ERSAP: Towards Better HEP/NP Data-Stream Analytics With Flow-Based Programming

Environment for Real-time Streaming,
Acquisition and Processing

V. Gyurjyan, D. Abbott, N. Brei, M. Goodrich, G. Heyes, E. Jastrzembski,
D. Lawrence, B. Raydo, C. Timmer

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

Outline

- Looking forward to future experiments at the JLAB and EIC
  - Reevaluate existing readout systems and their interface to the back-end
- Reactive, actor-model based programming model
- Results from ERSAP based pilot applications.

Streaming system components: Tiered storage model

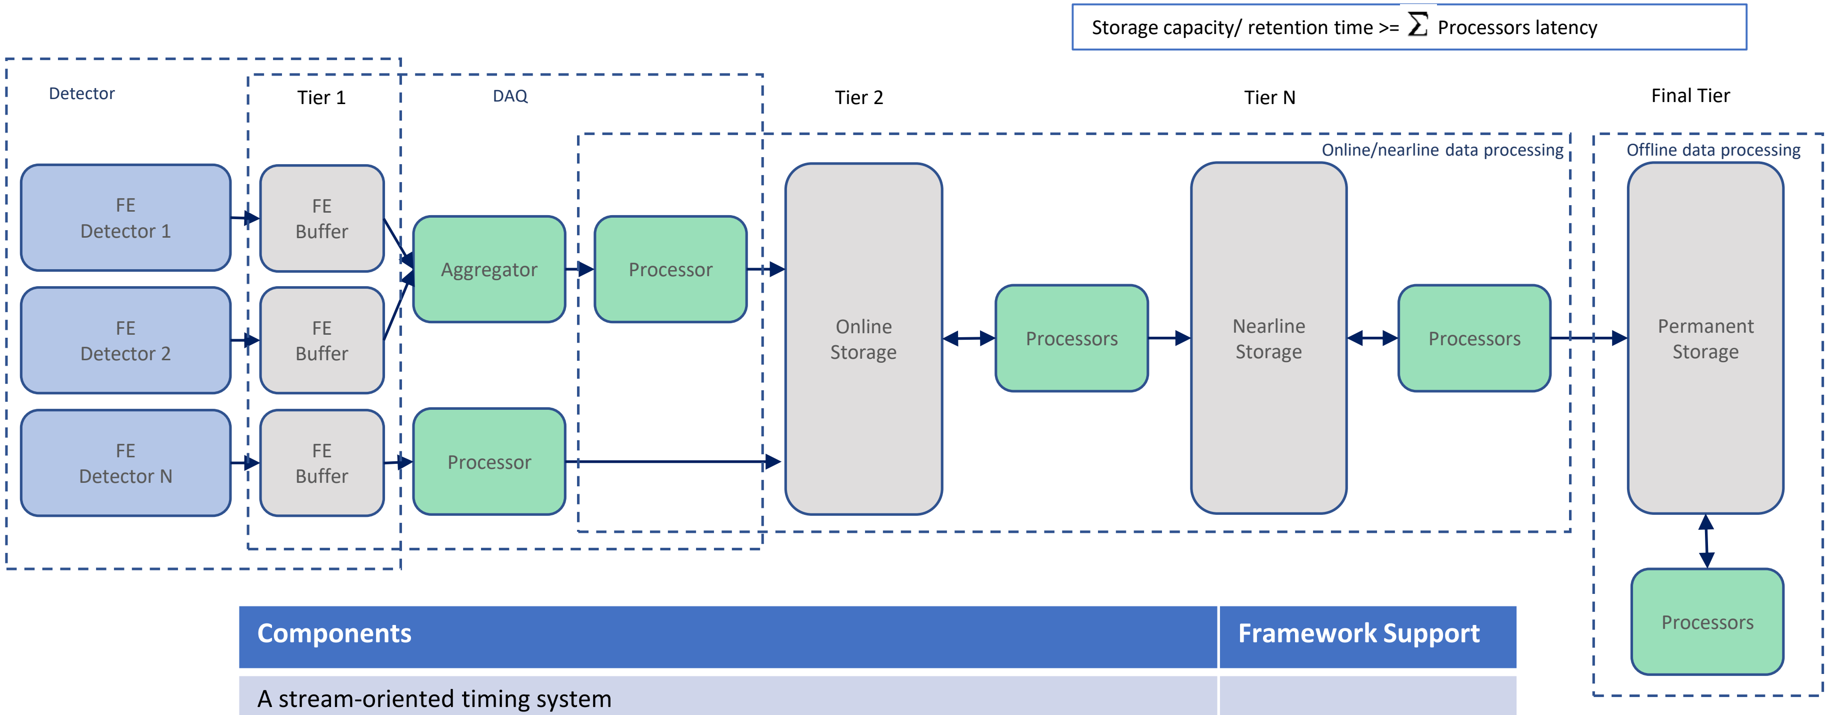

Components	Framework Support
A stream-oriented timing system	
A standard stream data format	✓
Front-end electronics that outputs time-based streams of data	
Efficient streaming data transport	✓
A stream oriented random-access data storage tier	✓
A framework for data processing tasks (virtual triggers, calibration, reconstruction, monitoring, data storage, etc.)	✓
A framework for data-flow orchestration and data-stream processing application design and deployment	✓

Πάντα ῥεῖ Flow-Based Programming Paradigm

- Proposed in the late 60s by J. Paul Rodker Morrison
- “Assembly line” data processing
- Data flows through asynchronous, concurrent processors (“black box” actors)
- Actors communicate via data chunks (called information packets or data-quanta)
- Data-quanta are traveling across predefined connections (conveyor belts), where connections are specified externally to the processors.
- Data is pushed through actors, while actors are reacting on passing data quantum.
- Actors are performing independent, well-defined functions
- Simple reconfigure
- Minimizes side-effects

ERSAP architecture

- Reactive event driven actor, data-stream pipe, and orchestrator.
- Stream of data quanta, flowing through directed graph of actors.
- Application is a network of independent “black box” actors.

