Streaming Readout and Data-Stream Processing at JLAB

Παντα ρει

V. Gyurjyan

Jefferson Lab

U.S. DEPARTMENT OF ENERGY | Office of Science

JSA

“Enable full offline analysis chains to be ported into real-time, and develop frameworks that allow non-expert offline analysis to design and deploy physics data processing systems.”

A Roadmap for HEP Software and Computing R&D for the 2020s. HEP Software Foundation, Feb. 2018

JLAB Experimental Halls

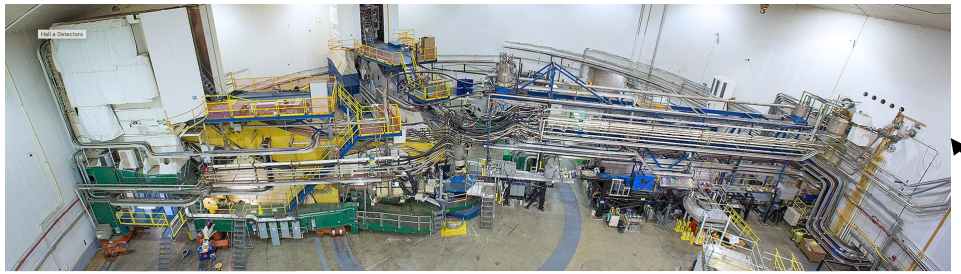

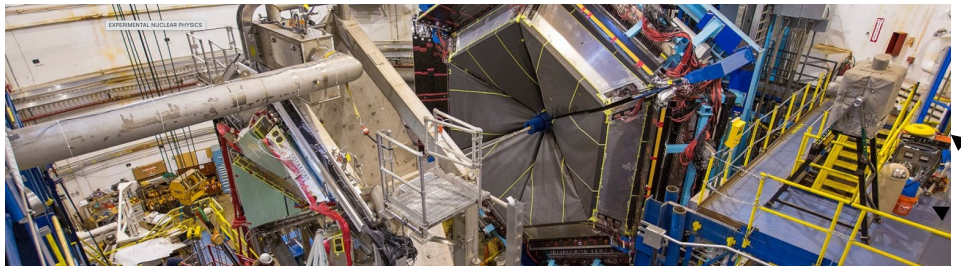

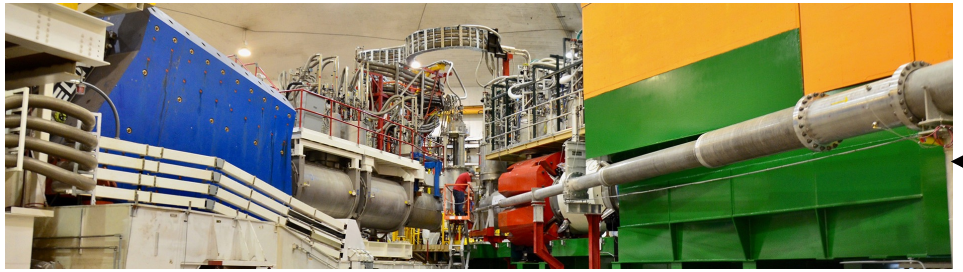

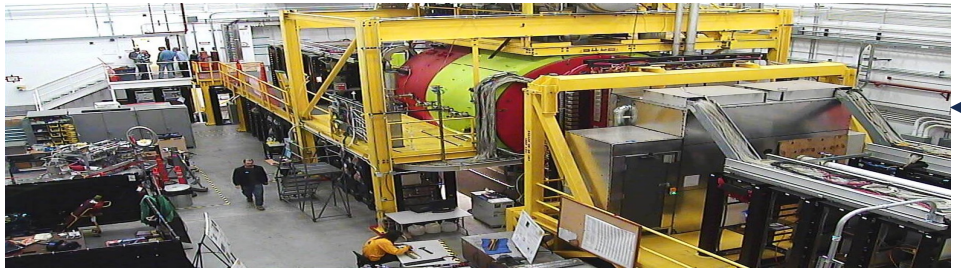

- Four experimental end-stations with different experimental equipment.
- Custom electronics for digitization
- Current and future experiments present increased DAQ requirements
  - Including streaming readout, FPGA-based edge processing, etc.

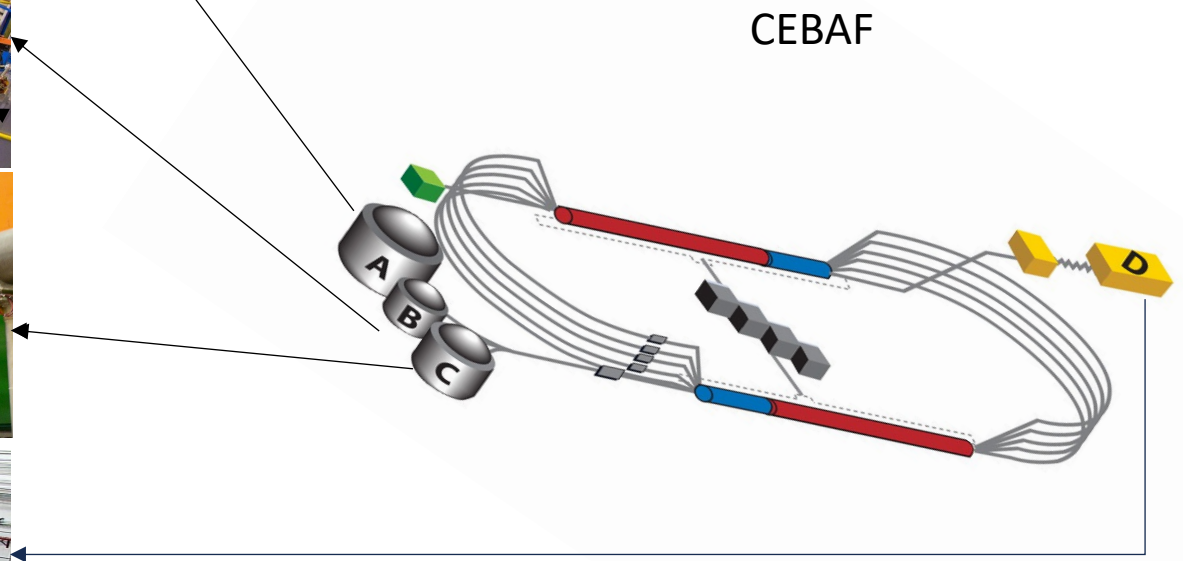

JLAB Grand Challenge in Readout and Analysis for Femtoscale Science

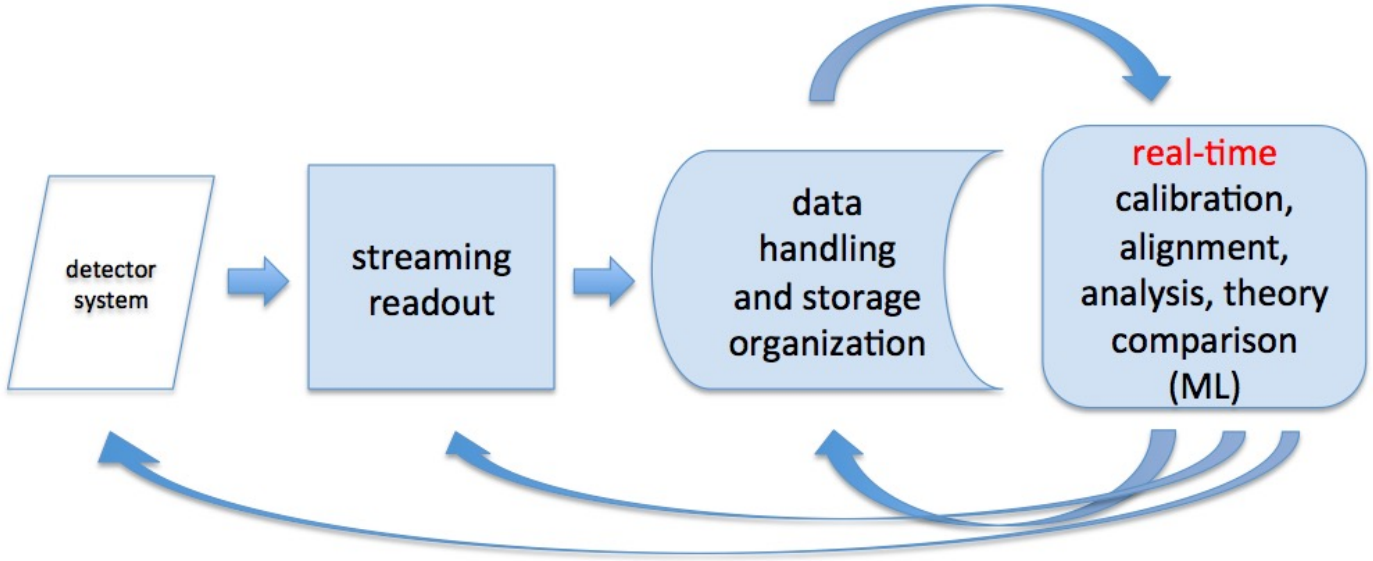

Integrated whole-experiment approach to detector readout and analysis toward scientific output.

Courtesy of : Amber Boehnlein, Markus Diefenthaler, Rolf Ent,
Graham Heyes, Cynthia Keppel, David Lawrence, Brad Sawatzky

JLAB Grand Challenge Related Projects.

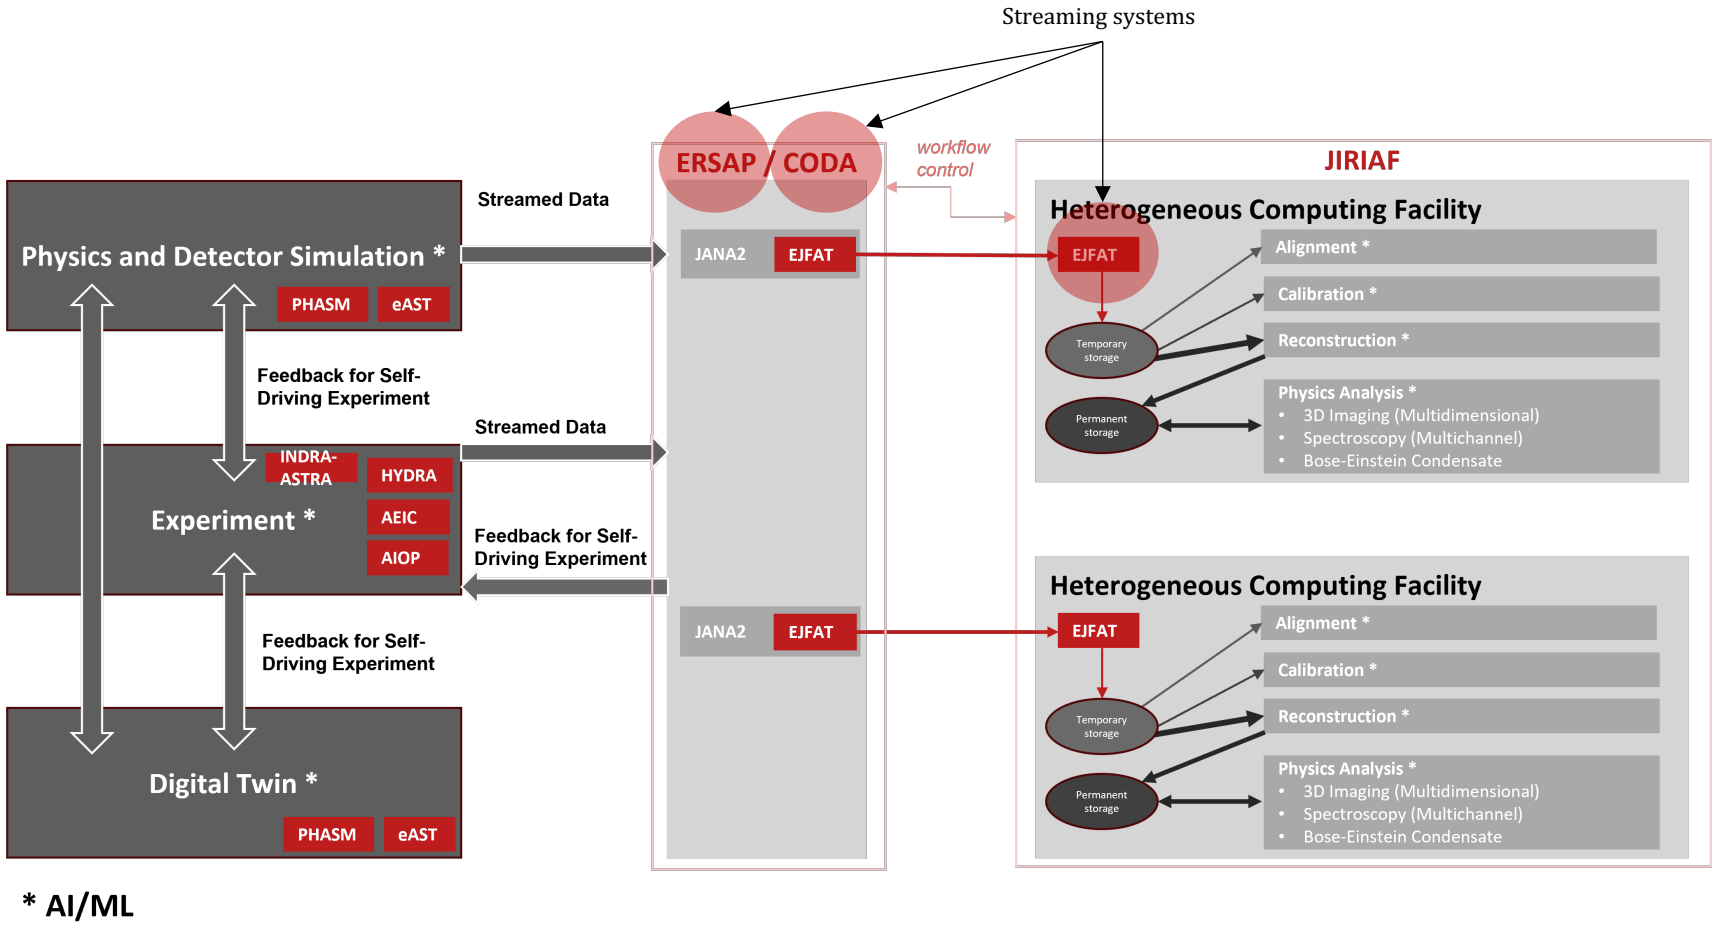

Slide courtesy of David Lawrence,

CEBAF Online Data Acquisition System: CODA

- CPU runs a software component ROC. It is responsible for payload board configuration and readout, as well as data formatting and passing it to the next stage.

- VTP relieves the ROC of all the "Readout" tasks and implements them in the FPGAs.
- Triggered or Streaming readout from ALL payload modules in parallel
- The Software ROC is now primarily responsible for configuring, controlling, and monitoring the VTP-based DAQ.

