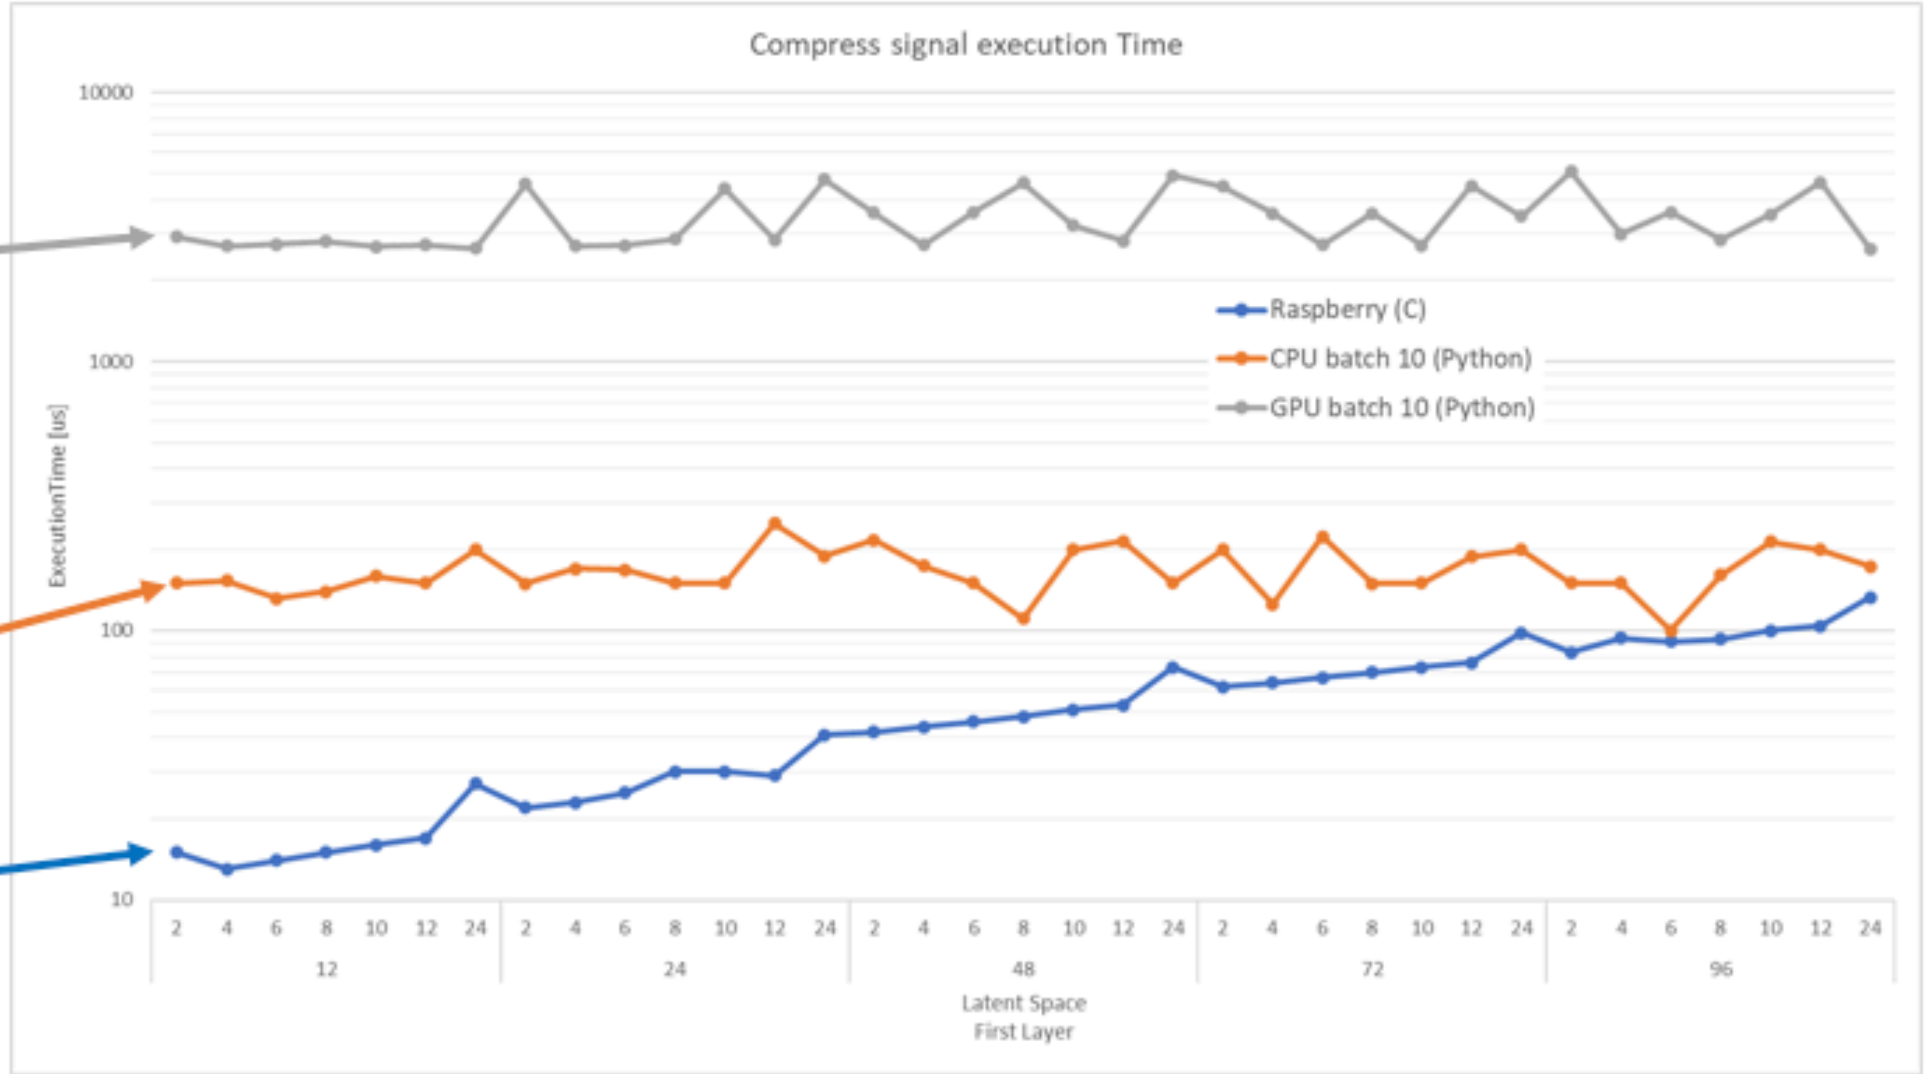

Execution time detailed analysis

Custom implementation with zlib to compress data in memory

Autoencoder perform better both in low cost hardware and on server.