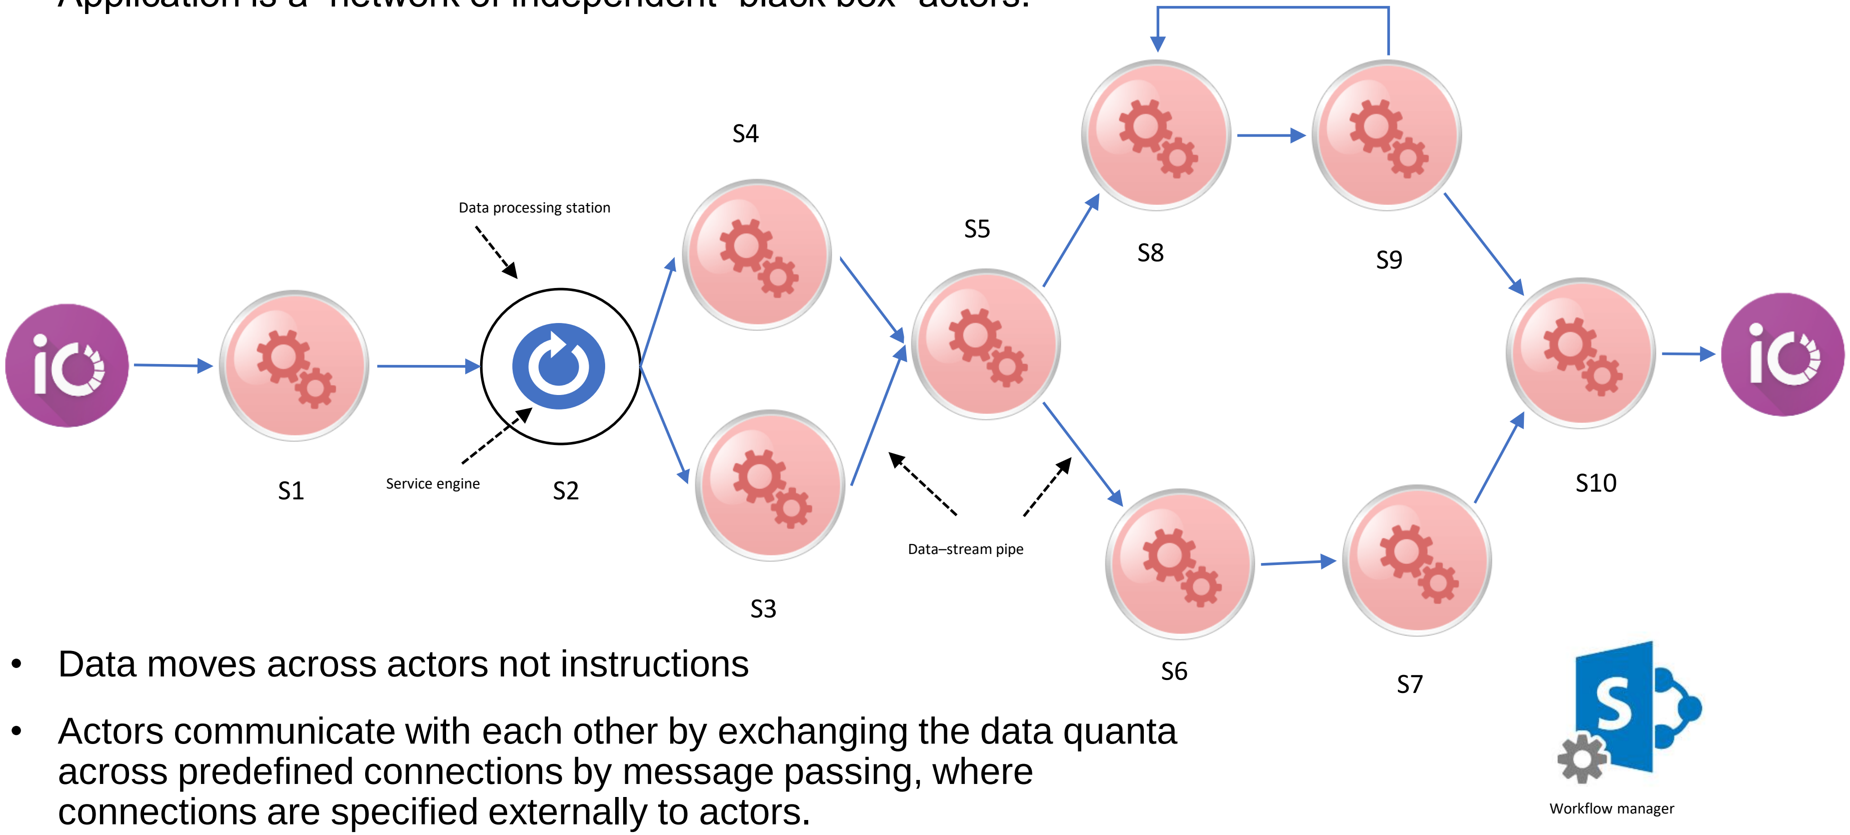

- Data moves across actors not instructions
- Actors communicate with each other by exchanging the data quanta across predefined connections by message passing, where connections are specified externally to actors.
- User provided data processing single-threaded algorithms (engines) are presented as fully scalable actors in the framework.