- TI Trigger interface card, responsible for trigger and clock distribution.

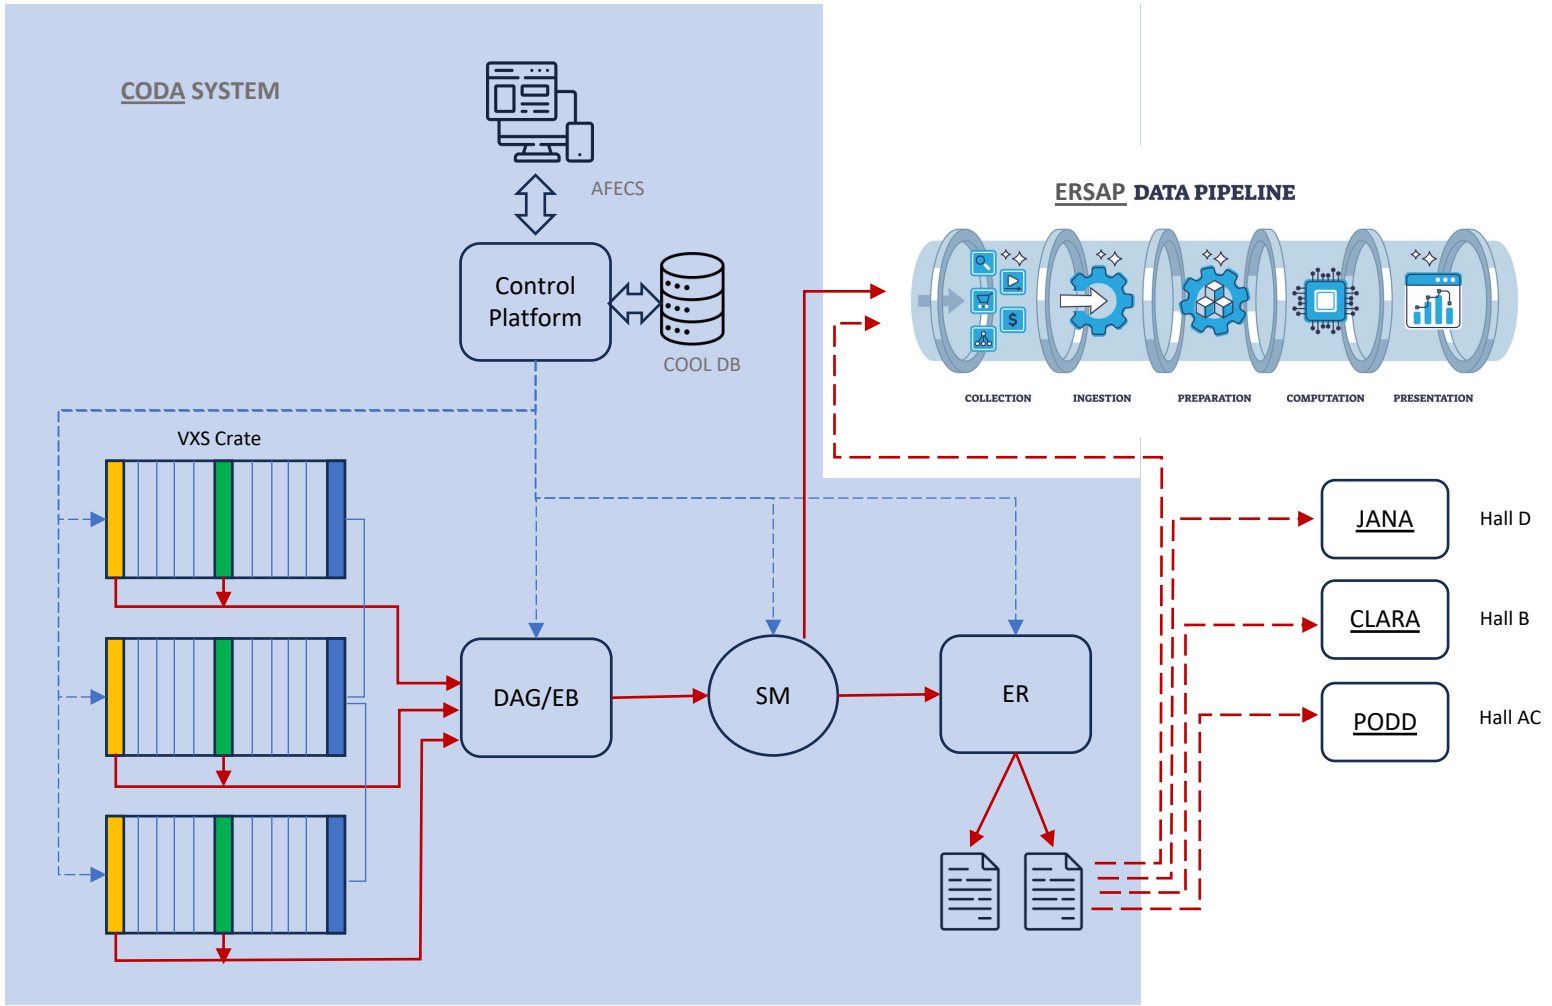

CODA Configuration: AFECS COOL

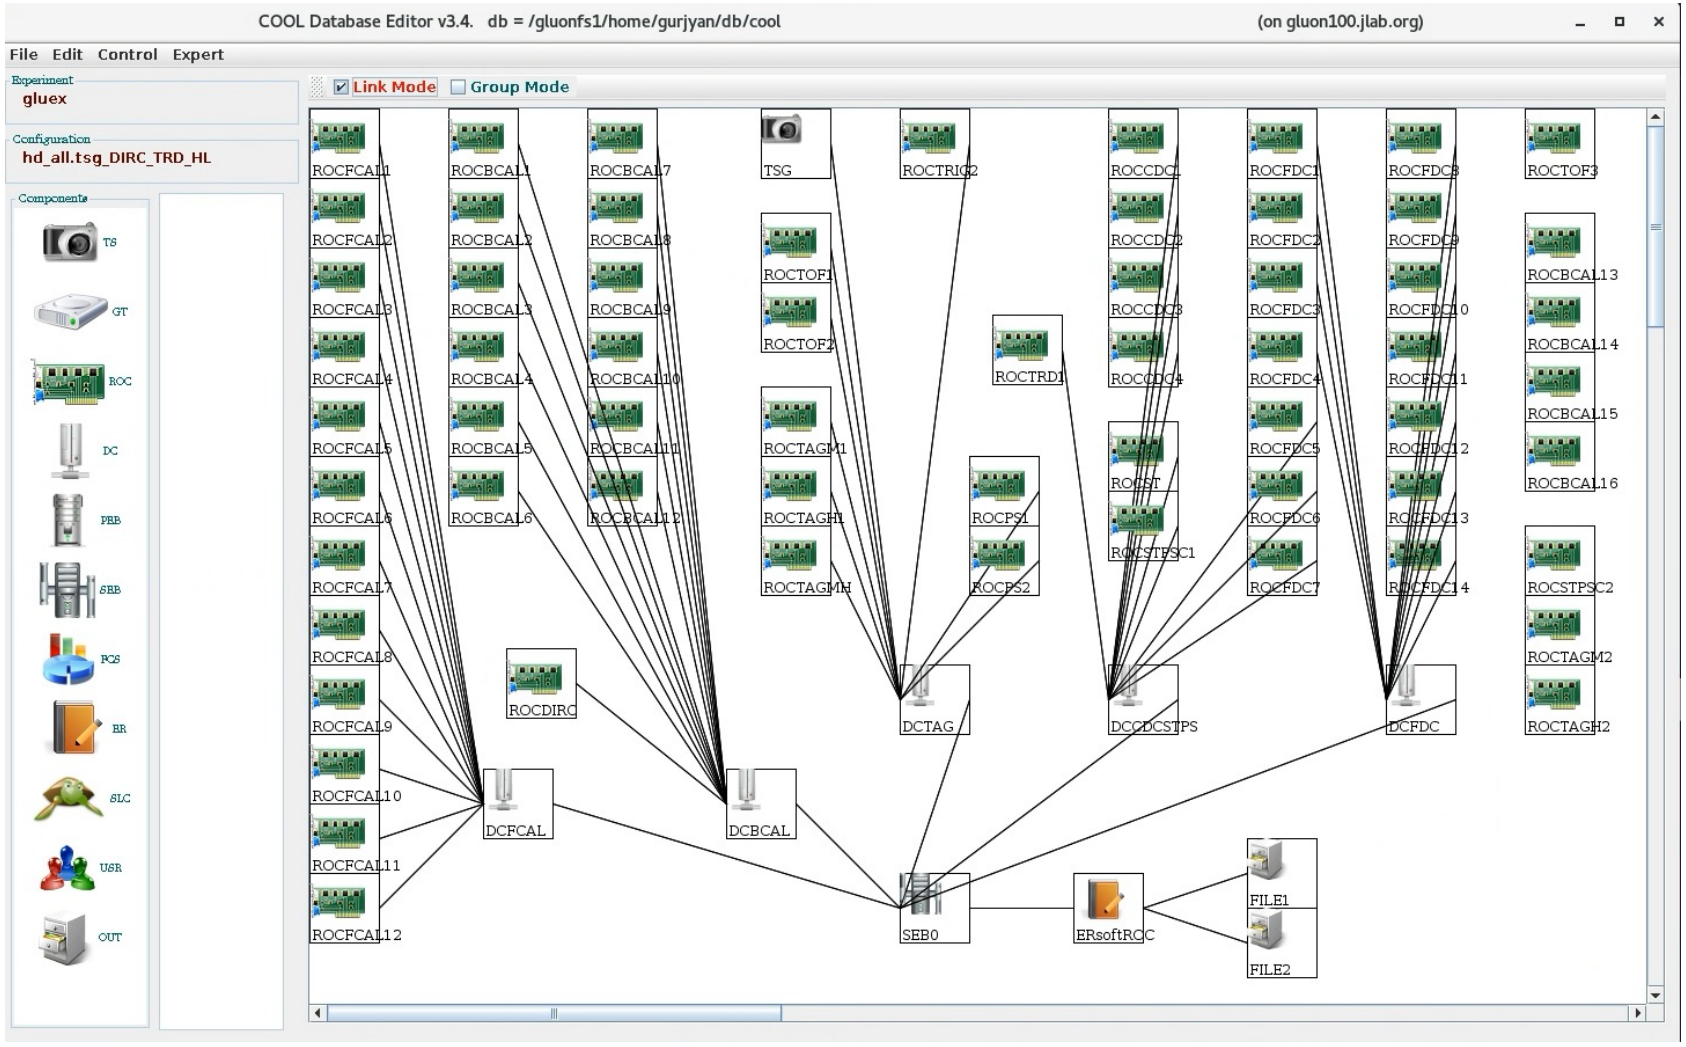

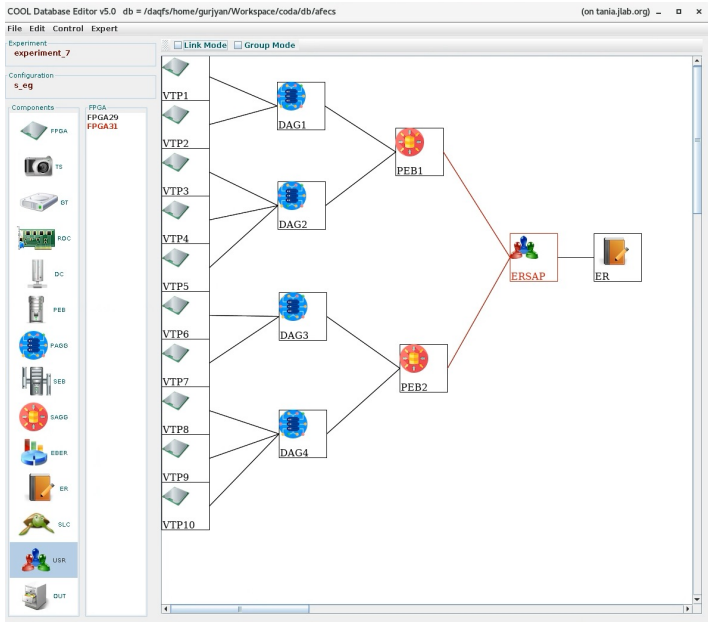

VXS Trigger Processor: VTP

Linux OS on the Zync-7030 SoC
(2-core ARM 7L , 1GB DDR3)
10/40Gbps Ethernet option
(runs the CODA ROC)

Xilinx Virtex 7 FPGA
Serial Lanes from both the VXS backplane and the Front panel
4GB DDR3 RAM