2,47 times faster on low cost hardware

2,51 times faster on server

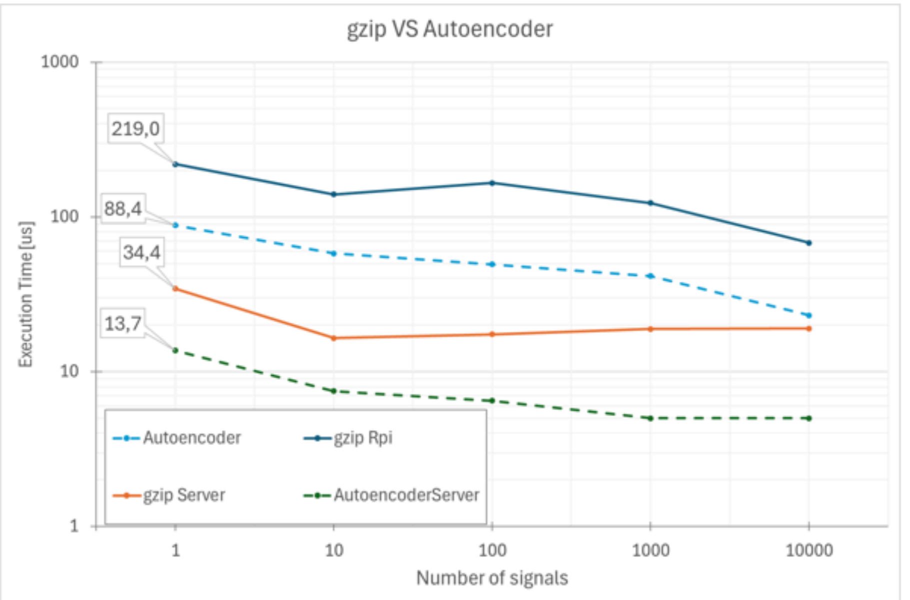

JLab SRO R&: Real-Time detector calibrations

Hall-D GlueX Central Drift Chambers

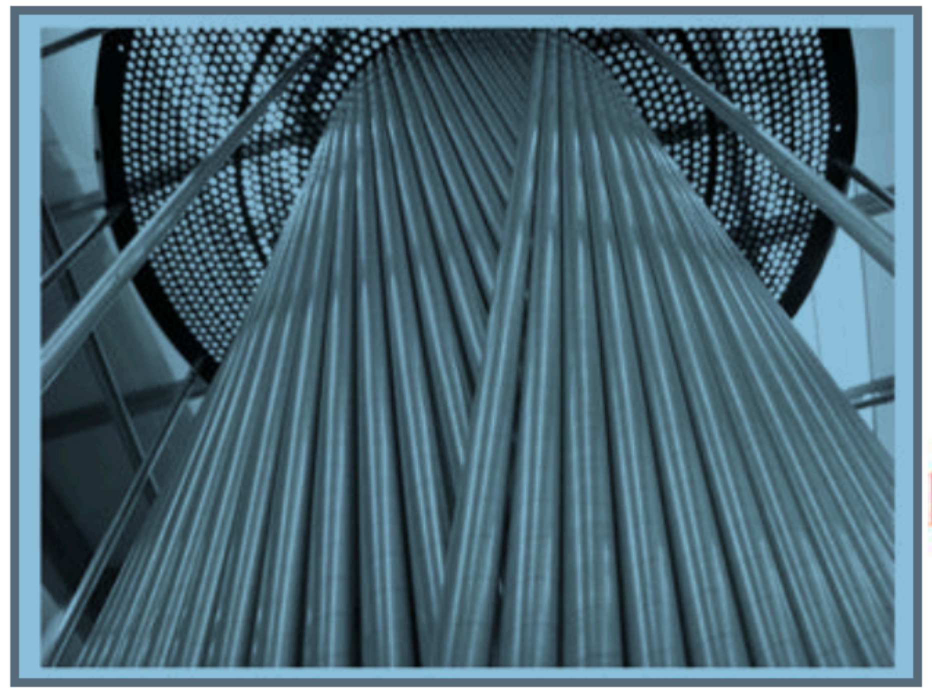

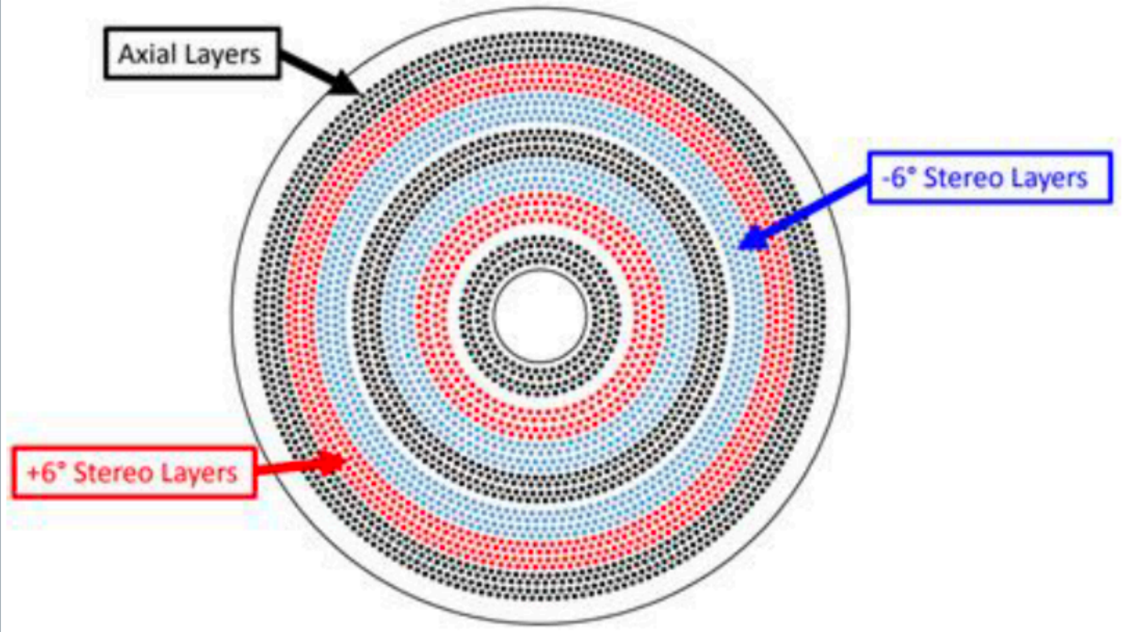

Used to detect and track charged particles with momenta p > 0.25 GeV/c

- 1.5 m long x 1.2 m diameter cylinder
- 3522 anode wires at 2125 V inside 1.6 cm diameter straws
- 50:50 Ar/CO2gas mix

Requires two calibrations: chamber gain and drift time-to-distance

- Gain Correction Factor (GCF): have most variation +/-15%
- Has one control: operating voltage

ML Technique: Gaussian Process (GP)

Target: Provide traditional Gain Correction Factor (GCF)

- atmospheric pressure within the hall
- temperature within CDC
- CDC high voltage board current

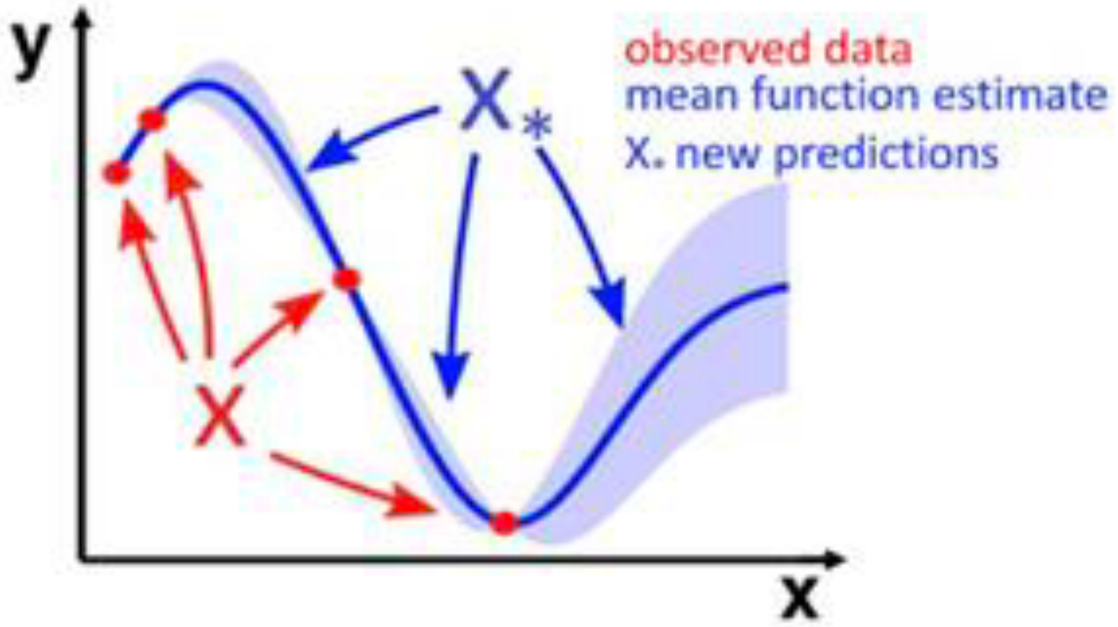

- GP calculates PDF over admissible functions that fit the data
- GP provides the standard deviation (uncertainty quantification -UQ)
- GP kernel: Radial Basis Function + White

It works!