ERSAP 3-layer structure

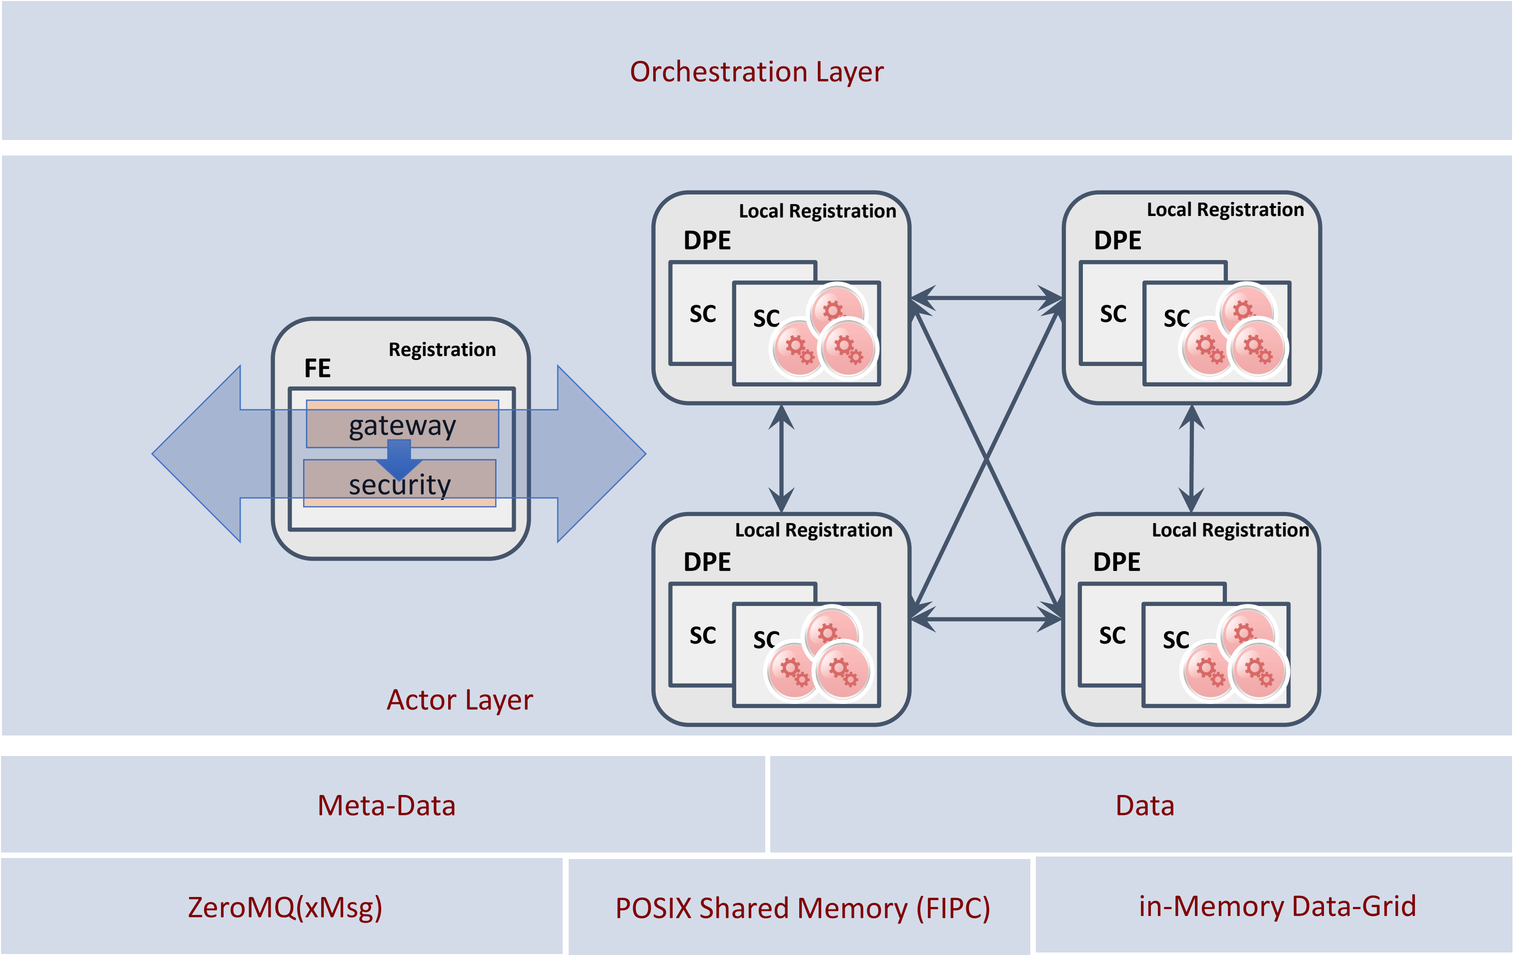

Data processing station: actor

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

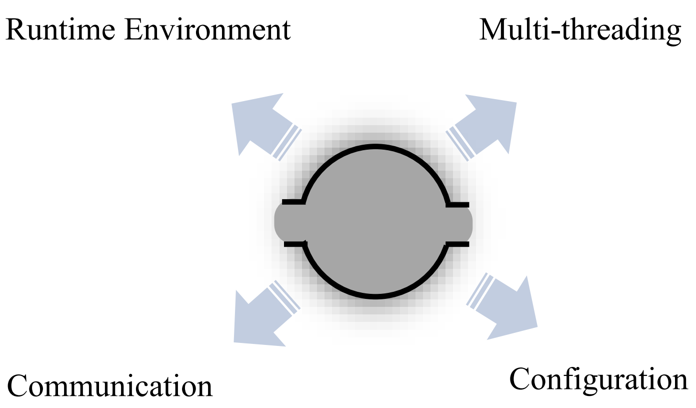

Runtime Environment

Multi-threading

Communication

Configuration

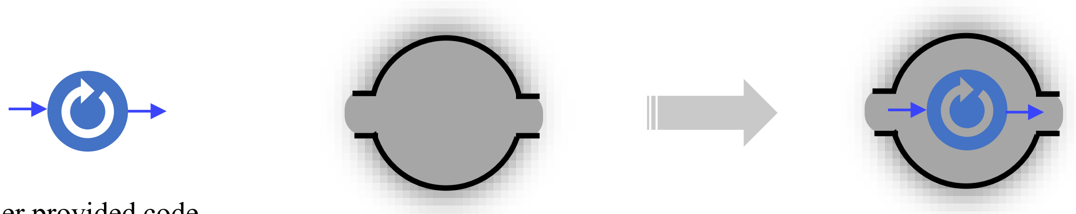

User provided code
Engine

Data Processing Station

Data Processing Actor

Streaming data transport (conveyer belt)