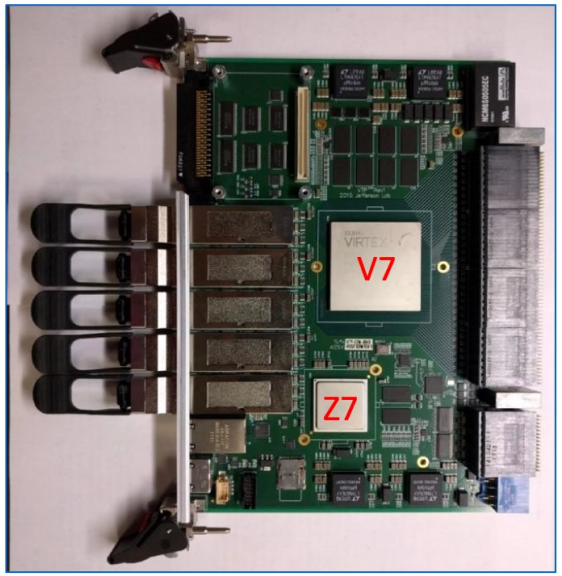

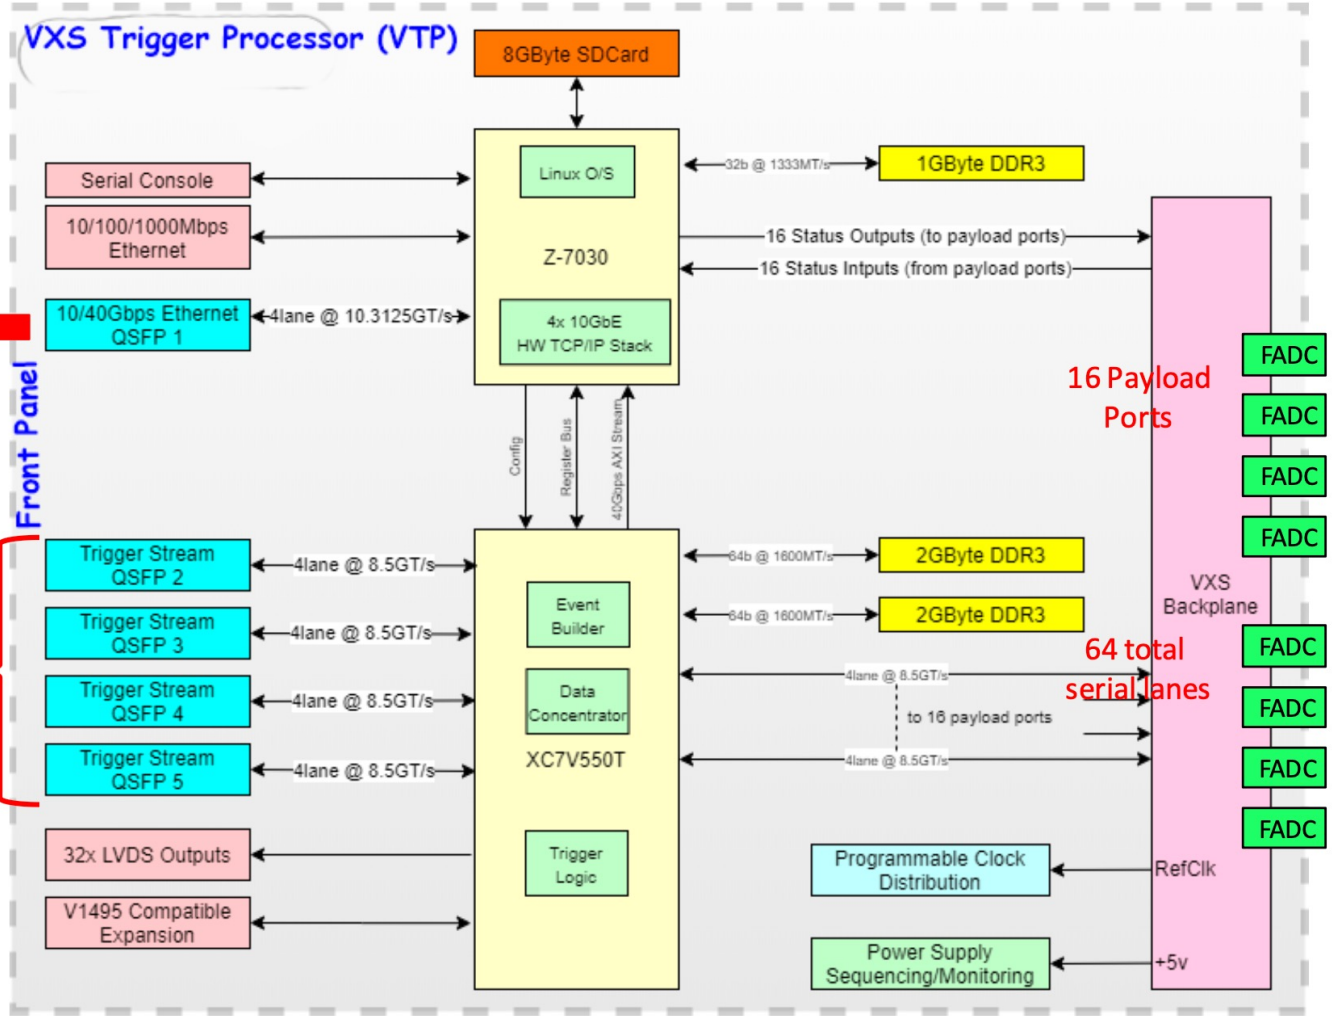

Slide courtesy of Ben Raydo, 23rd IEEE NPSS Real Time Conference August 1-5, 2022

Simultaneous Triggered and Stream Readout

VTP Trigger mode

At trigger arrival, we look back and read:

- FADC waveform values
- Threshold crossing times (hit time)
- Trigger absolute timestamp

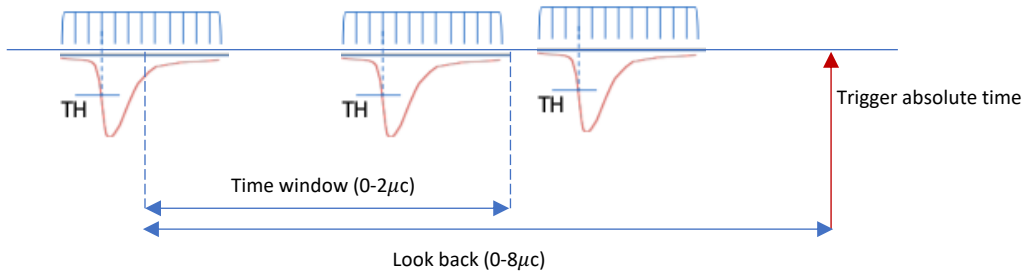

VTP Streaming mode

FPGA does pedestal subtracted sum over every hit over the threshold and reports:

- Threshold crossing: fine time stamp for each hit
- Aggregates all for N clock cycles (programmable up to 16bit, currently 65536 ns) into a frame and sends it with the absolute frame timestamp and frame number.

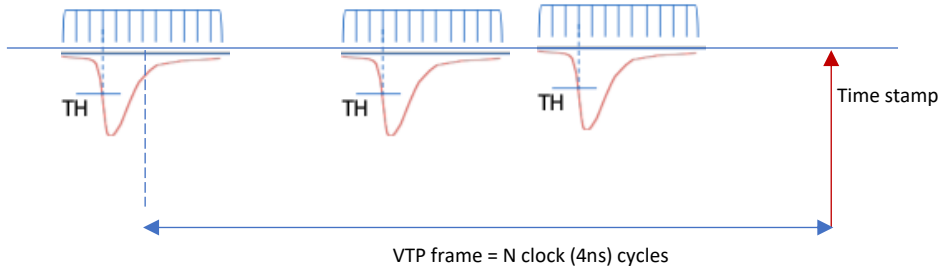

Drawing courtesy of Ben Raydo

FADC250

- It can work in triggered and streaming mode simultaneously. Generates a 12-bit sample every 4ns. That's 3 Gb/s for one channel. 16 channels is 48 Gb/s.
- Accepts external trigger and/or timestamp

Environment for Real-time Streaming, Acquisition, and Processing Framework: ERSAP

- Event reactive actors, networked by data pipelines.
- Compositional actors with conditional data routing at runtime.
- Flow-based programming paradigm

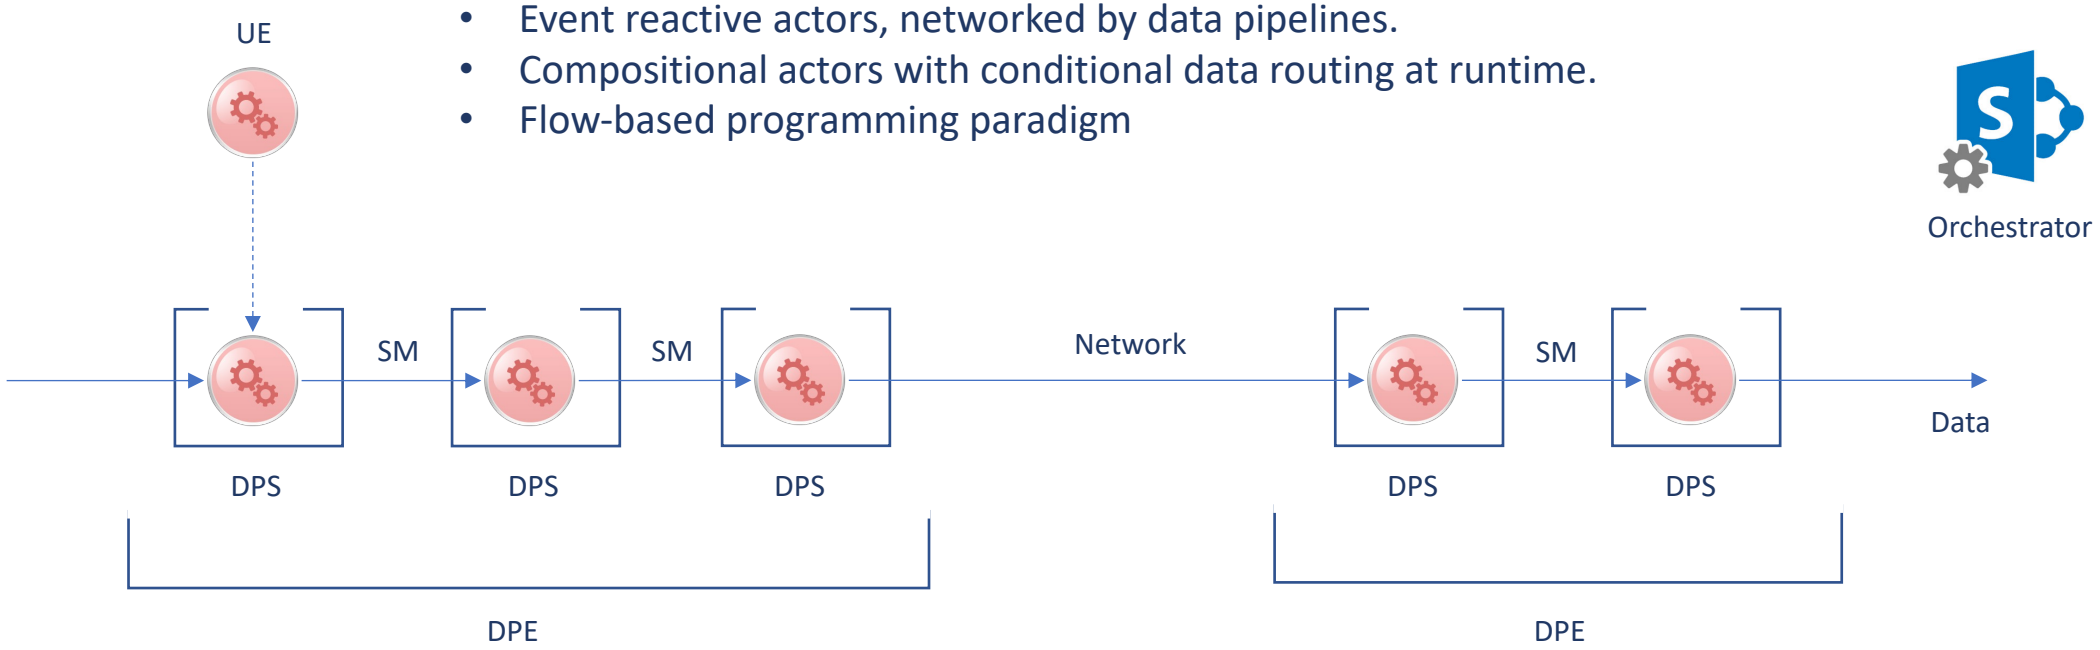

DPE : Data Processing Environment

SM : Shared Memory

DPS : Data Processing Station

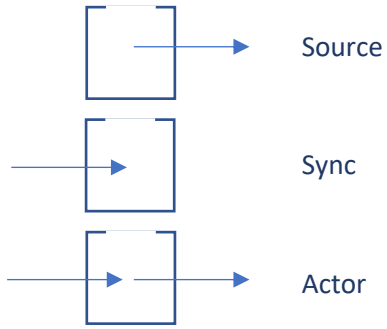

UE : User Engine

Init(Object O)

Object process(Object O)