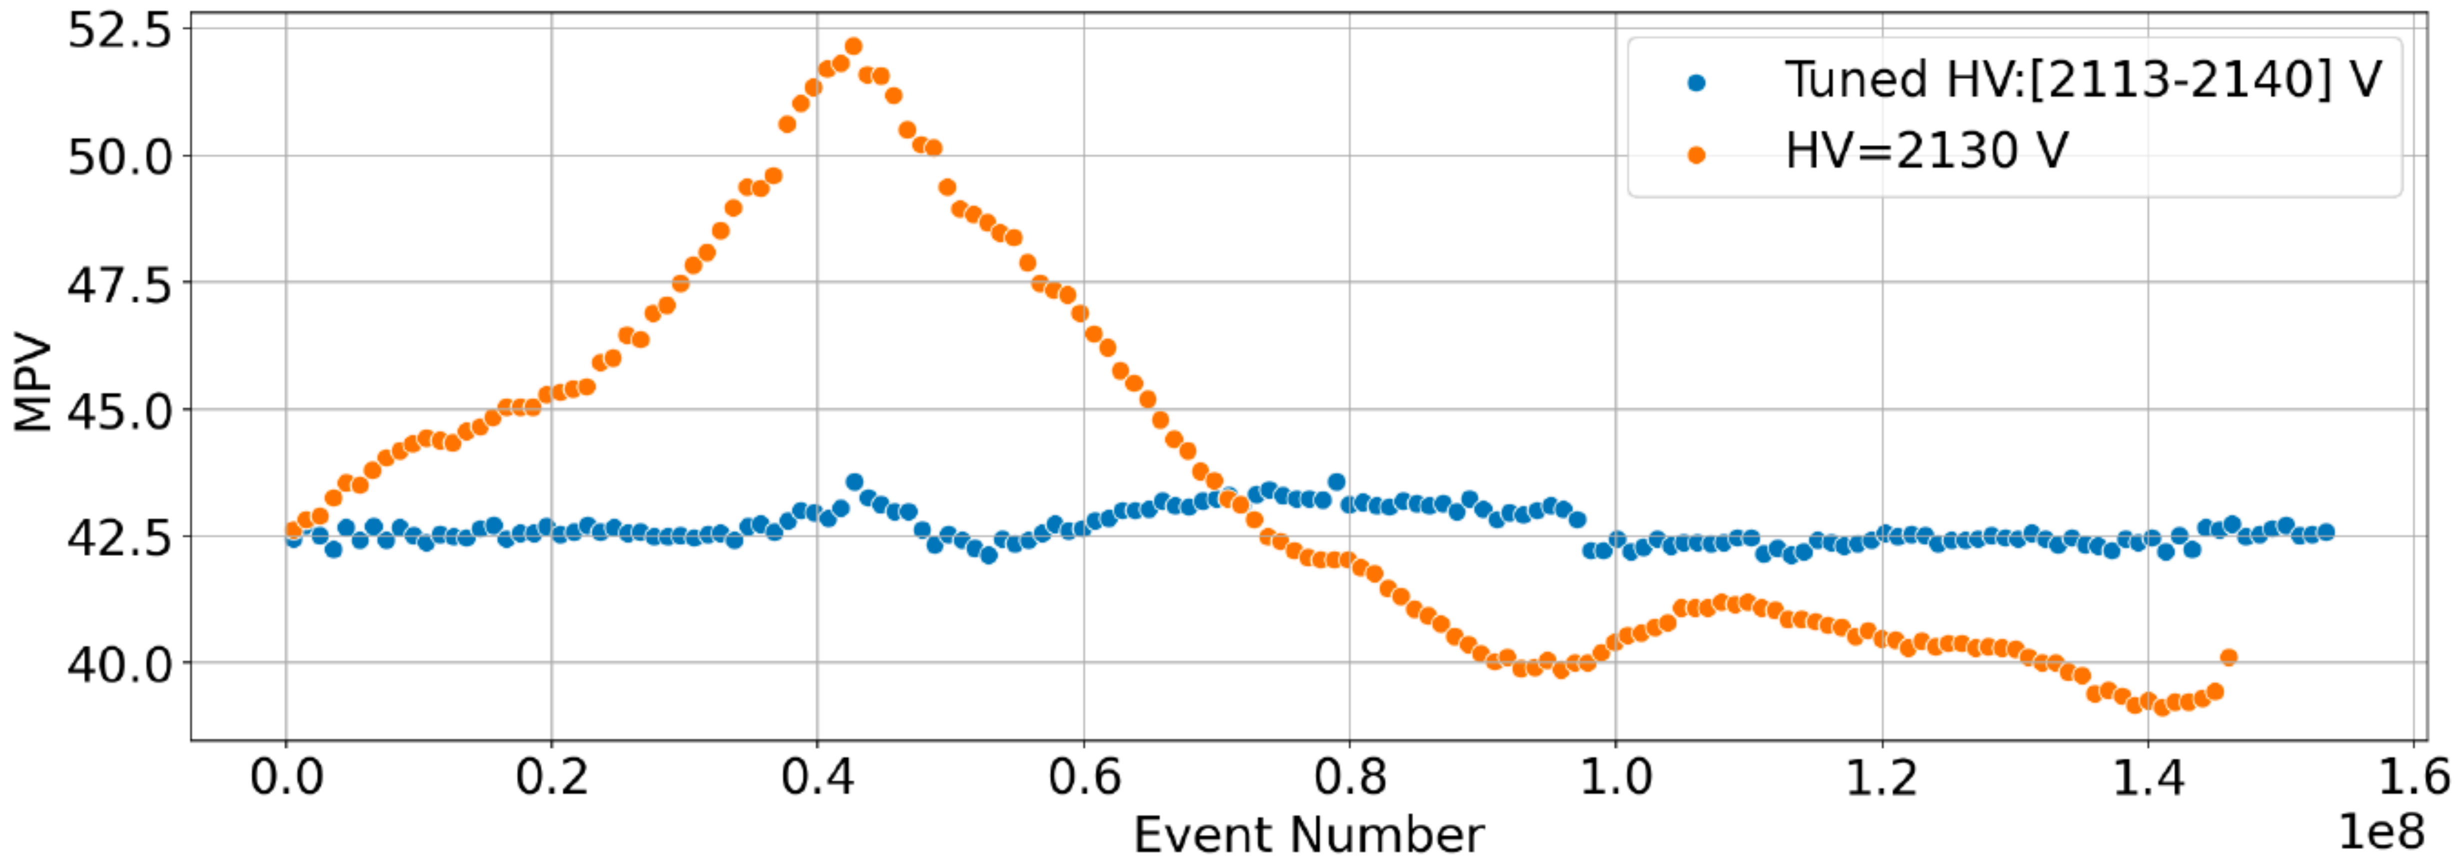

- Half the CDC (orange) at fixed HV, t he other half (blue) had its high voltages adjusted every 5 minutes

JLab SRO R&D: GEANT simulations

Credit: Maurizio Ungaro

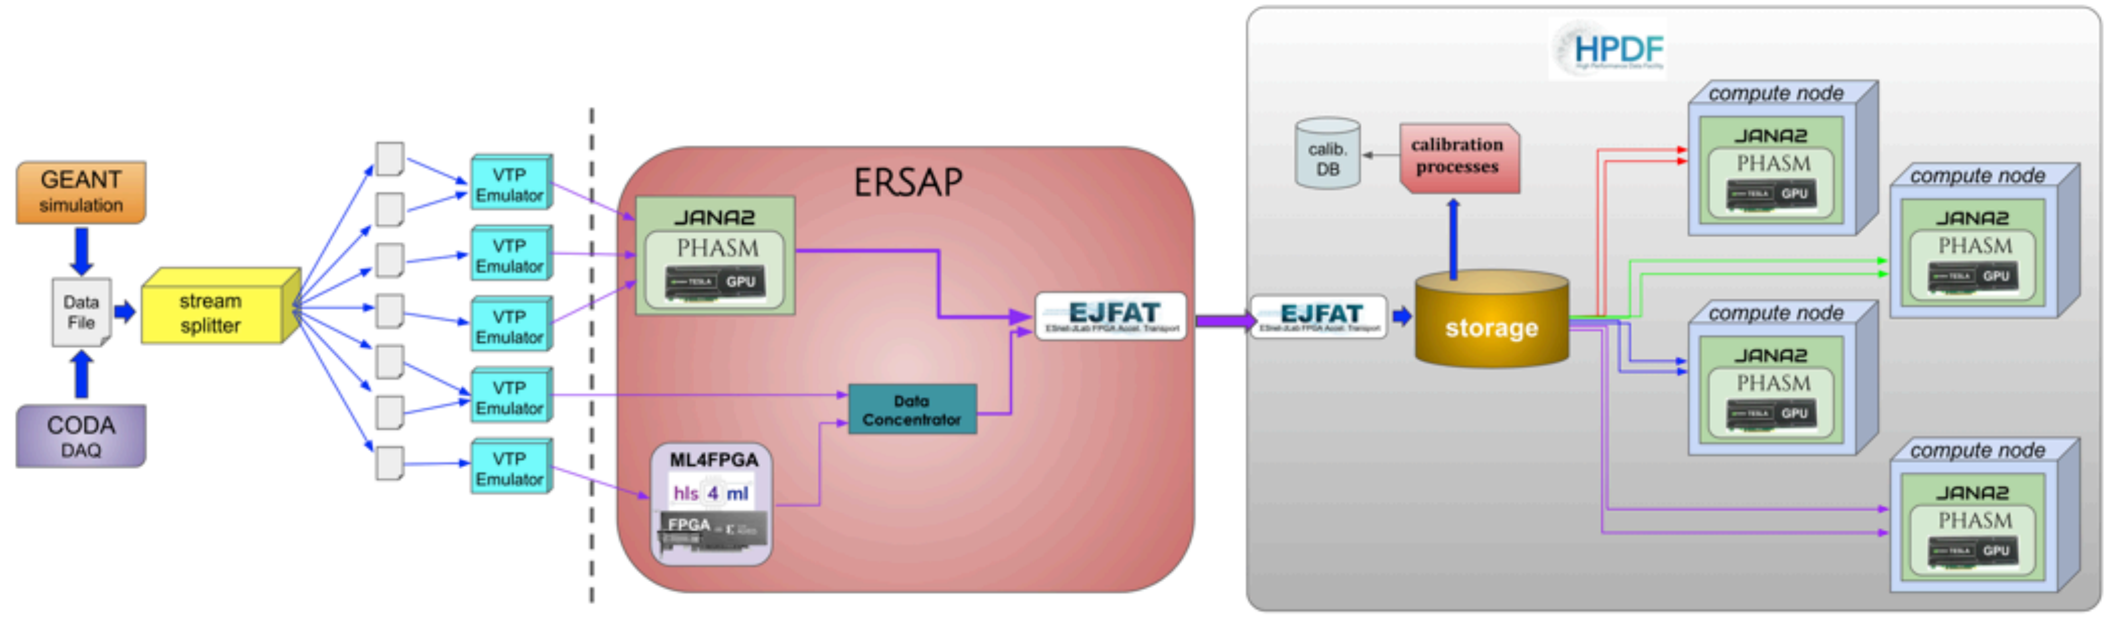

Highly configurable multi-stream source allows realistic streaming simulations

Onsite components will implement first stages of data filtering/reduction

Offsite processing must incorporate built-in calibration latencies and storage. This will also help inform HPDF design