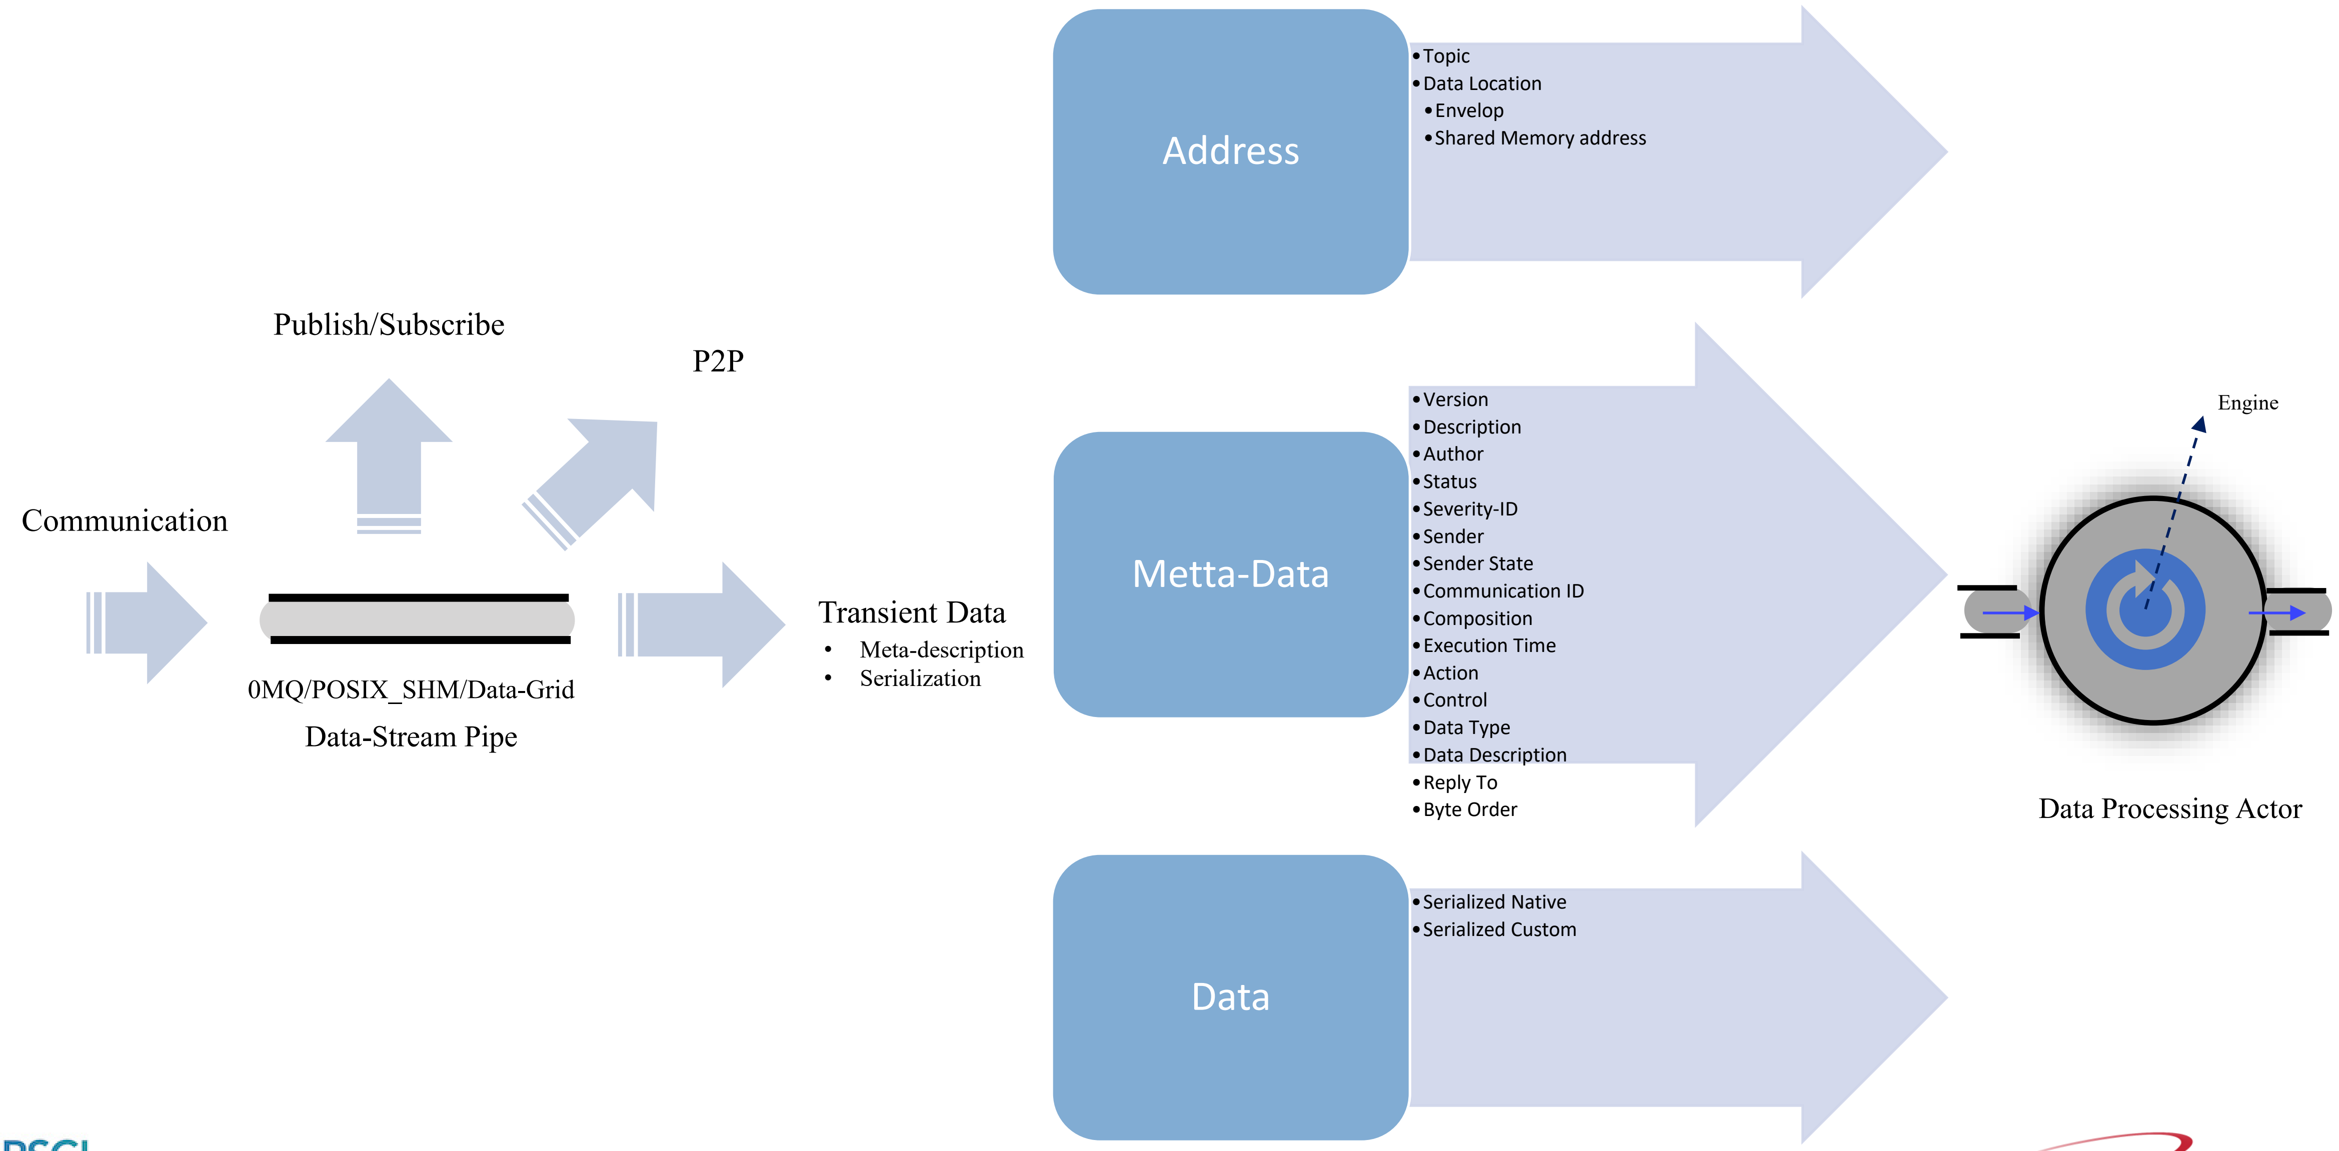

Publish/Subscribe

P2P

Communication

0MQ/POSIX_SHM/Data-Grid

Data-Stream Pipe

Transient Data

- Meta-description
- Serialization

Address

- Topic
- Data Location
    - Envelop
    - Shared Memory address

Metta-Data

- Version
- Description
- Author
- Status
- Severity-ID
- Sender
- Sender State
- Communication ID
- Composition
- Execution Time
- Action
- Control
- Data Type
- Data Description
- Reply To
- Byte Order