ERSAP 3-layer structure

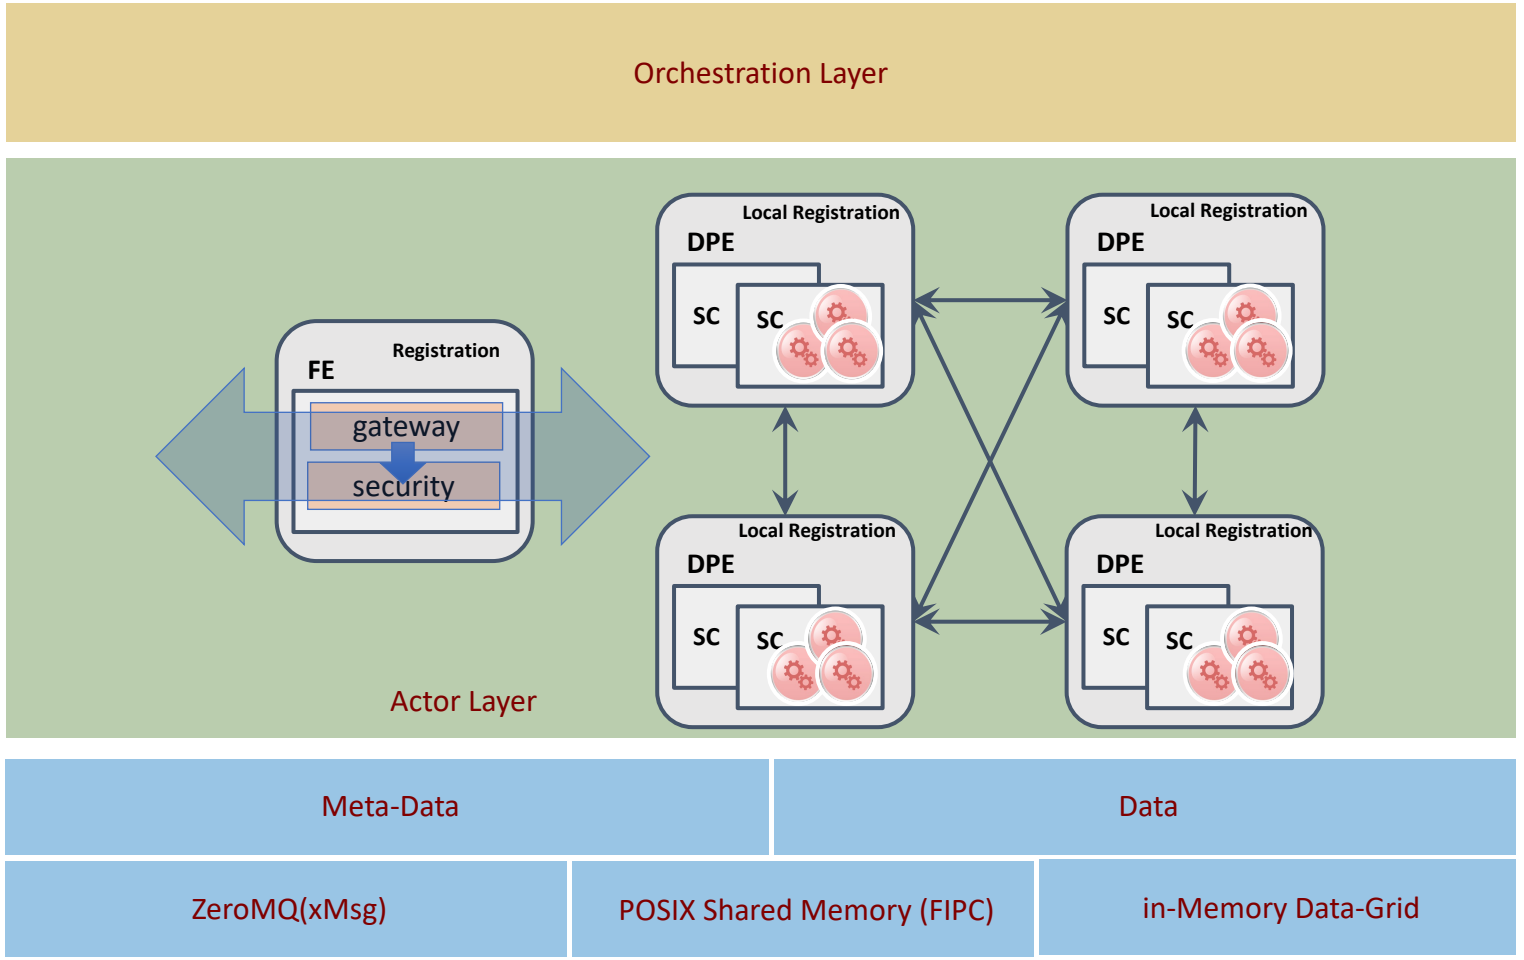

EIC prototype calorimeter SRO pipeline at DESY. CODA & ERSAP

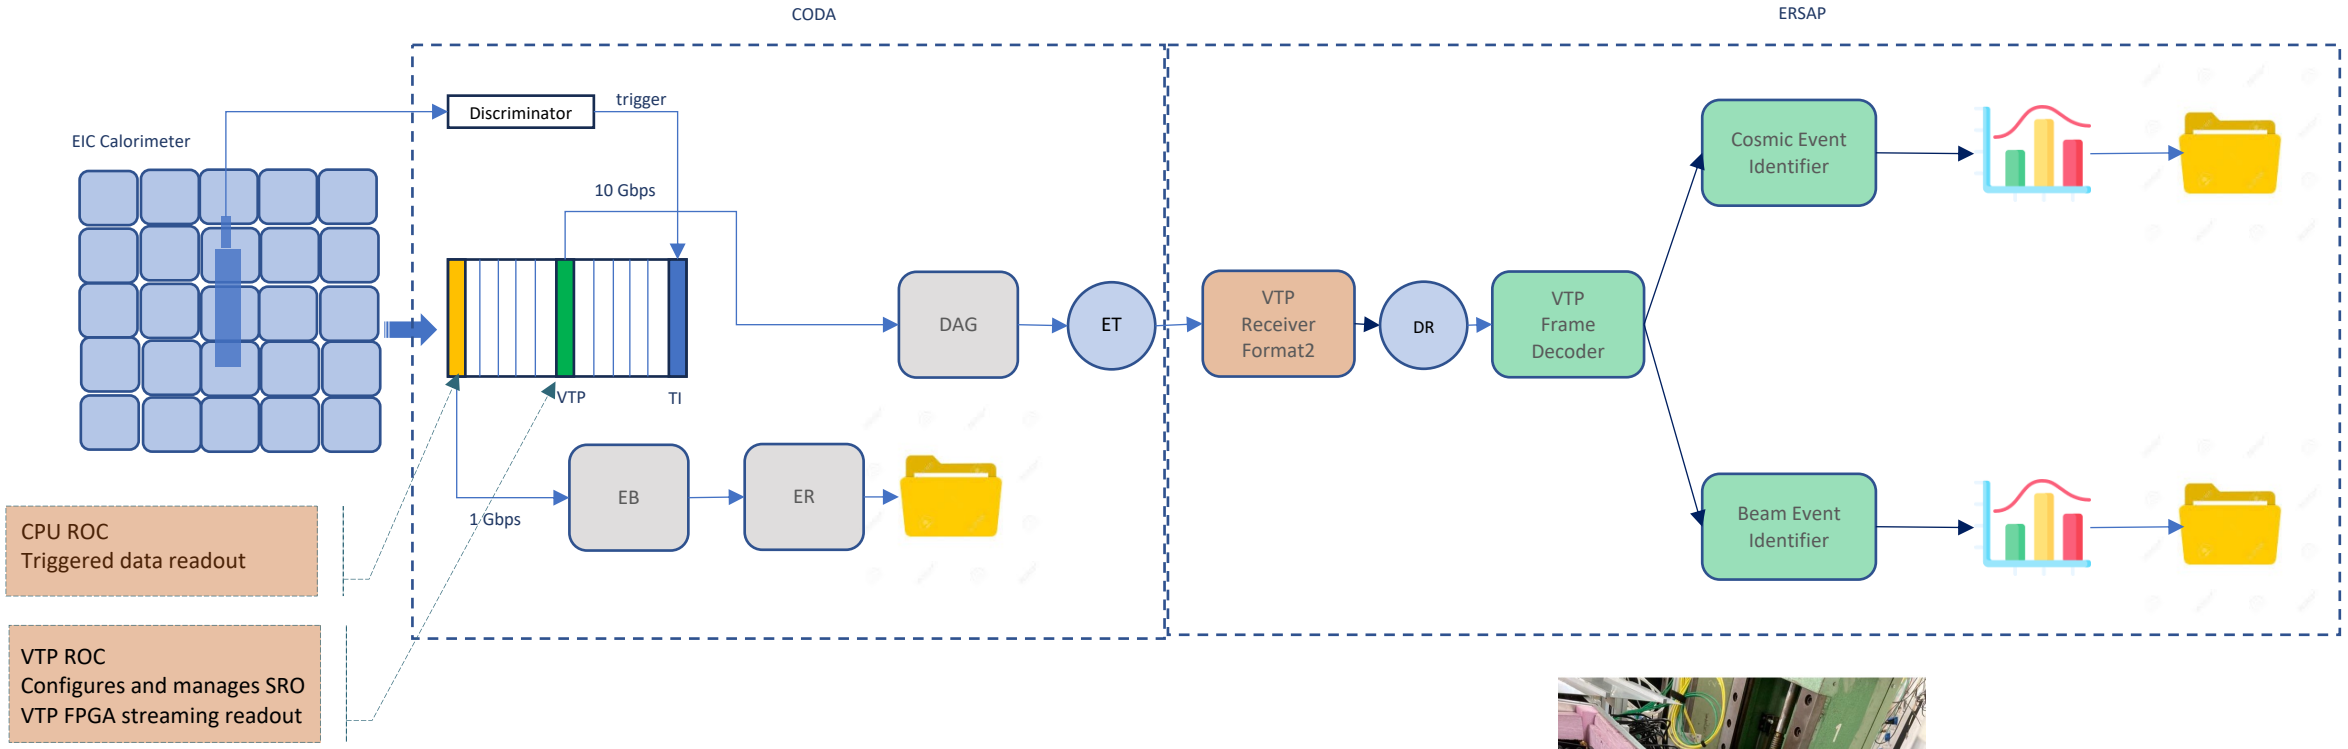

Triggered data are waveforms read out over the VME bus.
Stream data are integrated sums and times of all hits over a threshold in the calorimeter regardless of the trigger status.

Event Identification. Fixed-size sliding window technique

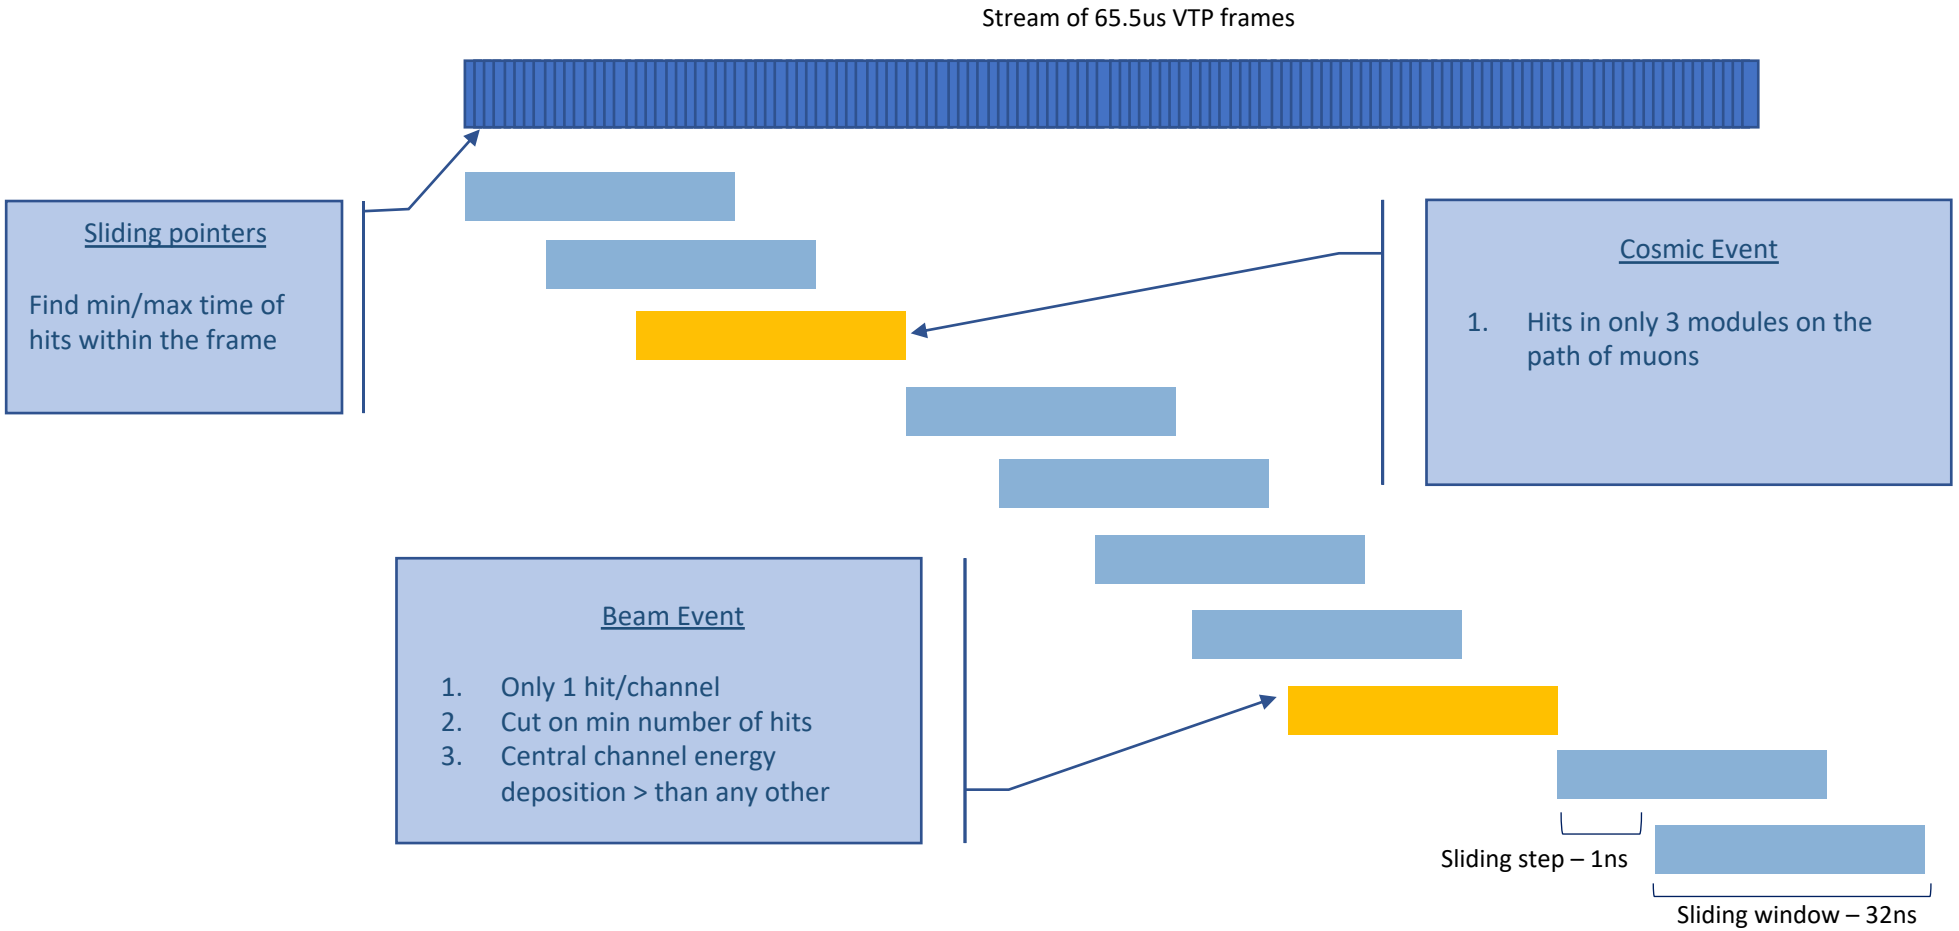

Adaptive Window Streaming Event Time-Based Clustering Algorithm

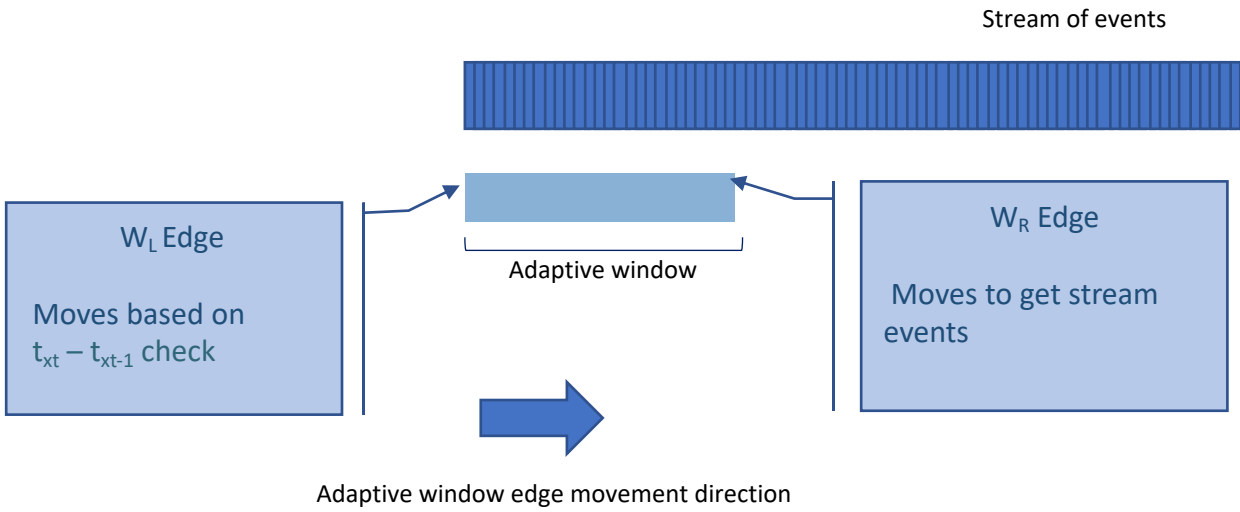


```
Initialize Window W<-W U {x1}
Loop {
  W←W ∪{xt}  // Move WR+1
  If ( | txt – txt-1 |) <= Th ) {
    cluster_event_count++
     channel[i++] = txt_channel
      if ( cluster_event_count >= Nn
            &&
           no duplicates in channel[]) {
          cluster is identified
          move WL+cluster_event_count
          cluster_event_count = 0
           clear channel[]
       }
   } else {
     move WL+cluster_event_count
     cluster_event_count = 0
     clear channel[]
   } }
```