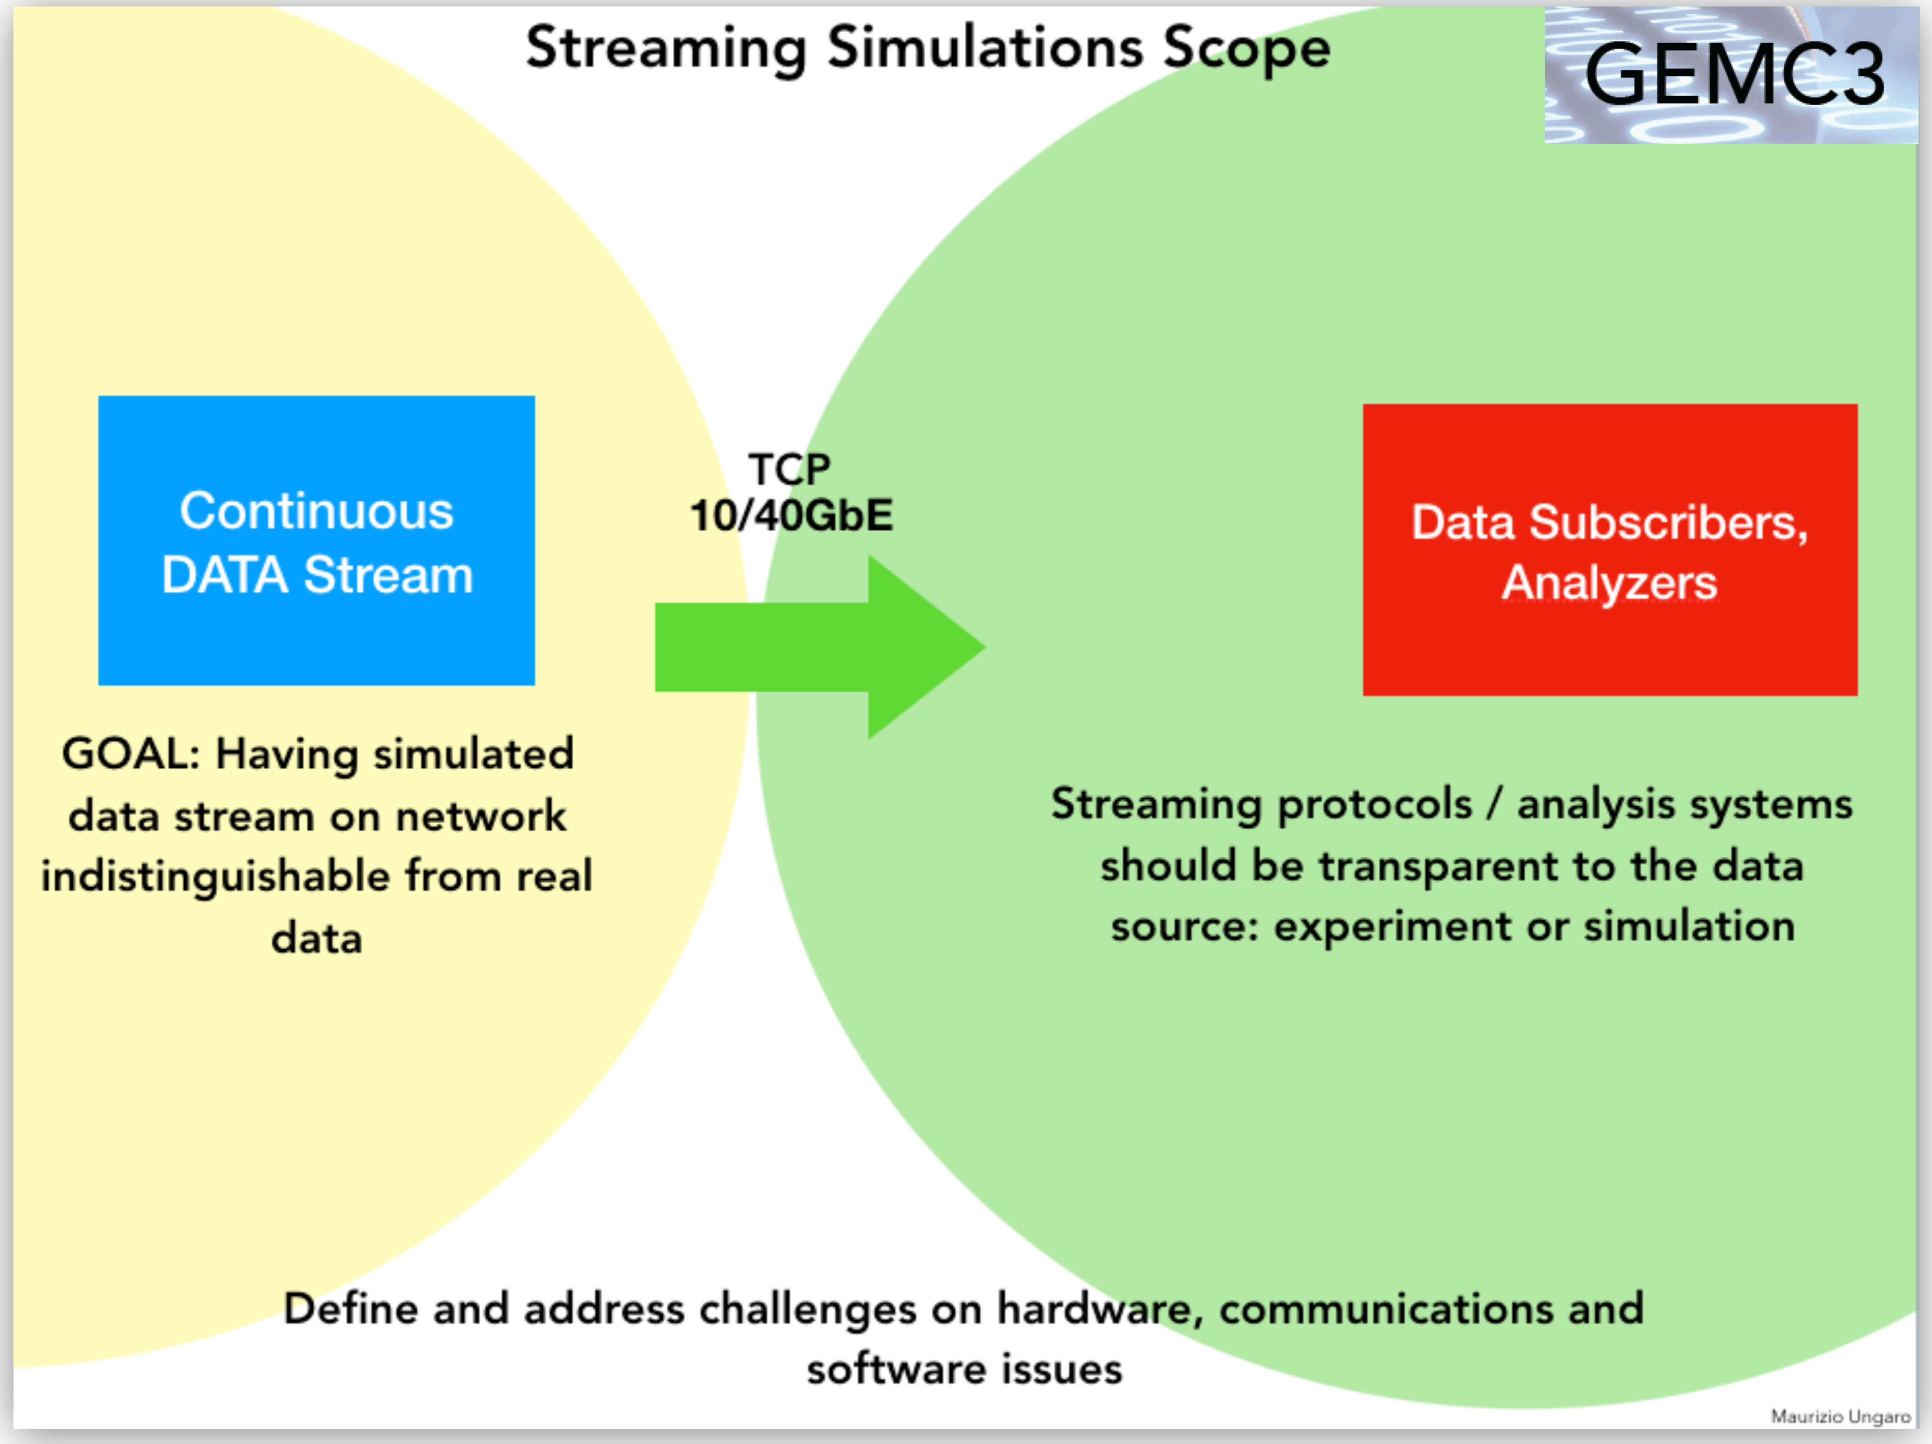

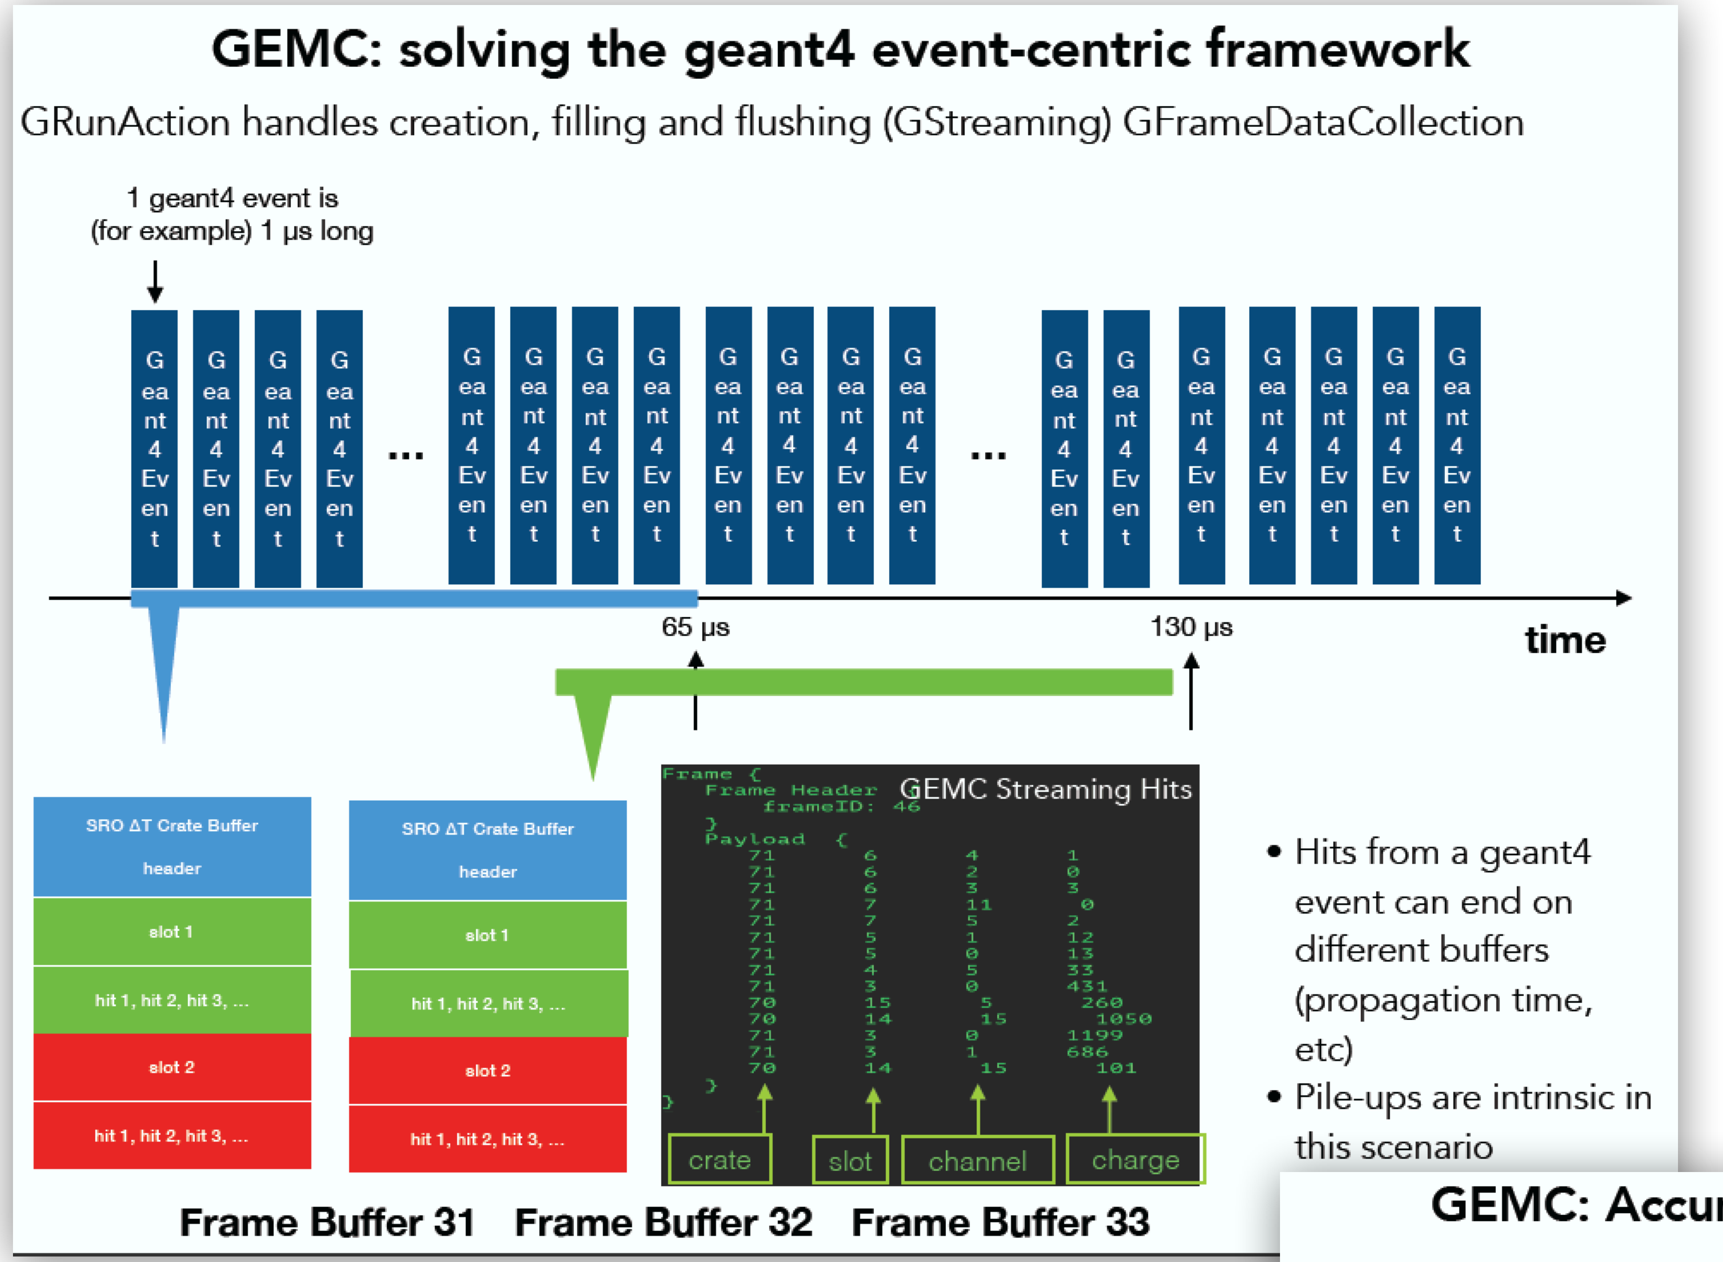

- **Development of a MC GEANT-based toolkit to implement SRO in detector simulations**
  - Transform event-based to stream-based G4 logic
  - Develop libraries share same on-line data format
  - Emulate TCP output to feed to ERSAP
  - Milestone Nov 2021, FT Calorimeter streaming