Data

- Serialized Native
- Serialized Custom

Data Processing Actor

Hall-D EIC prototype calorimeter setup

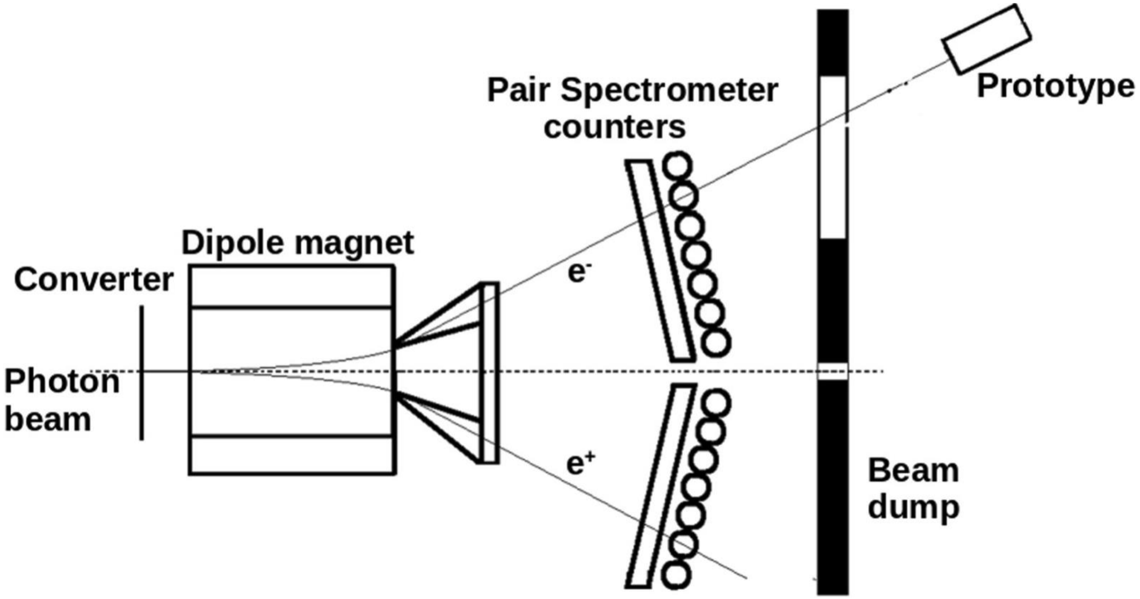

Hall-D prototype calorimeter ERSAP SRO pipeline