Streaming Data in Real-Time

Raw data spectrum

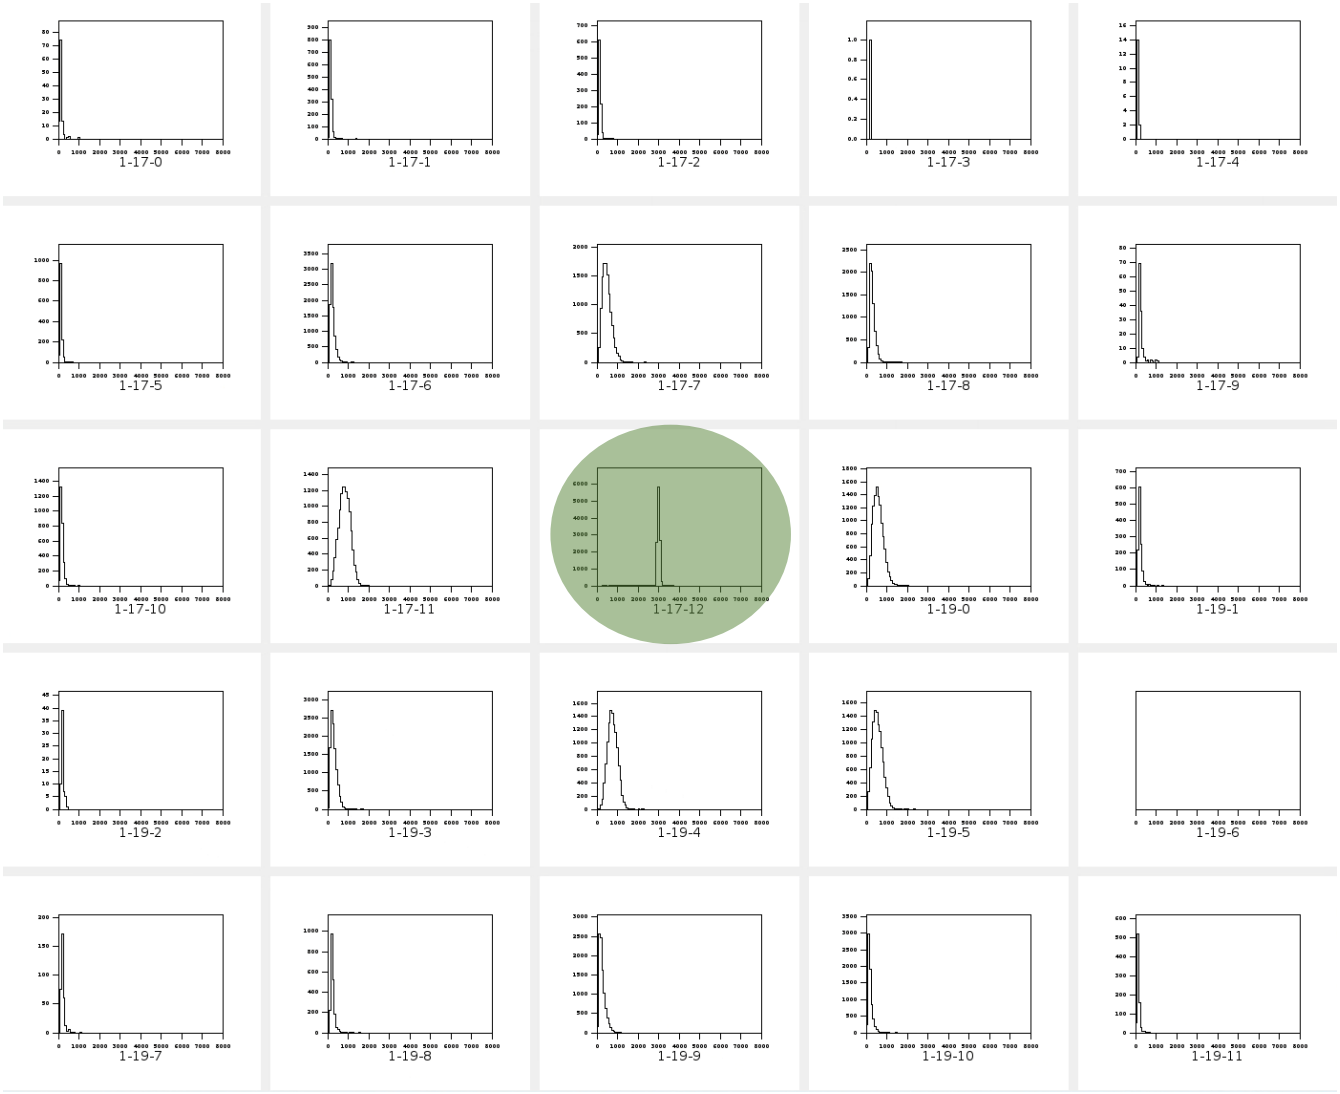

Hit Identification

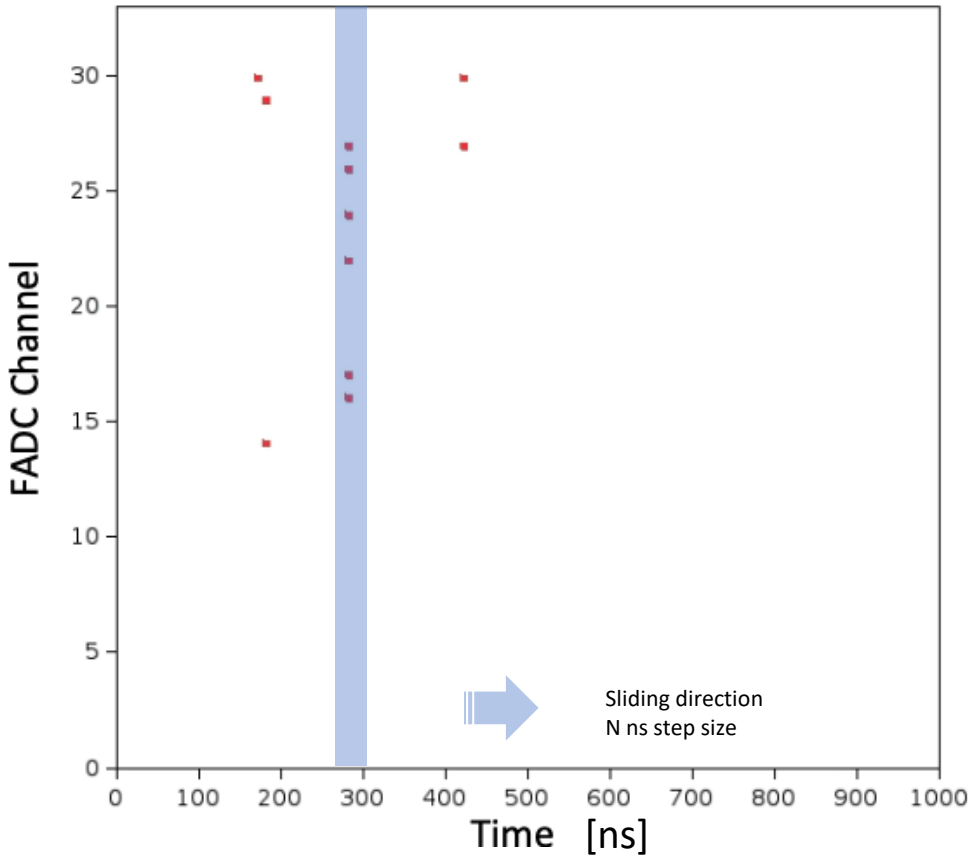

EPSCI Jefferson Lab

Identified Beam Events

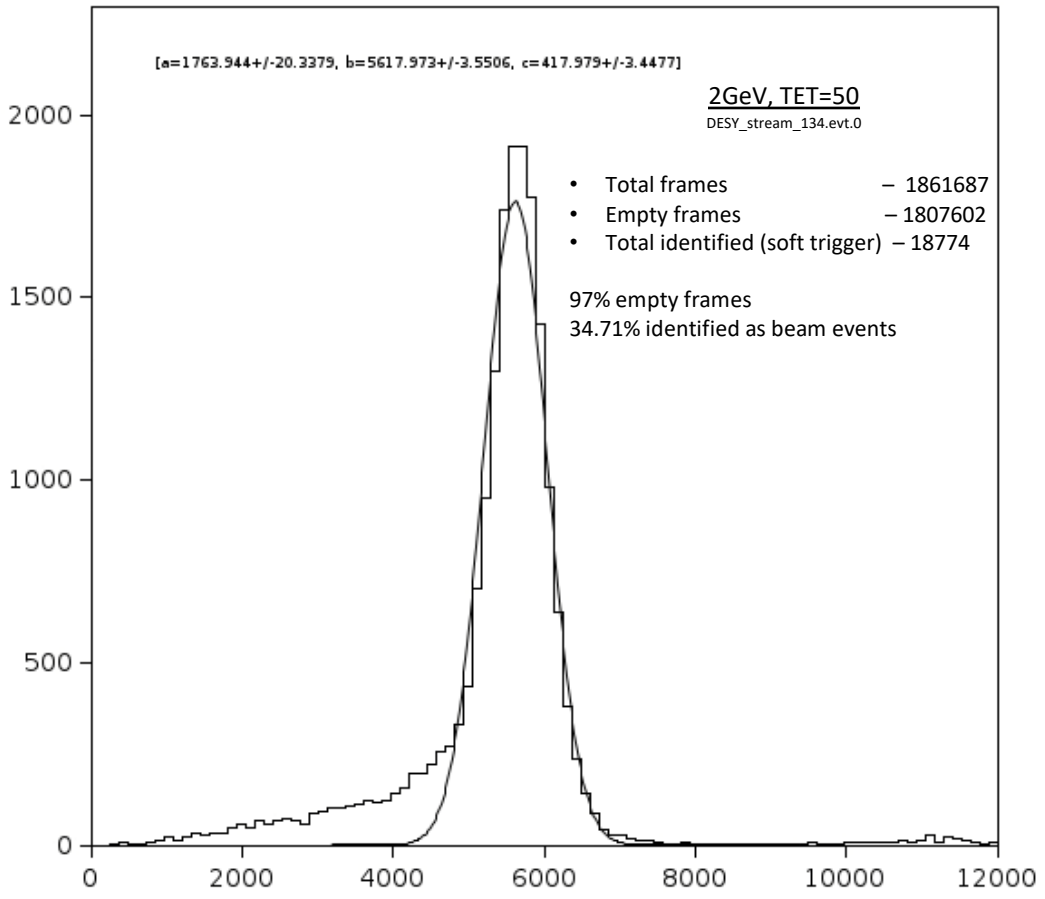

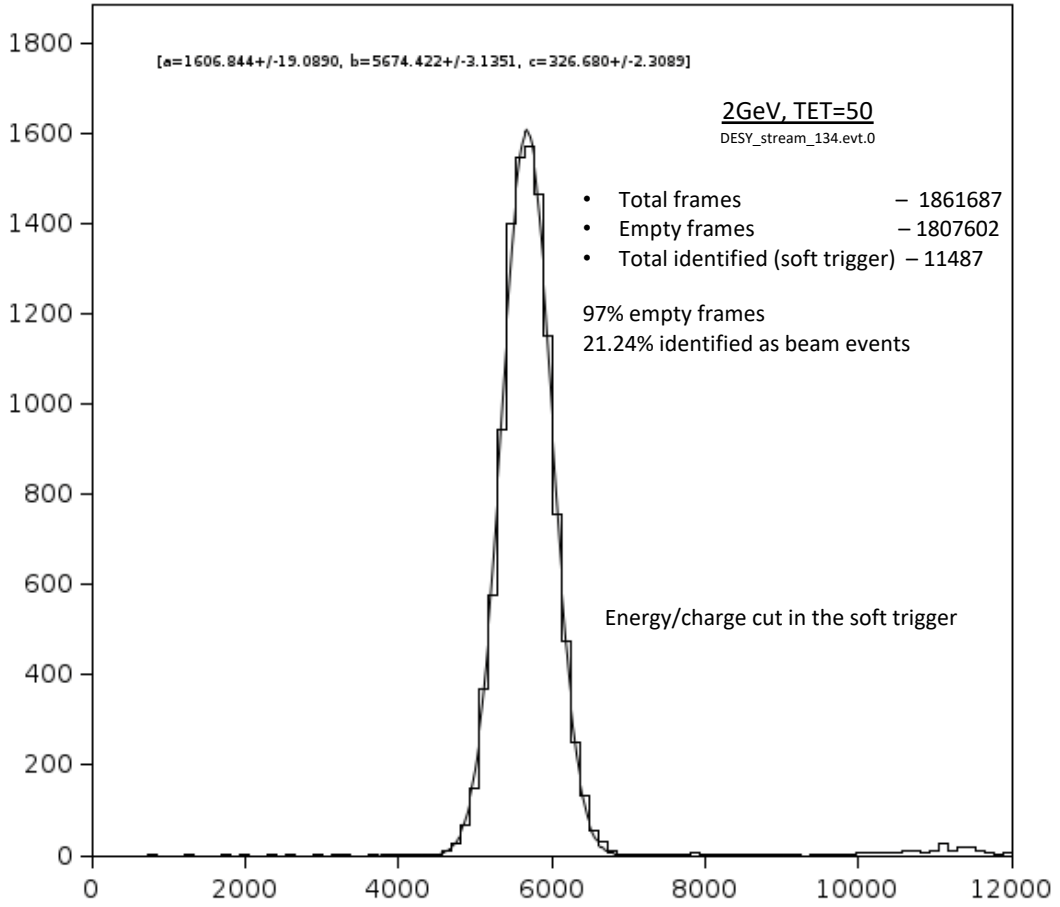

Stream Identification Efficiency

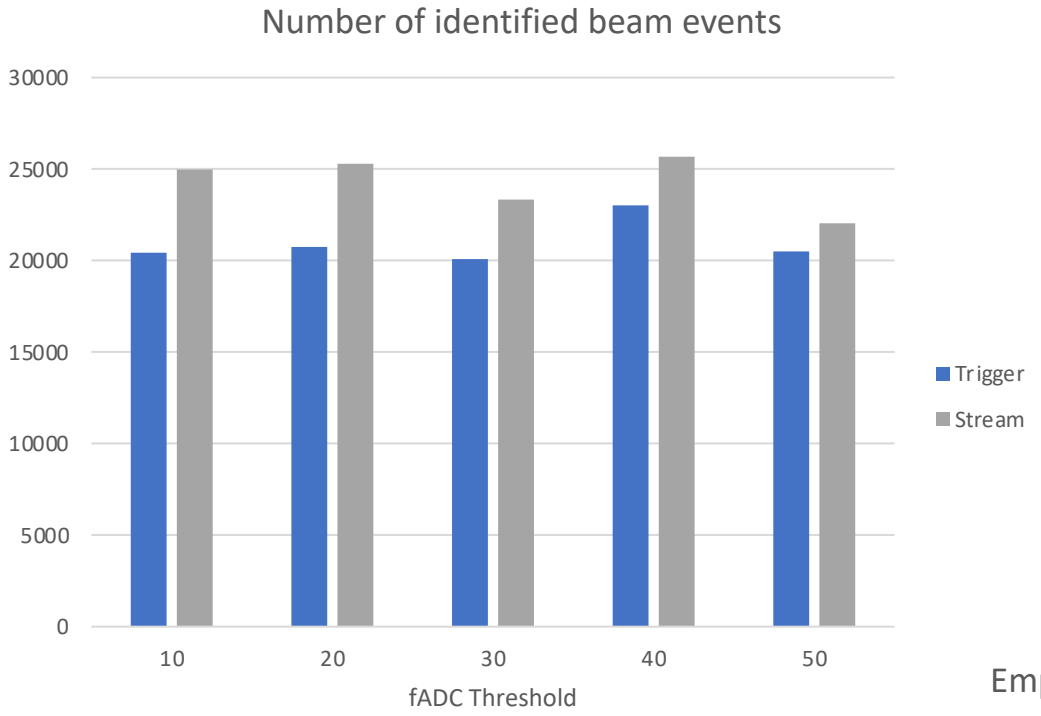

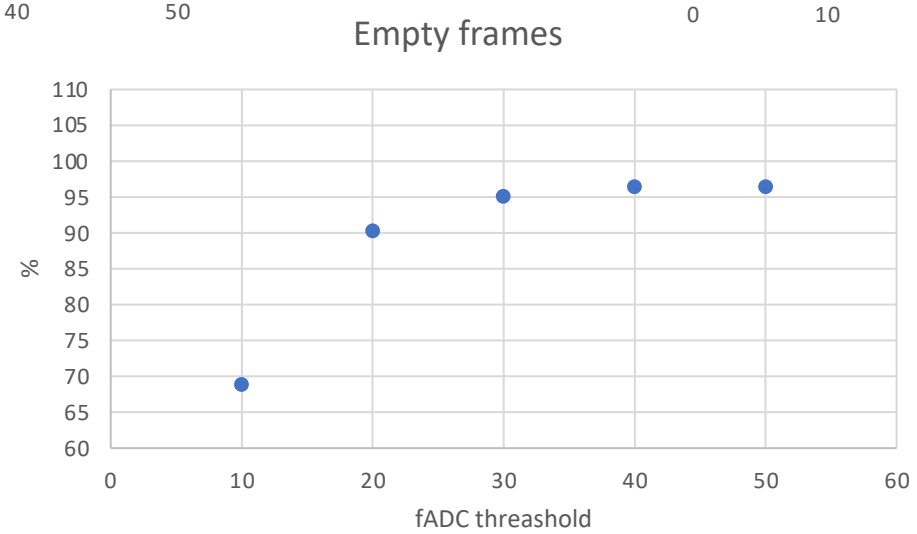

Calorimeter Response Linearity

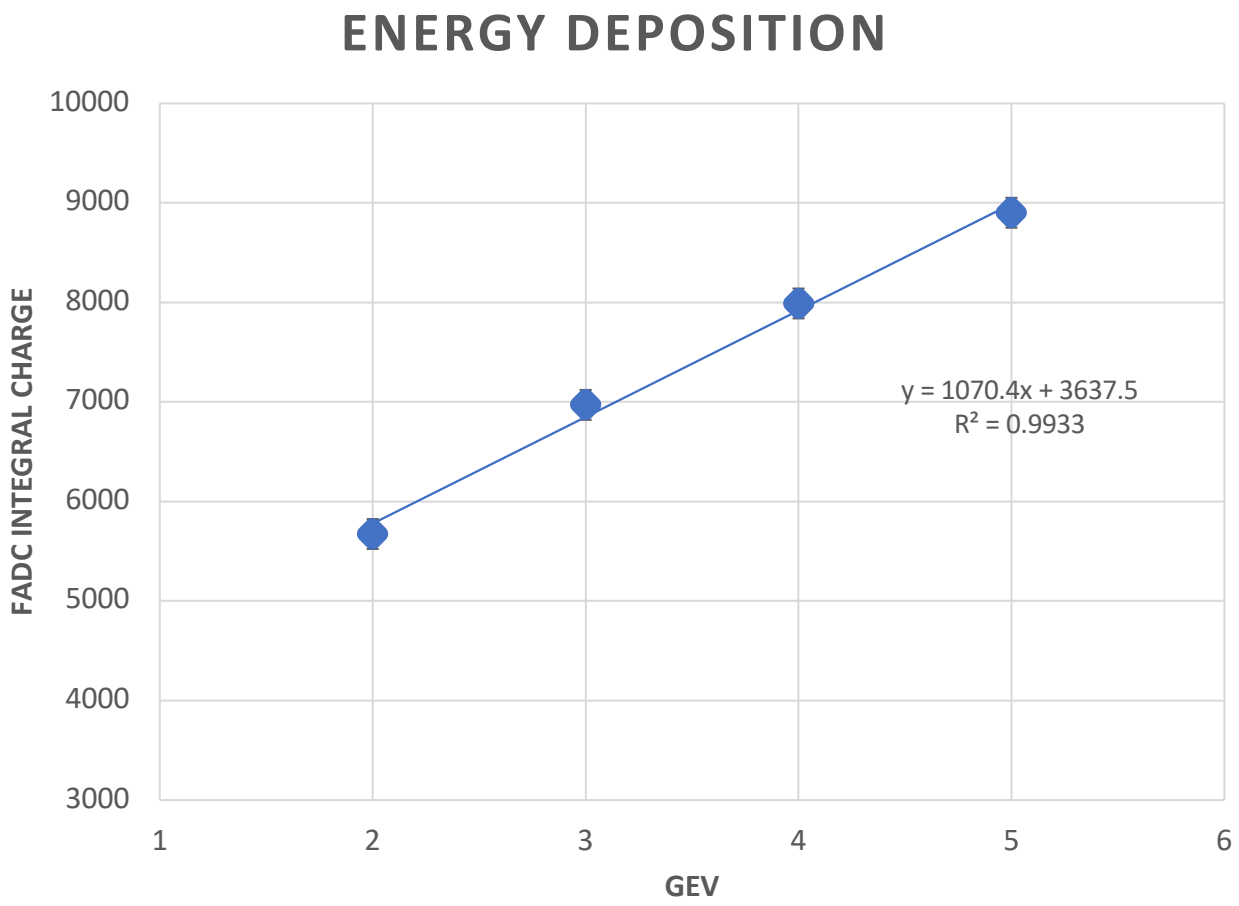

ESnet FPGA Accelerated Transport System: EJFAT

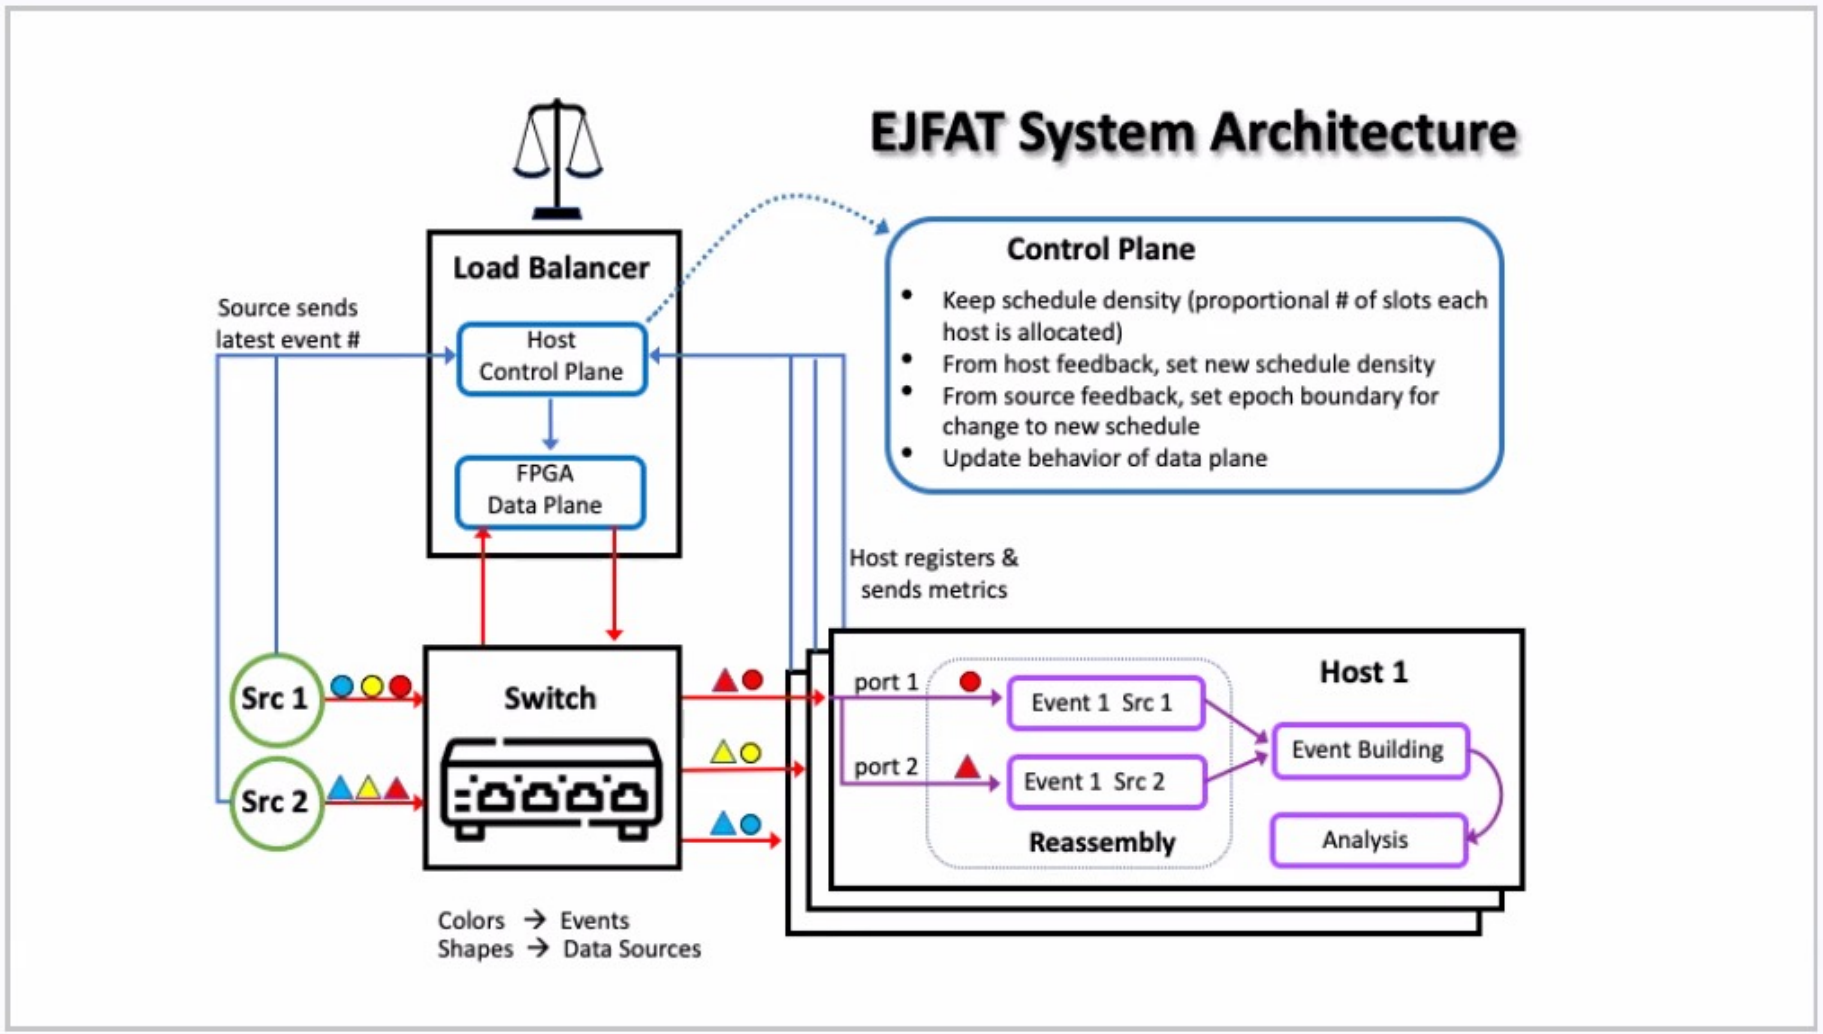

Slide courtesy of Mike Goodrich

CLAS12 Raw (triggered) Data-Stream Processing. EJFAT & ERSAP

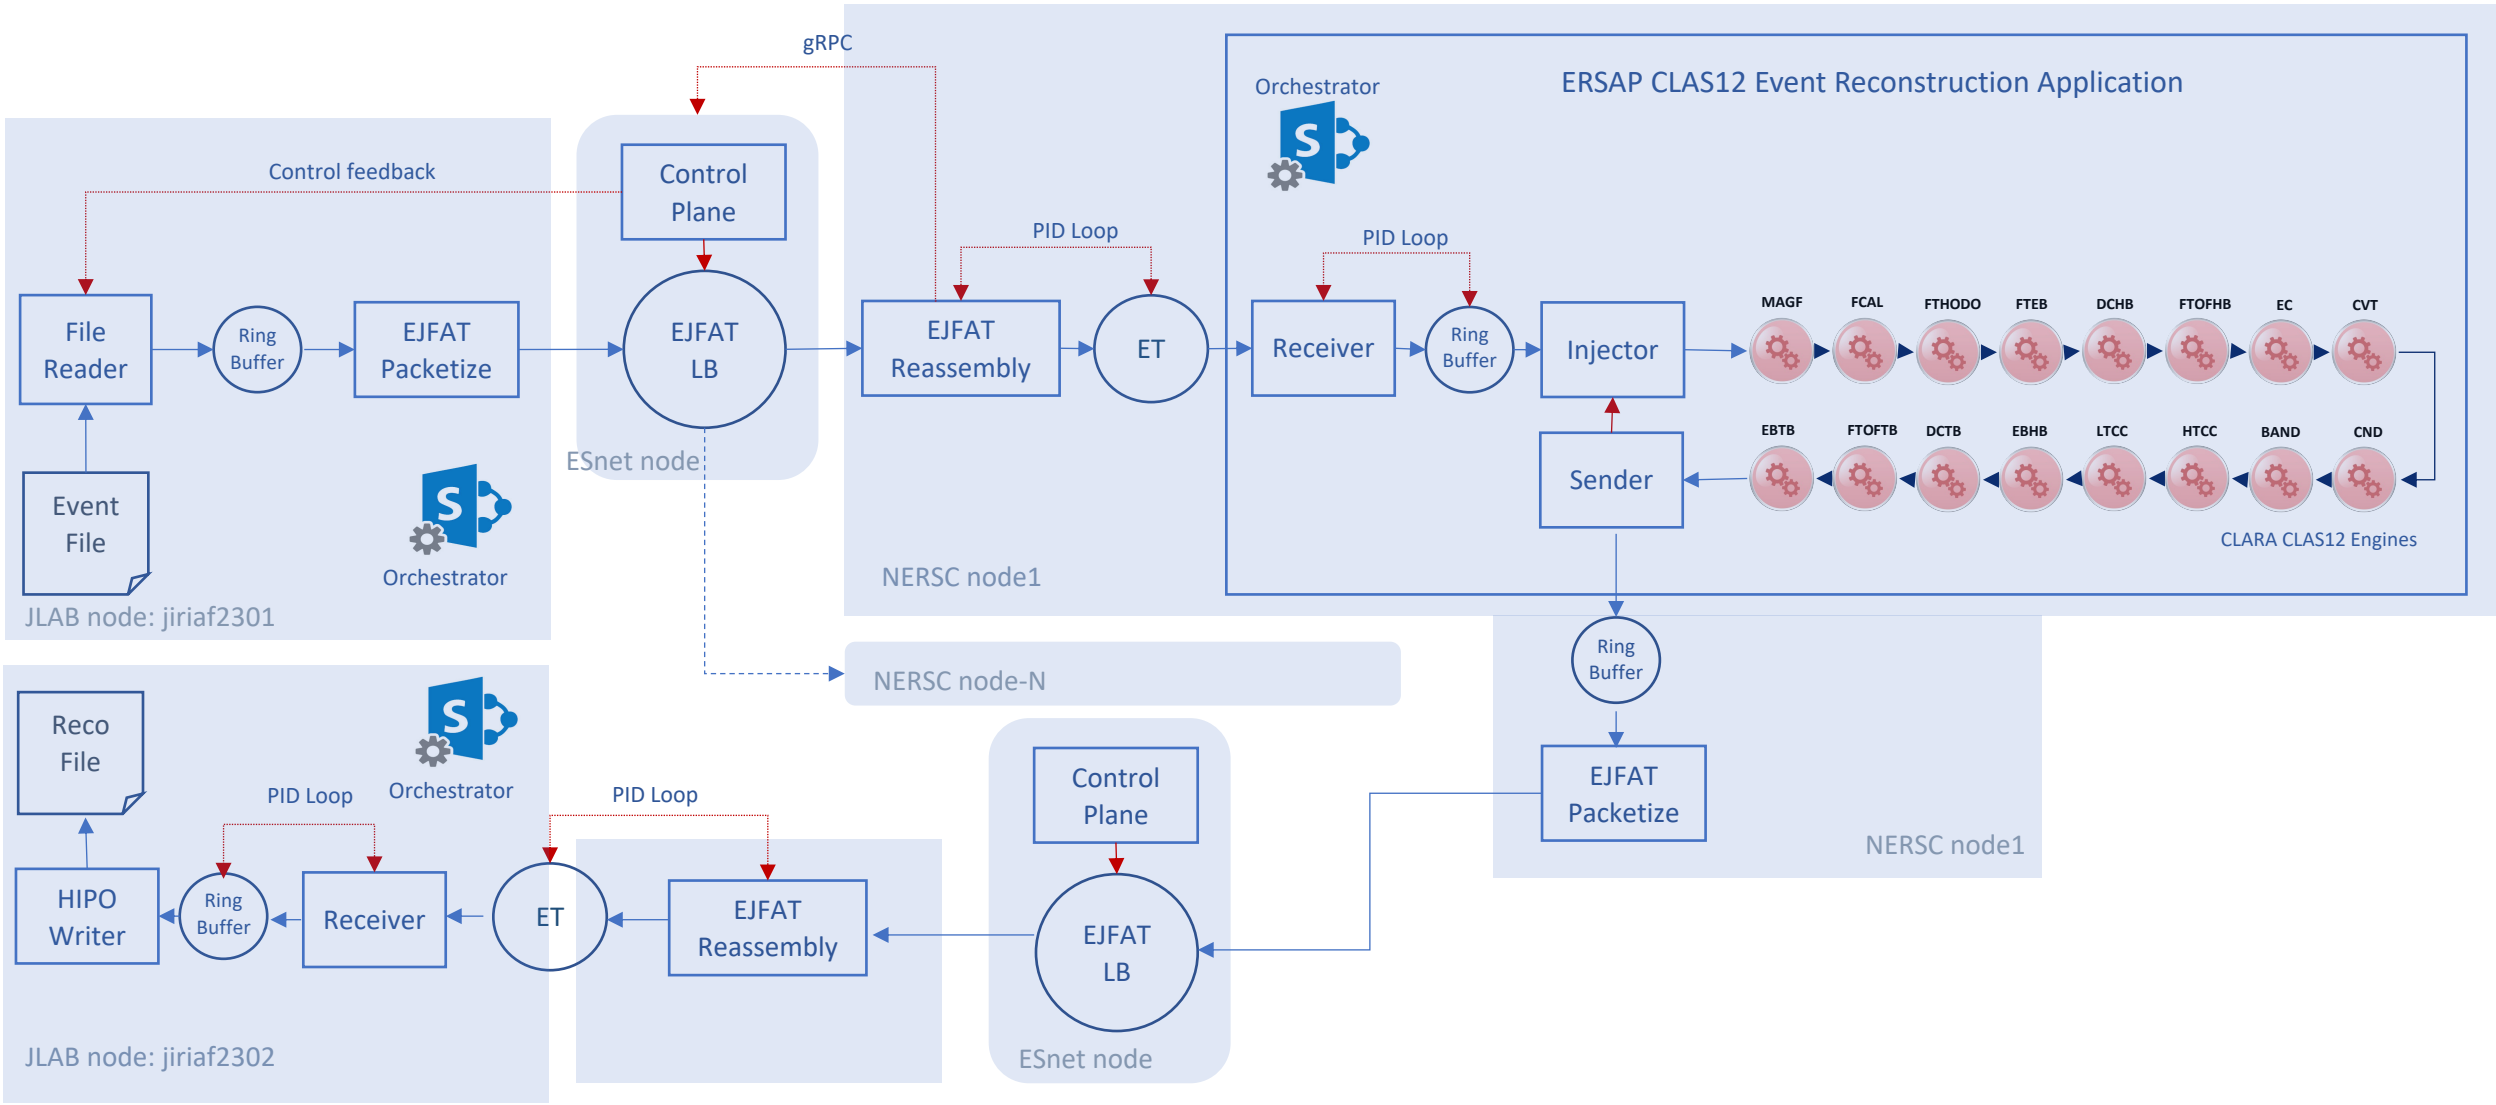

CLAS12 Stream Event Reconstruction: JLAB – Esnet - NERSC

JLAB gurjyan@tania:~

```
Events:        476.2 Hz,  475.6 Avg, total 9788220

Packets:       1780 Hz,    1771 Avg, time: diff = 4000138 usec, abs = 1328662276 epoch msec
Data (+hdrs):  13.8 (13.83) MB/s,  13.69 (13.73) Avg
Events:        475 Hz,  475.6 Avg, total 9790120

Packets:       1754 Hz,    1771 Avg, time: diff = 4000120 usec, abs = 1328666276 epoch msec
Data (+hdrs):  13.6 (13.64) MB/s,  13.69 (13.73) Avg
Events:        475.5 Hz,  475.6 Avg, total 9792022

Packets:       1774 Hz,    1771 Avg, time: diff = 4000119 usec, abs = 1328670276 epoch msec
Data (+hdrs):  13.73 (13.76) MB/s,  13.69 (13.73) Avg
Events:        476 Hz,  475.6 Avg, total 9793926