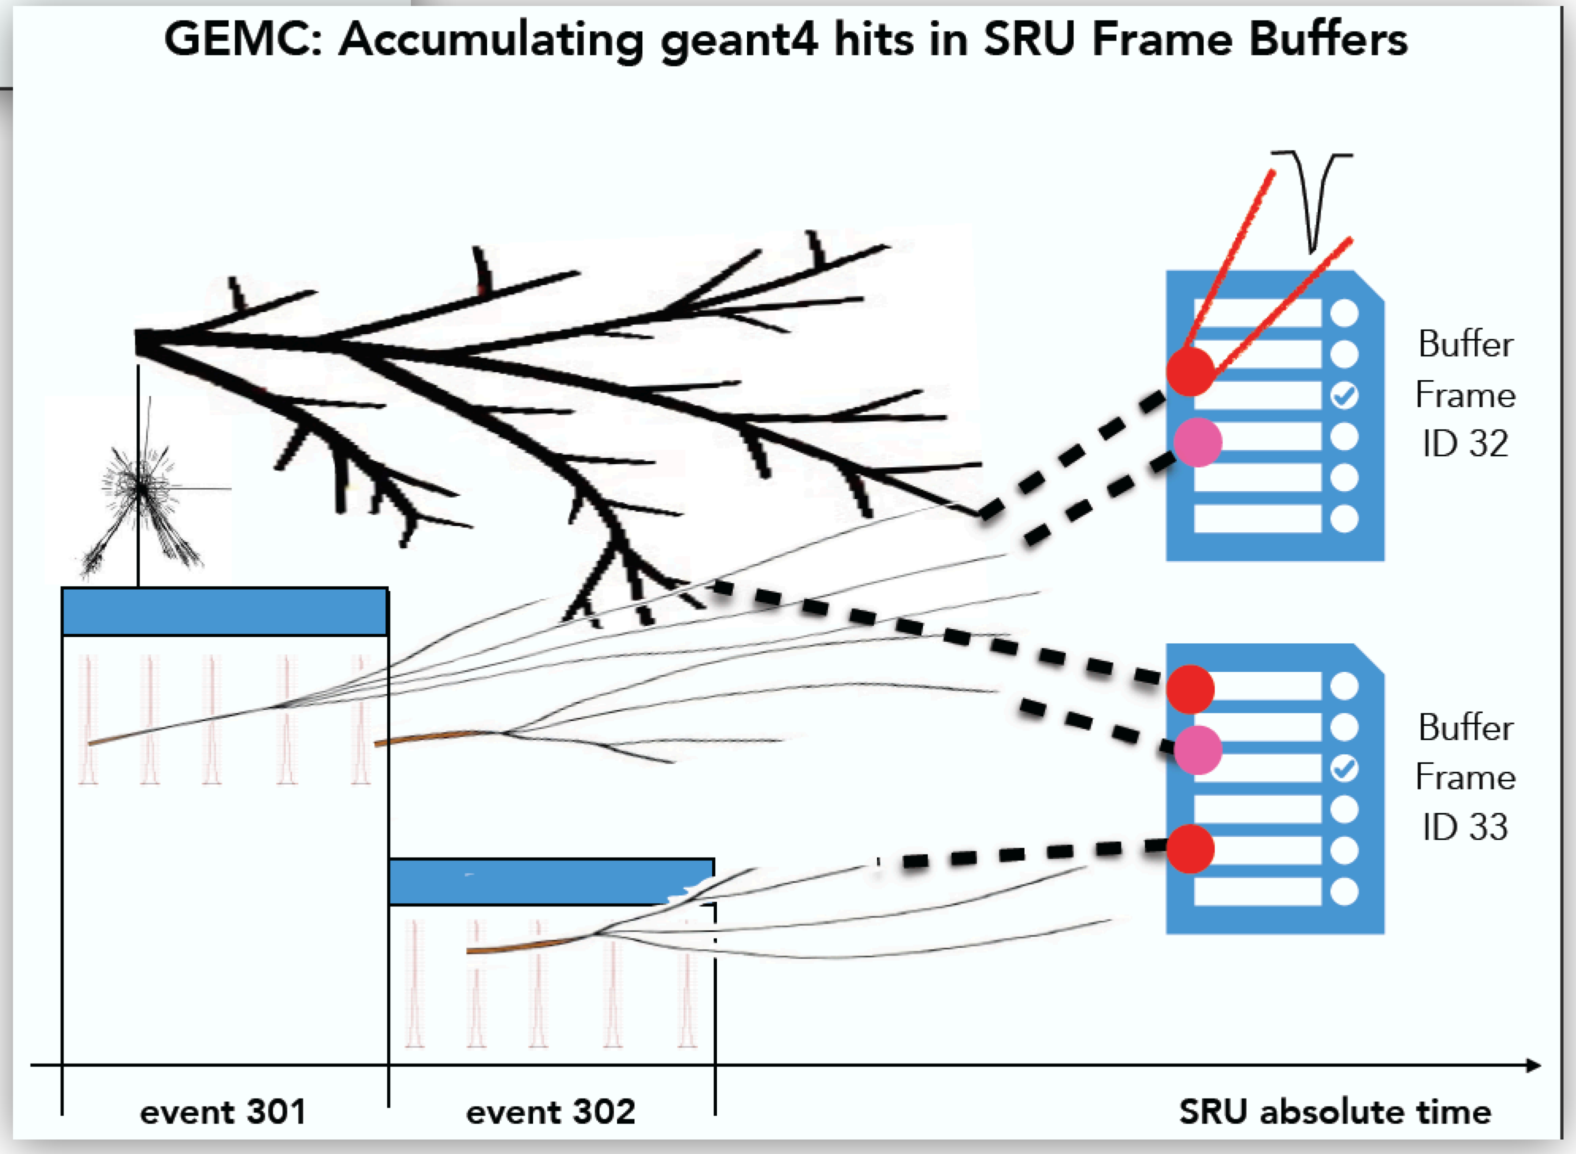

ePIC Streaming Computing

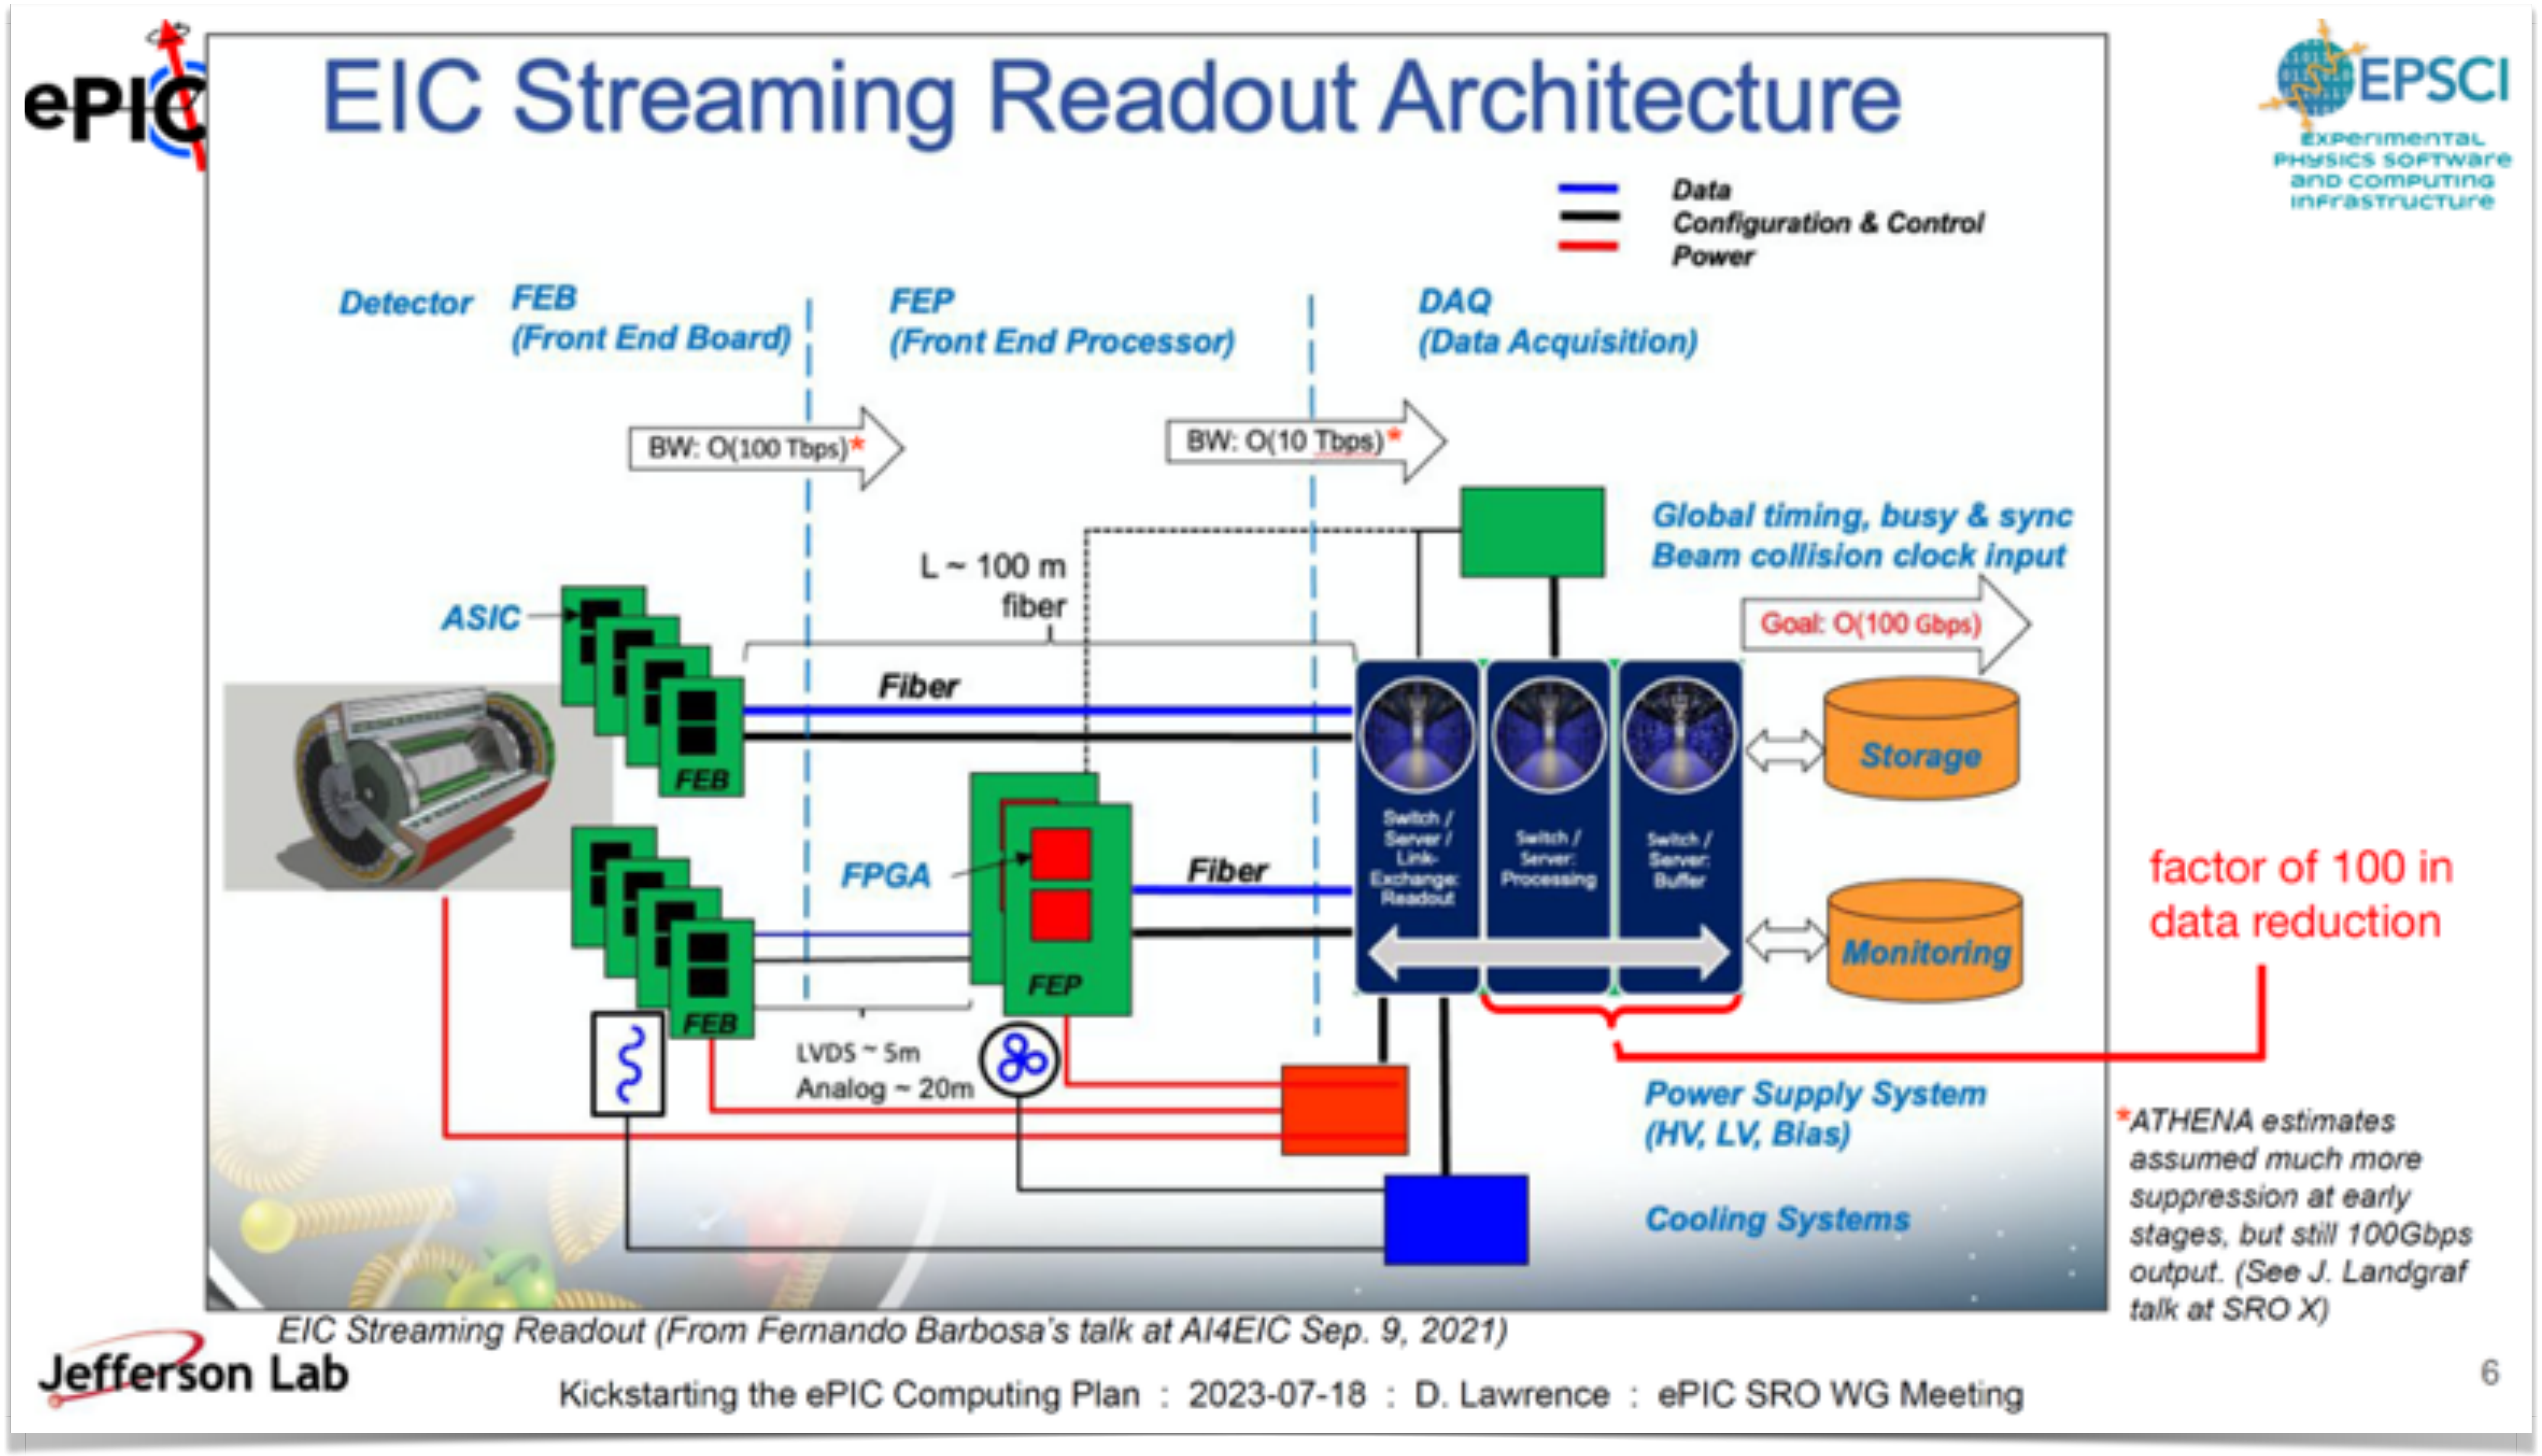

Streaming RO for ePICS

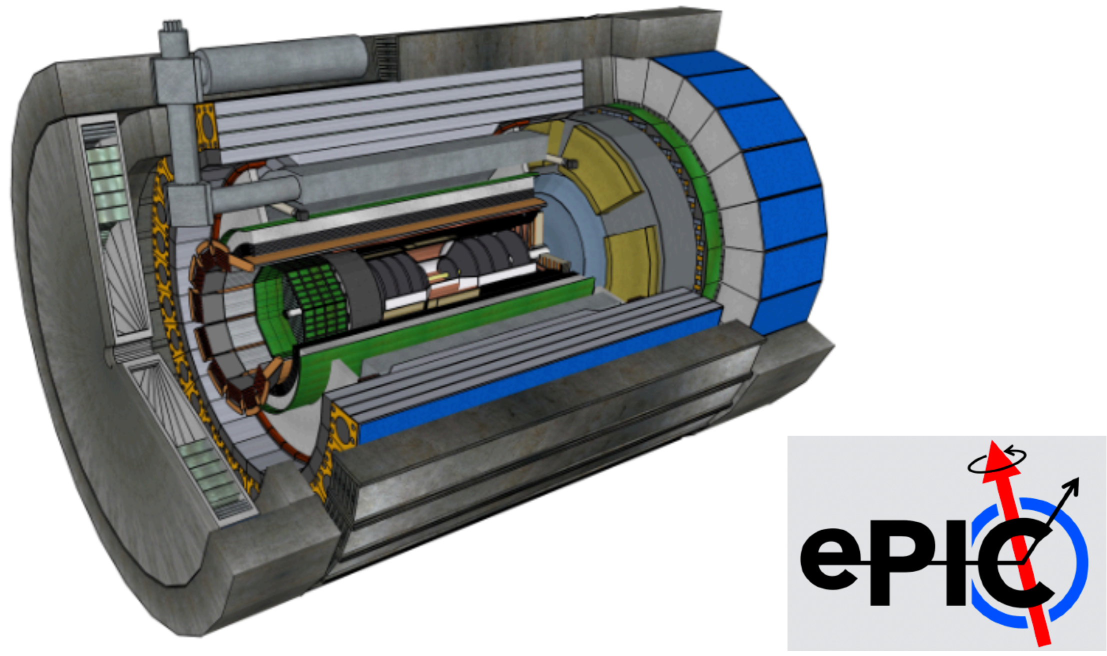

- Full consensus for SRO within the EIC community (Yellow Paper, ECCE, ATHENA, …)
- Rates at ePICS not comparable to LHC HI-LUMI but the advantages of SRO remain:
  - multiple channels to trigger on
  - Holy Grail: to manage (storage) an unbiased (un-triggered) data set for further analysis
  - on/off-line event selection with full detector information