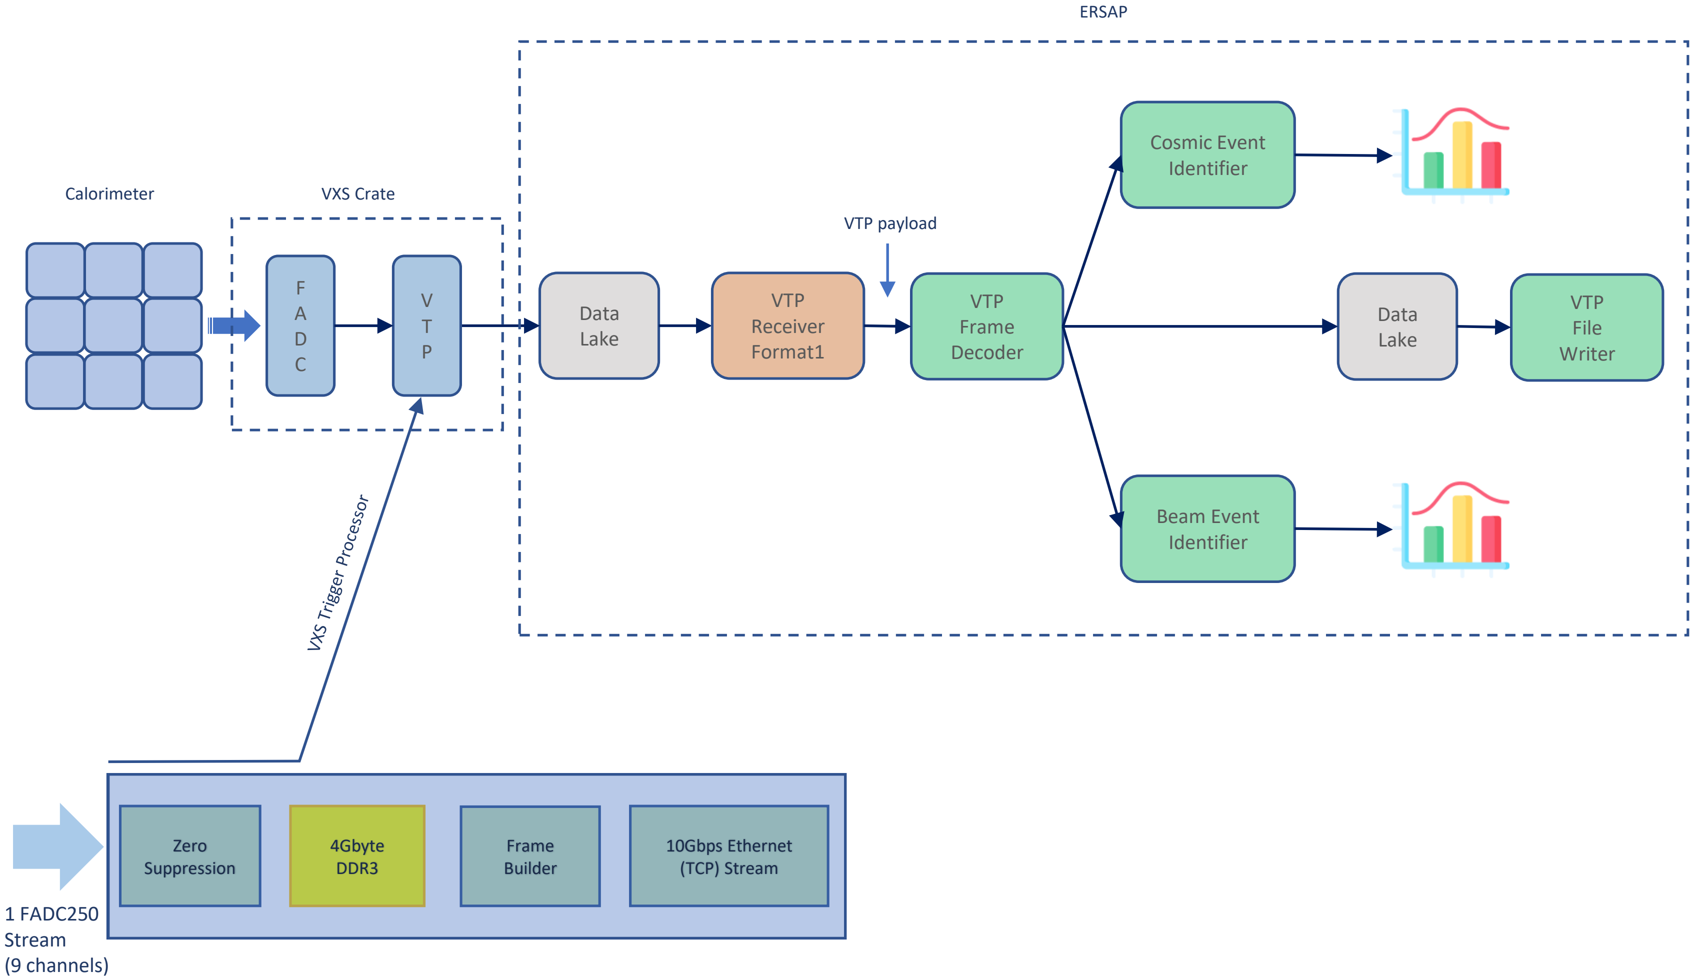

VTP fADC data stream frames

65.5us VTP frame

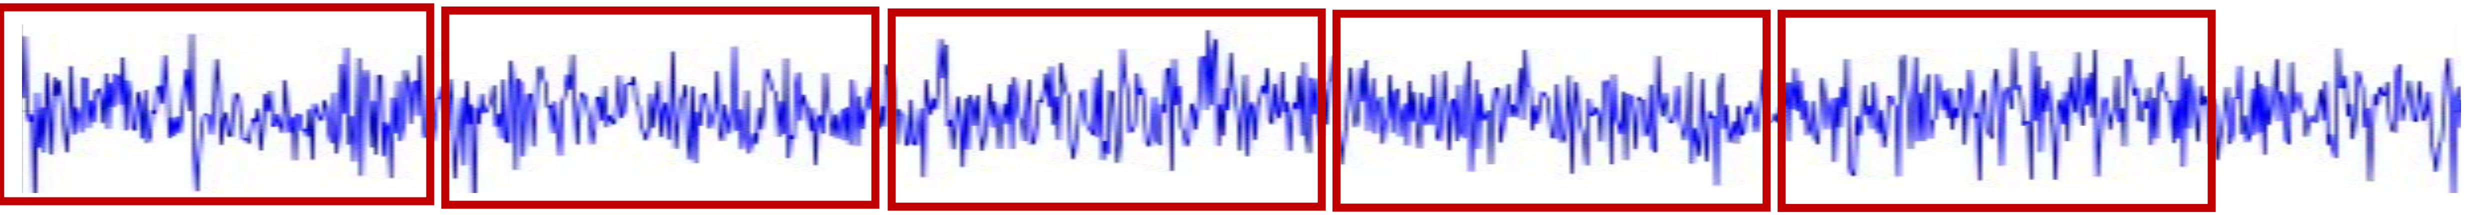

Event Identification. Sliding window technique

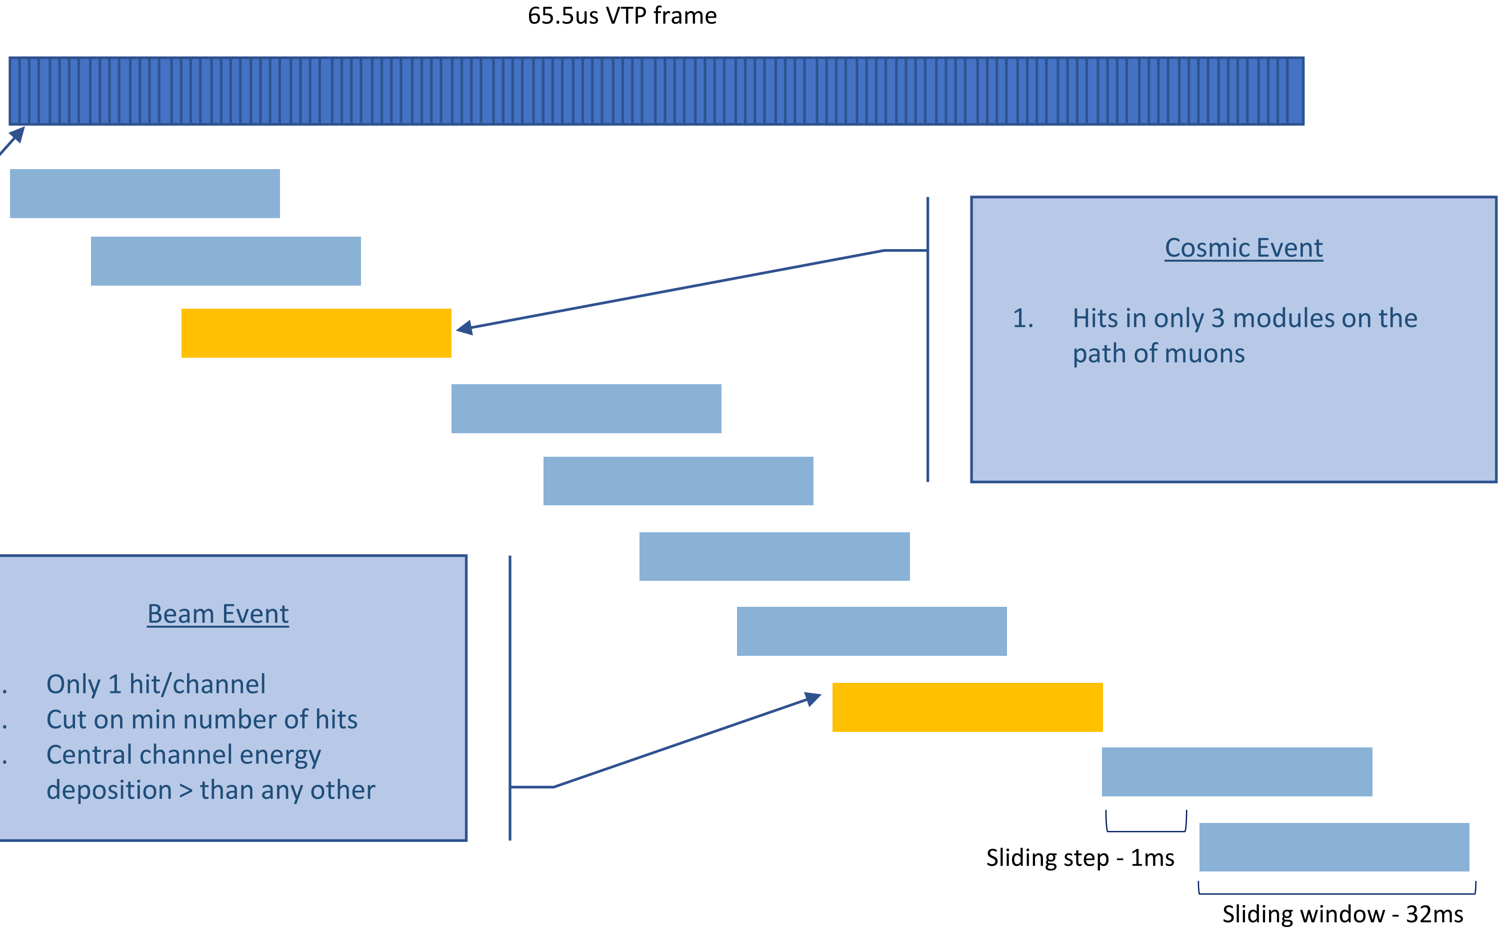

Sliding pointers