Packets:       1788 Hz,    1771 Avg, time: diff = 4000119 usec, abs = 1328674276 epoch msec
Data (+hdrs):  13.82 (13.85) MB/s,  13.69 (13.73) Avg
Events:        475 Hz,  475.6 Avg, total 9795826
```

NERSC gurjyan@login08:~

```
Events:        475.4 Hz,  313.5 Avg, total 96574
Dropped: evts: 0, 0 total, pkts: 0, 0 total

2023-11-09 17:44:58.951: Processed  500 events in   9.81 s   average event time =  19.62
2023-11-09 17:44:59.166: Processed  500 events in  10.83 s   average event time =  21.67
     Fifo level 950  Avg:  818.27,  86.13%,  pid err -1.299115

2023-11-09 17:45:00.215: Processed  500 events in   9.25 s   average event time =  18.50
Packets:       1759 Hz,    1175 Avg, time: diff = 4000100 usec, abs = 1561582856 epoch ms
Data (+hdrs):  13.67 (13.7) MB/s,  9.087 (9.11) Avg
Events:        476.2 Hz,  315.6 Avg, total 98479
Dropped: evts: 0, 0 total, pkts: 0, 0 total

     Fifo level 826  Avg:  920.73,  96.92%,  pid err -1.380536

2023-11-09 17:45:05.476: Processed  500 events in   9.22 s   average event time =  18.43
Packets:       186.5 Hz,    1163 Avg, time: diff = 4000110 usec, abs = 1561586856 epoch m
Data (+hdrs):  1.452 (1.456) MB/s,  8.99 (9.013) Avg
Events:        50.25 Hz,  312.2 Avg, total 98680
Dropped: evts: 0, 0 total, pkts: 0, 0 total

2023-11-09 17:45:07.785: Processed  500 events in   8.83 s   average event time =  17.67 ms   [ total 12000 events   292.43 s ]
2023-11-09 17:45:08.021: Processed  500 events in   8.86 s   average event time =  17.71 ms   [ total 10000 events   289.94 s ]