Credit: Fernando Barbodsa, Markus Diefentaler

ePIC Streaming Computing

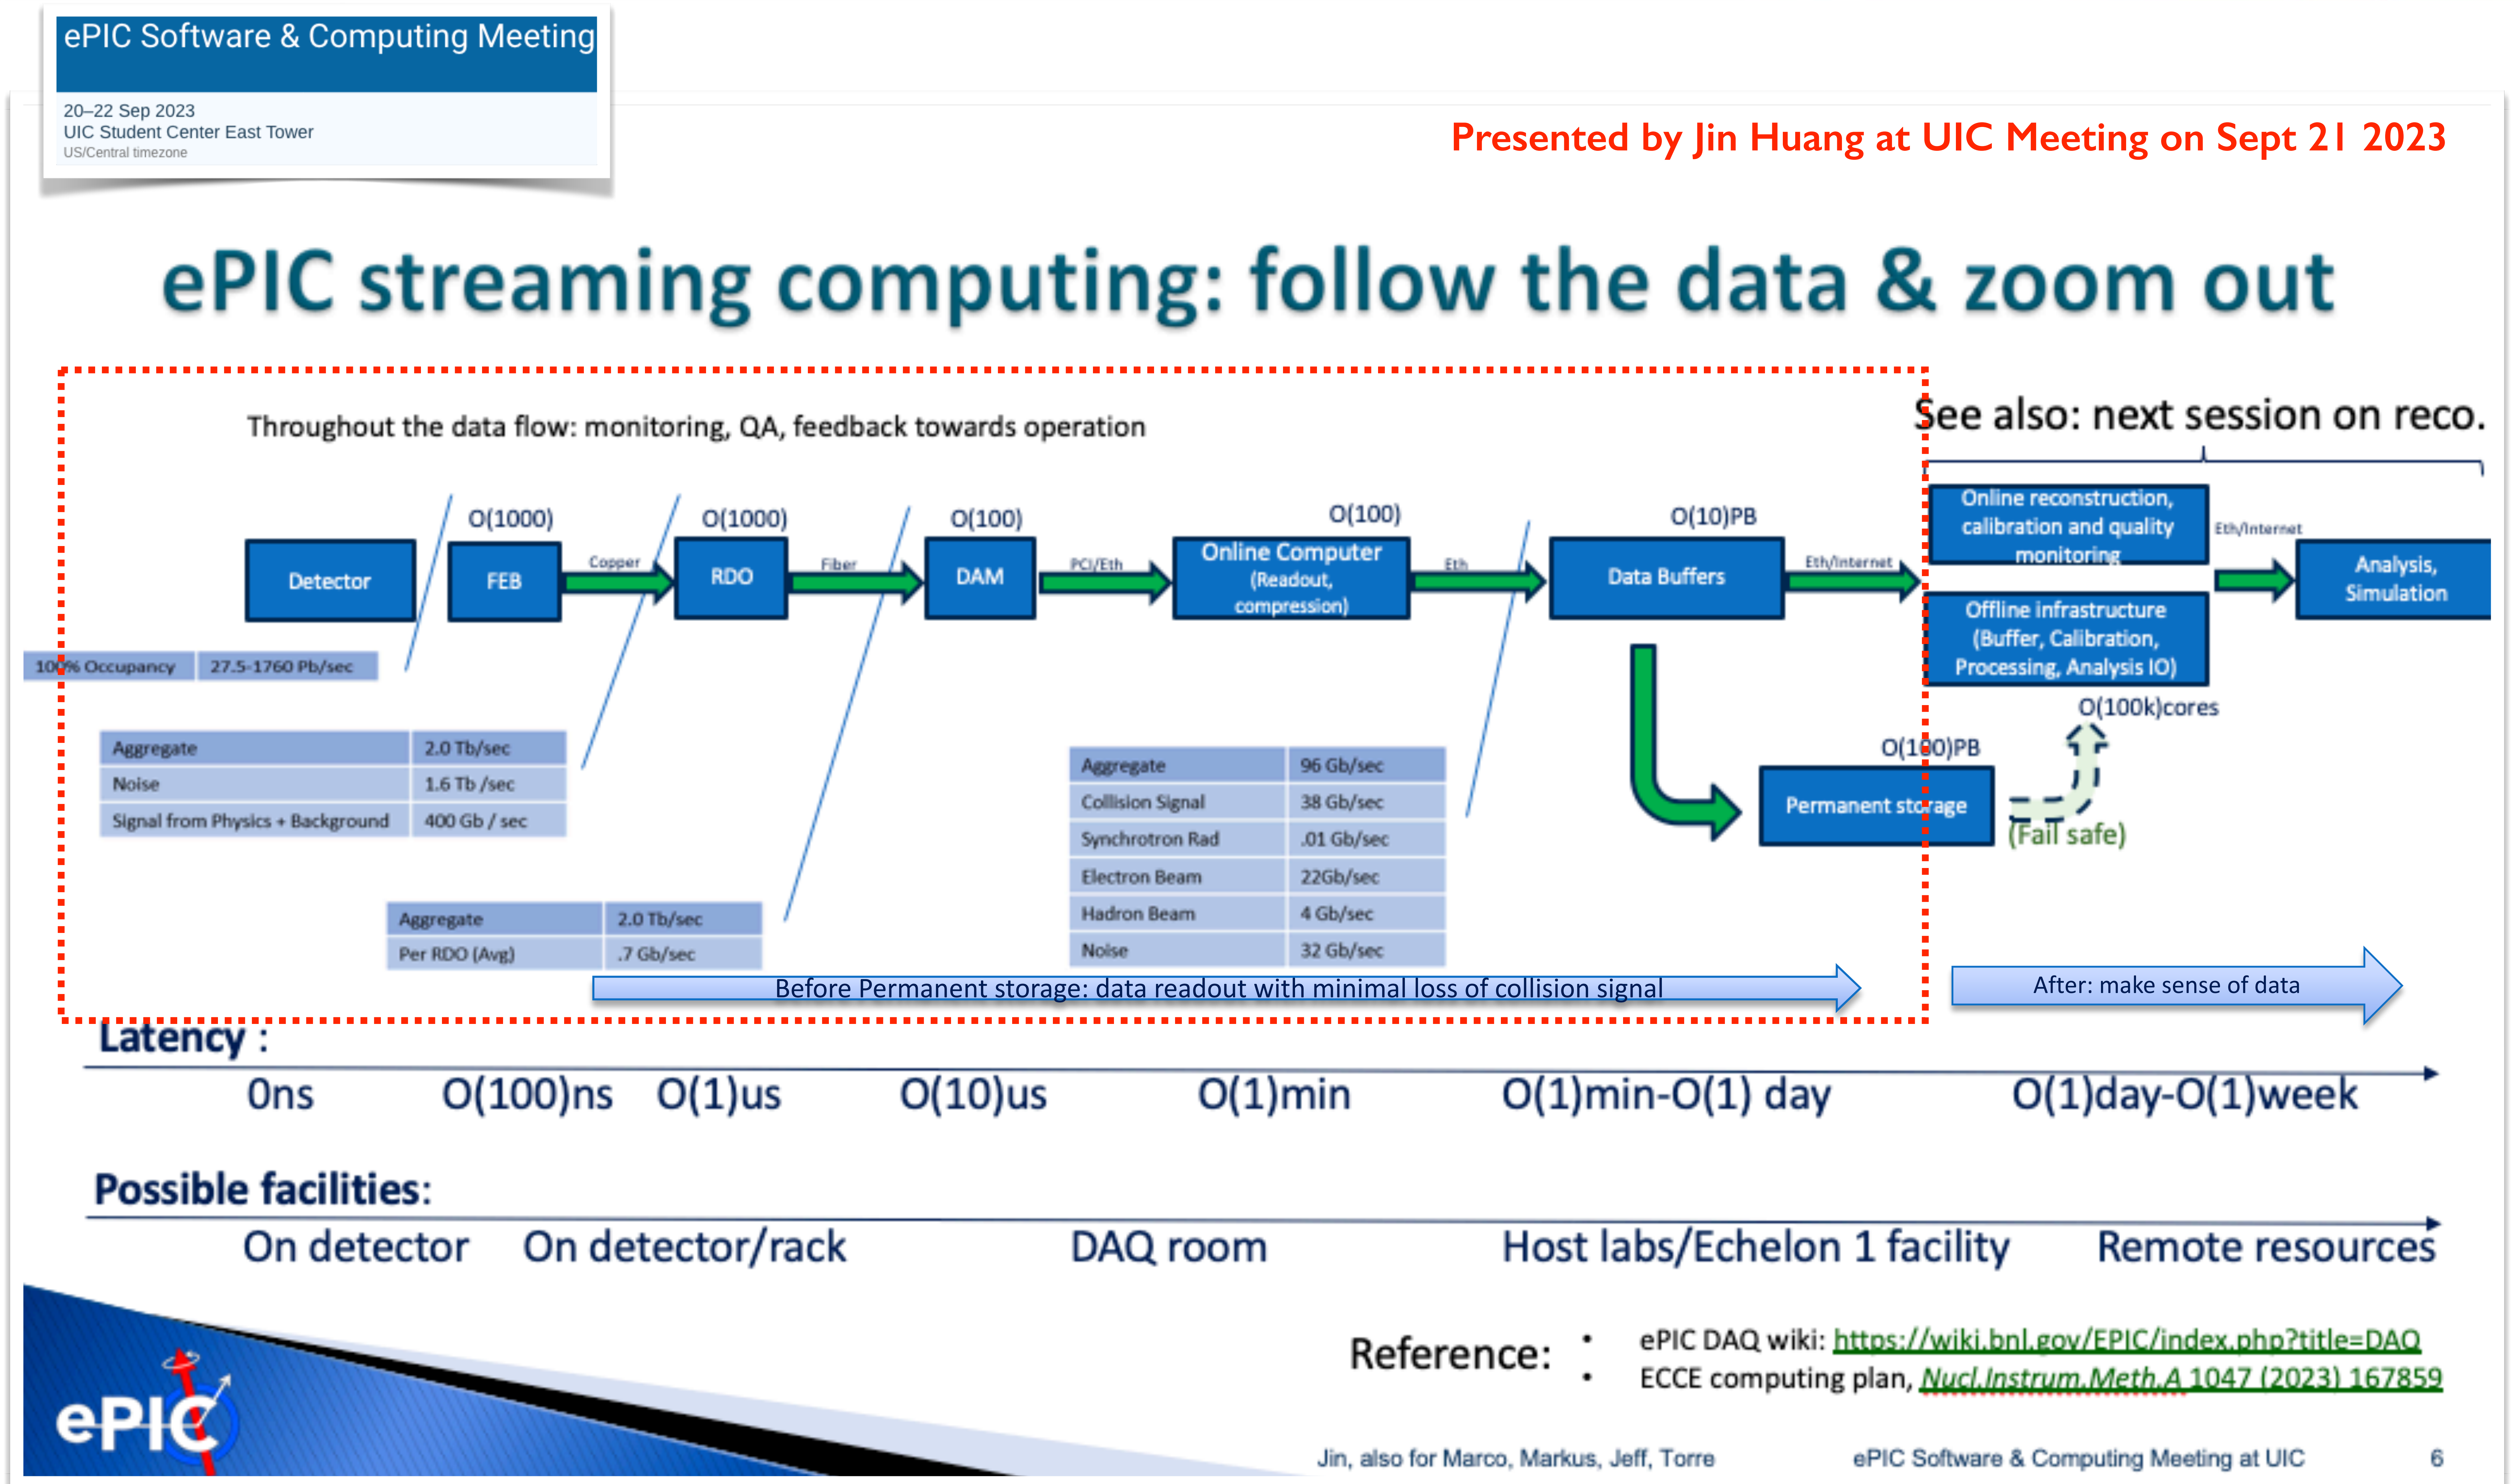

Interfaces

- Each step in the workflow has a different latency
- Identify interfaces for a 'service-oriented' approach

Within the 'control room'

- Each stage in data flow requires IO specs (based on CPU, GPU, FPGA reduction)
- 'control room' boundary based on permanent data storage

Outside the control room

- Networking
- CPU/GPU farm
- Local/remote resources
- on/off-line analysis

Credit: Jin Huang, Jeff Landgraf

Summary and Outlook

- Streaming RO is needed for high-luminosity experiments to explore increasingly rare reactions
- Streaming RO is 'THE' option for future electron beam experiments running at high luminosity (JLab, EIC, …)
- Take advantage of the full detector's information for optimal (smart) tagging/filtering
- So many advantages: performance, flexibility, scaling, upgrading …

  … but, has to demonstrate to be as effective (or more!) than triggered systems
- Streaming Readout on-beam tests performed in Hall-D and Hall-B at JLab
- First SRO chain (FE + SRO sw + ON-LINE REC) tested with existing hardware
- Deployment of JLab SRO framework based on micro-services architecture (ERSAP)
- Taking advantage of current JLab operations for on-beam tests
- Development of ancillary components such as G4 MC (GEMC3), real-time calibrations, …
- Developing solutions applicable to the future Electron Ion Collider
- Built a working SRO prototype and a work team!

Many thanks to the whole SRO team:
D.Abbott (JLab), F.Ameli (INFN), MB (JLab/INFN), V.Berdnikov (CUA), S.Boyarinov (JLab) M.Bondí (INFN), N.Brei (JLab),L.Cappelli (INFN) A.Celentano (INFN), T.Chiarusi (INFN), C.Cuevas (JLab), R. De Vita (INFN), M.Diefenthaler (JLab), J. Landgraf (BNL), W. Gu (JLab), C.Fanelli (MIT), R. Fang (ODU), A. Farhat (ODU), S.Furletov (JLab), F.Giacomini (INFN), V.Gyurjyan(JLab), W.Gu (JLab), G.Heyes (JLab), T.Horn (CUA), J.Huang (BNL), E. Jastrzembski (JLab), D.Lawrence (JLab), L.Marsicano (INFN), X.Mei (JLab), P.Moran (MIT), P.Musico (INFN), C.Pellegrino (INFN), E. Pooser (JLab), B.Raydo (JLab), A.Roy (JLab), F.Rossi (INFN), H. Szumila-Vance (JLab), C.Trimmer (JLab), M.Ungaro (JLab), S.Vallarino (INFN), J-Y Tsai (JLab), Y. Xu (ODU).