2023-11-09 17:45:09.108: Processed  500 events in   8.89 s   average event time =  17.79 ms   [ total 10500 events   290.89 s ]
     Fifo level 0  Avg:  0.04,  0.00%,  pid err -0.891927

Packets:       1018 Hz,    1161 Avg, time: diff = 4000159 usec, abs = 1561590856 epoch msec, cpu = 83
Data (+hdrs):  7.86 (7.88) MB/s,  8.976 (8.999) Avg
Events:        270 Hz,  311.7 Avg, total 99760
Dropped: evts: 1, 1 total, pkts: 4, 4 total
```

c (on ejfat-2)

File View Studio

EC Engine Monitoring

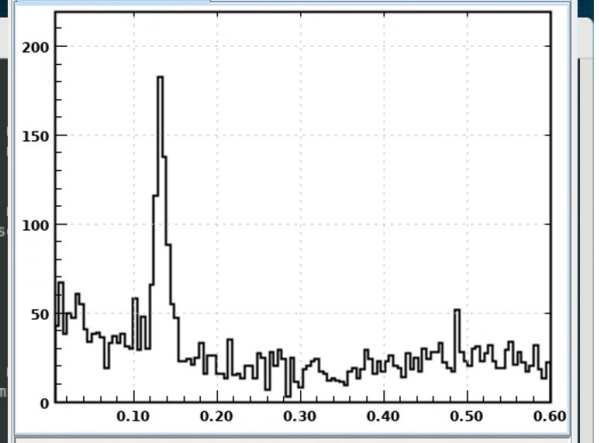

Status:

JIRIAF Dashboard at JLAB

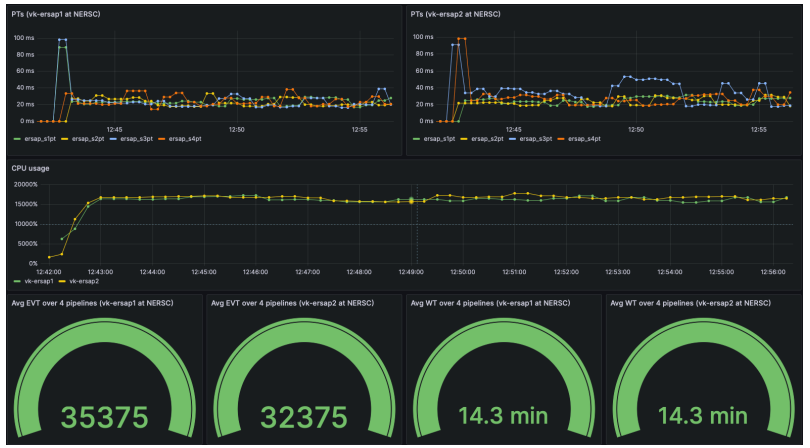

EPSCI Jefferson Lab

Summary

- Described "Grand Challenge" lab-wide initiative focused on streaming readout and data stream processing
- Introduced the JLAB common data acquisition system (CODA) upgrade to support streaming
  - CODA is now capable of triggering and streaming readout simultaneously (hybrid DAQ)
- Introduced ERSAP data-stream processing framework inspired by CODA's data-centric approach
  - ERSAP uses data application building blocks like LEGO bricks in a flow-based programming paradigm
  - Advantages of this approach include fault isolation, hardware and software heterogeneity support, and easy software evolution
- Presented results from streaming readout and data stream processing, including the EIC prototype calorimeter test at DESY and CLAS12 event reconstruction from streaming data at a distance.

Thank You

Raw data spectrum

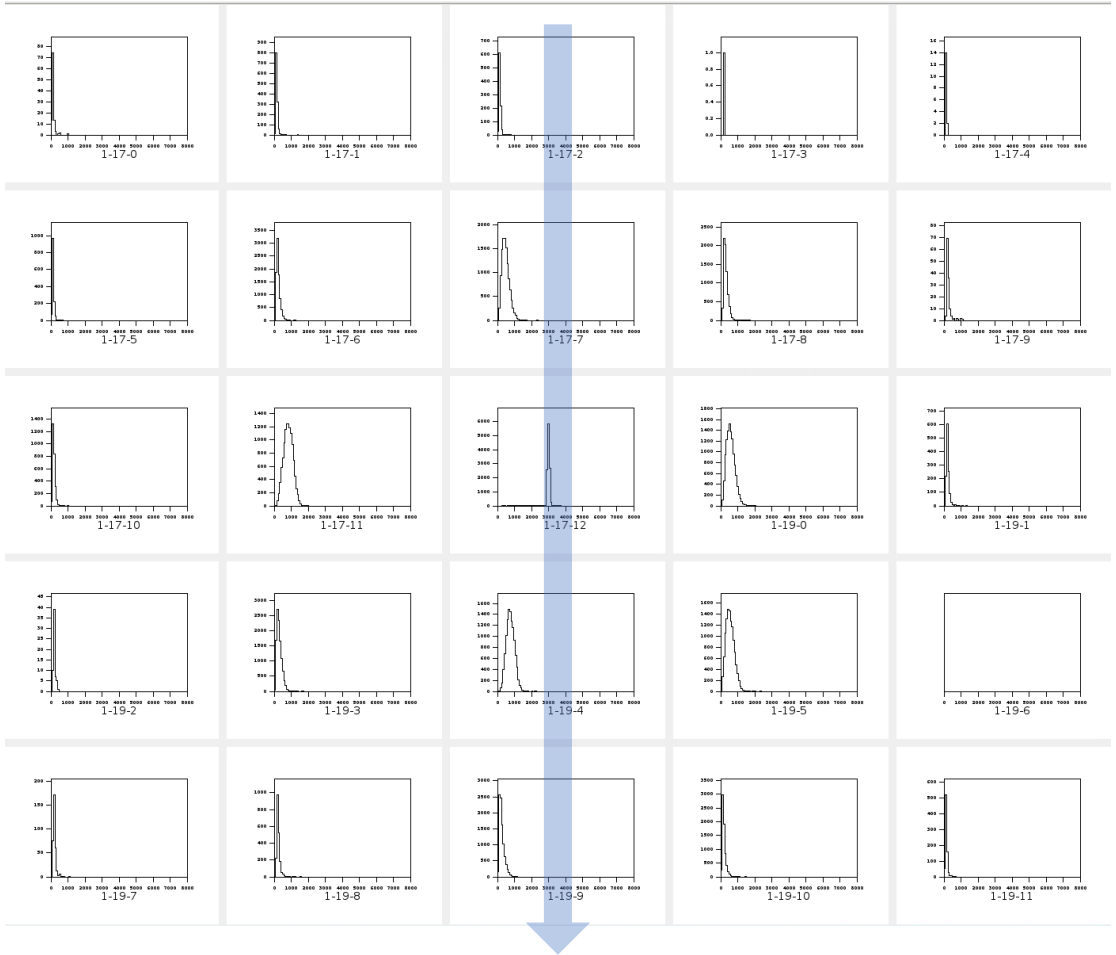

Cosmic Identification

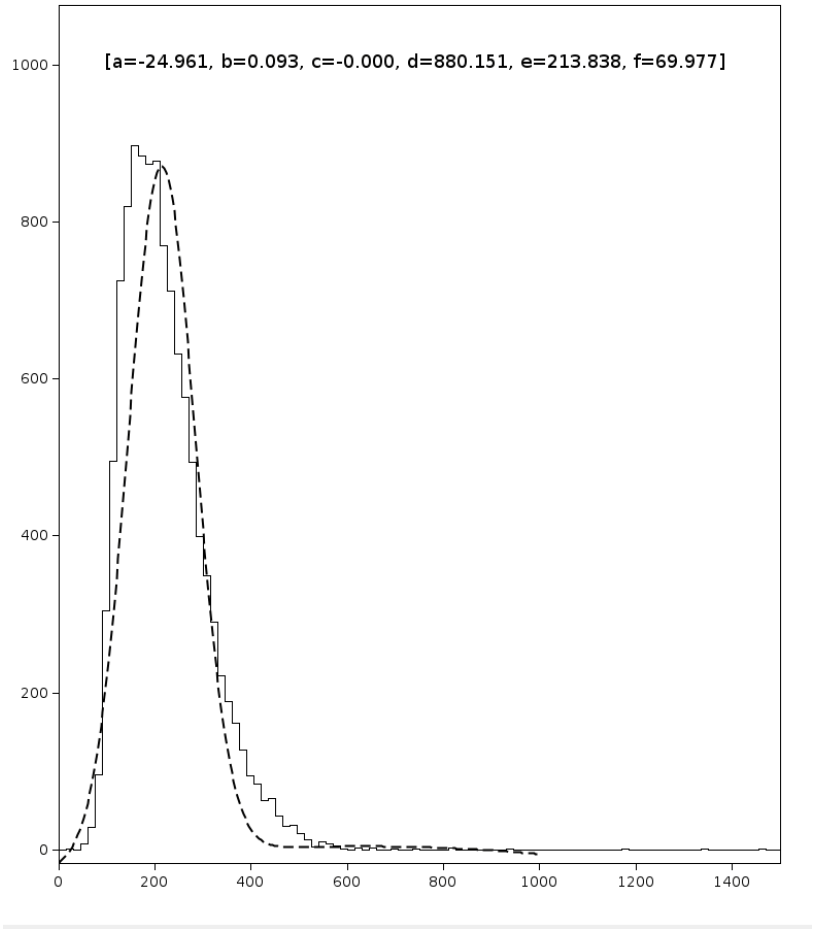

Data processing station: actor

- User *engine* run-time environment.
- Engine follows data-in/data-out interface.
- Engine gets JSON object for run-time configuration.

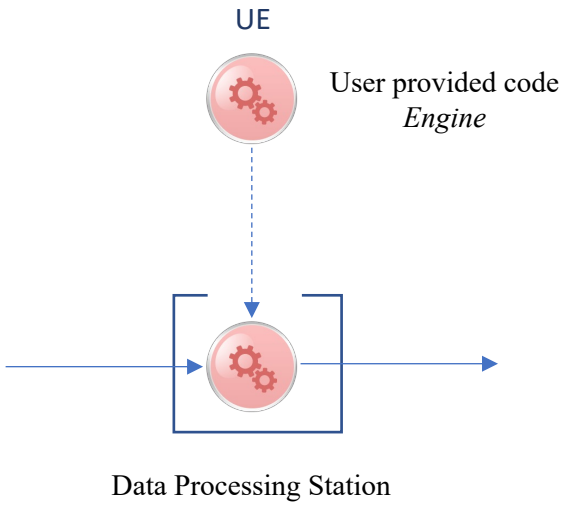

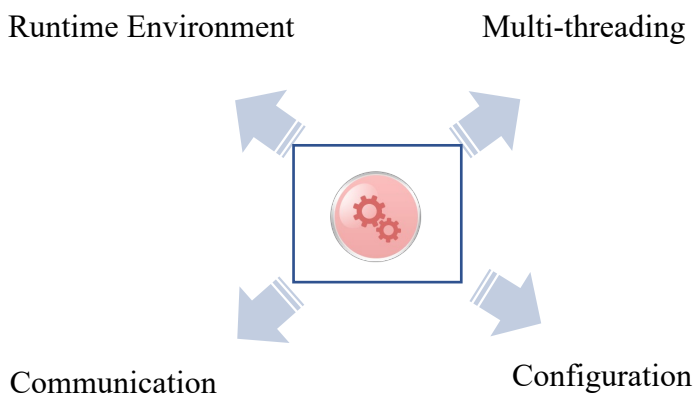

Streaming data transport (conveyer belt)

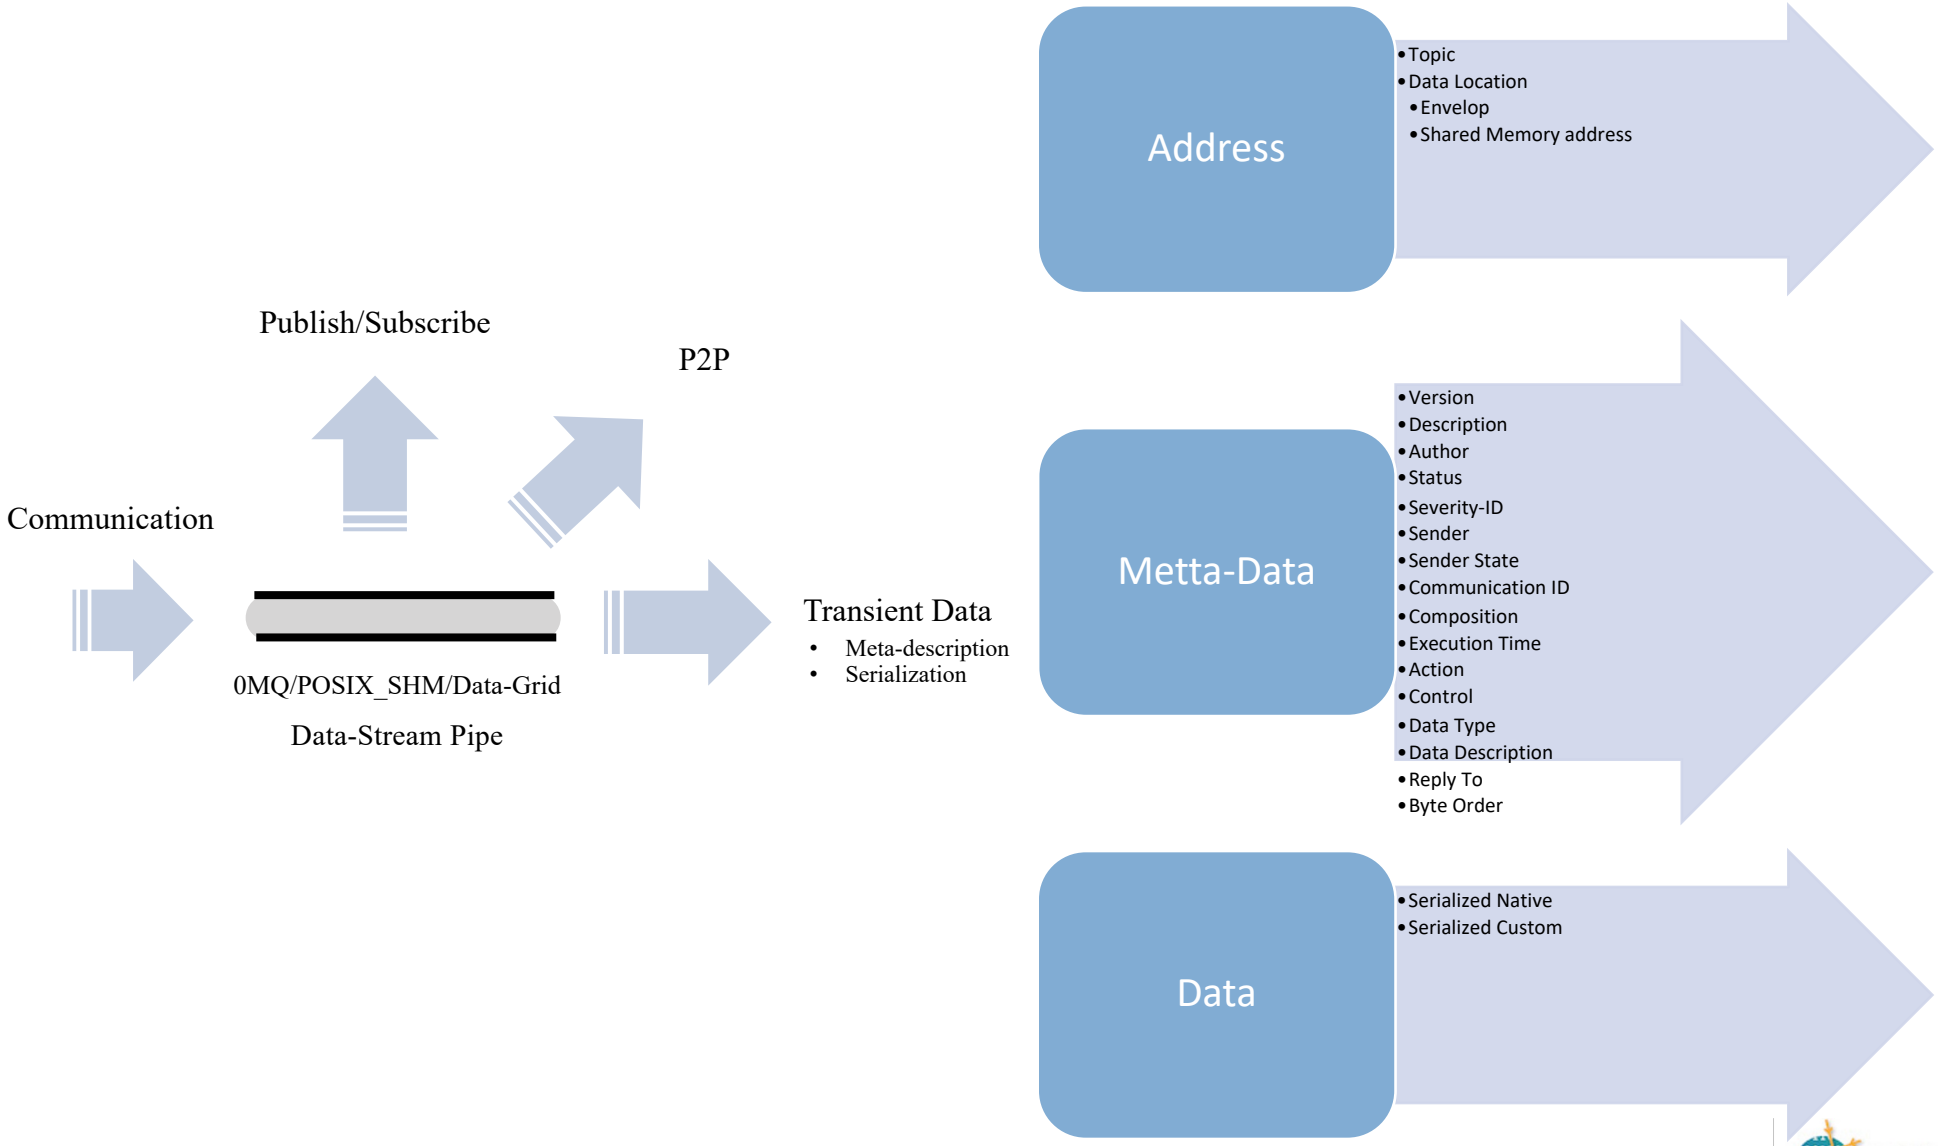

EIC prototype calorimeter SRO application design and configuration

```
---
io-services:
  reader:
    class: org.jlab.ersap.coda.engines.AggFileReaderEngine
    name: Source
  writer:
    class: org.jlab.ersap.coda.engines.AggStoreHistogramEngine
    name: Sync
services:
  - class: org.jlab.ersap.coda.engines.FAdcIdEngine
    name: Beam
  - class: org.jlab.ersap.coda.engines.FAdcCosmicIdEngine
    name: Cosmic
#  -class: KMeanEvtdentifier
#    name: cppBeam
#    lang: cpp
```

```
configuration:
 io-services:
   writer:
     frame_title: "ERSAP"
     frame_width: 1400
     frame_height: 1200
     #> hist_titles is a string containing the list of crate-slot-channel separated by ,
     hist_titles: "1-17-0, 1-17-1, 1-17-2, 1-17-3, 1-17-4, 1-17-5, 1-17-6, 1-17-7, 1-17-8, 1-17-9, 1-1
      hist_bins: 100
     hist_min: 0
     hist_max: 8000
     scatter_reset: true
     #> grid_size defines a layout for histogram visualization
     #> (e.g. 5 will plot 25 histograms in 5x5 matrix)
     grid_size: 5
 services:
   Beam:
     s_window: 32
     s_step: 1
     s_hits: 5
     #      t_slot: 17
     #      t_channel: 14
     b_thr: 20
     bc_slot: 17
     bc_channel: 12
     bc_qmin: 0
     bc_qmax: 8000
   Cosmic:
     s_window: 32
     s_step: 1
     s_hits: 5
mime-types:
  - binary/data-evio
  - binary/data-jobj
```