Find min/max time of hits within the frame

Cosmic Event

1. Hits in only 3 modules on the path of muons

Beam Event

1. Only 1 hit/channel
2. Cut on min number of hits
3. Central channel energy deposition > than any other

3x3 Calorimeter. Beam

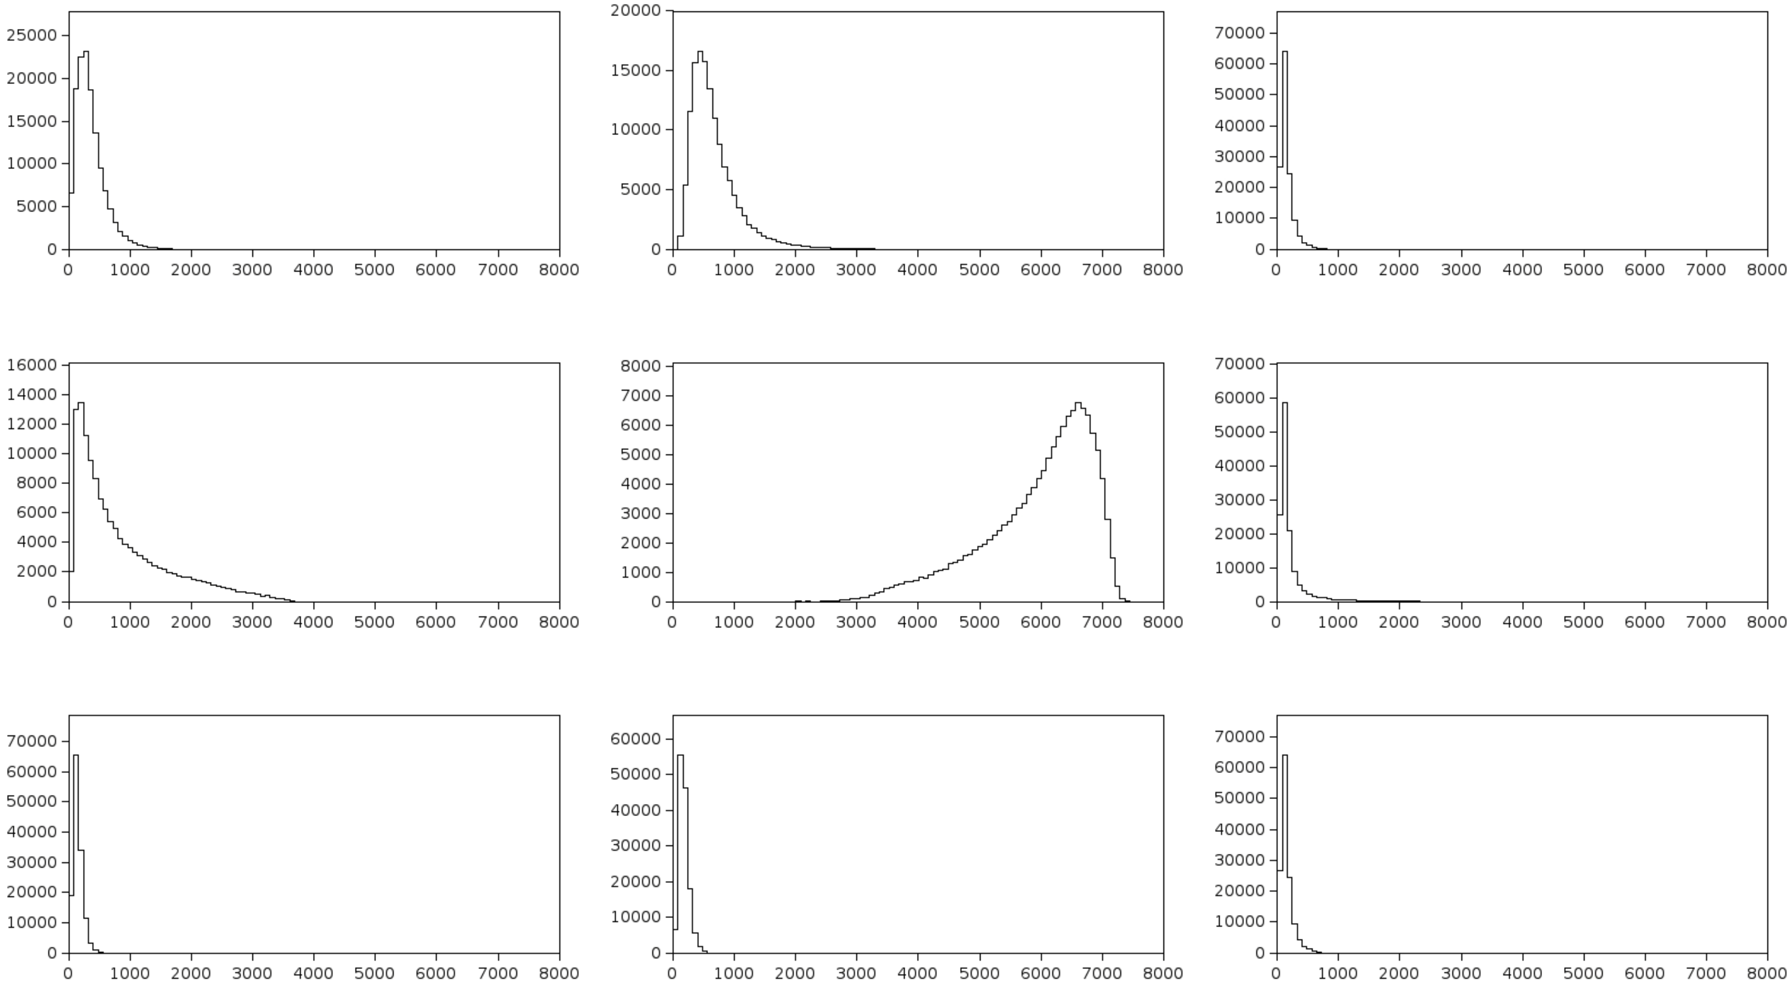

3x3 Calorimeter. Beam sum

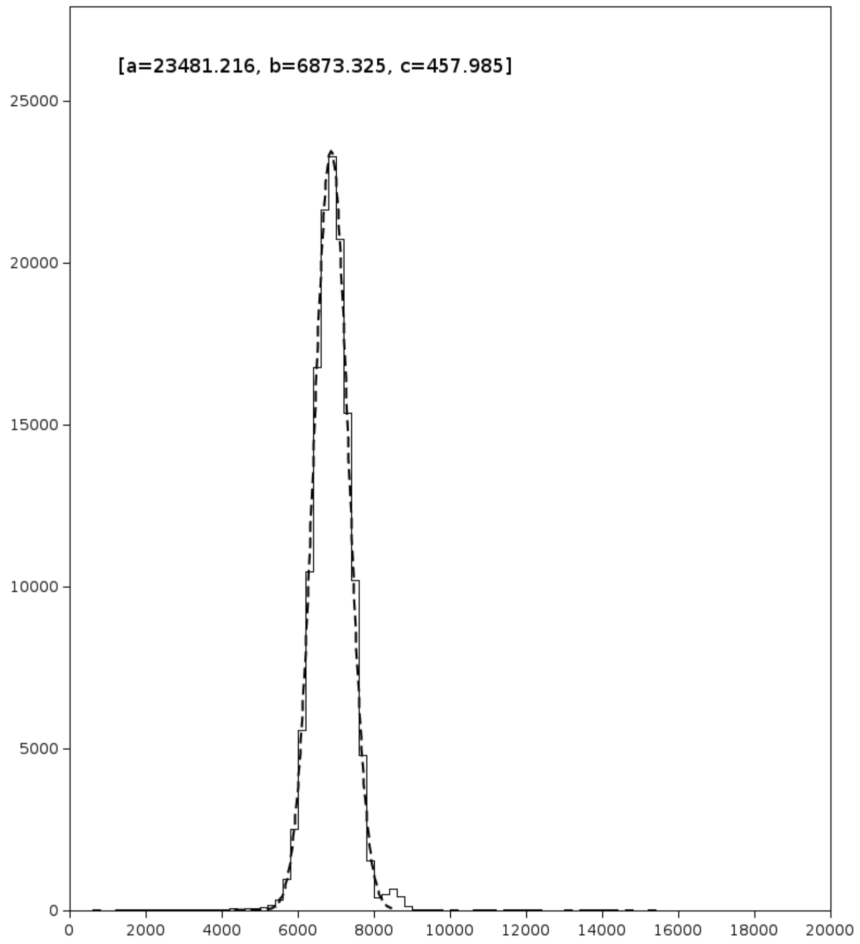

~2GeV
Based on cosmic muon deposition
(see next slide)

3x3 Calorimeter. Cosmic

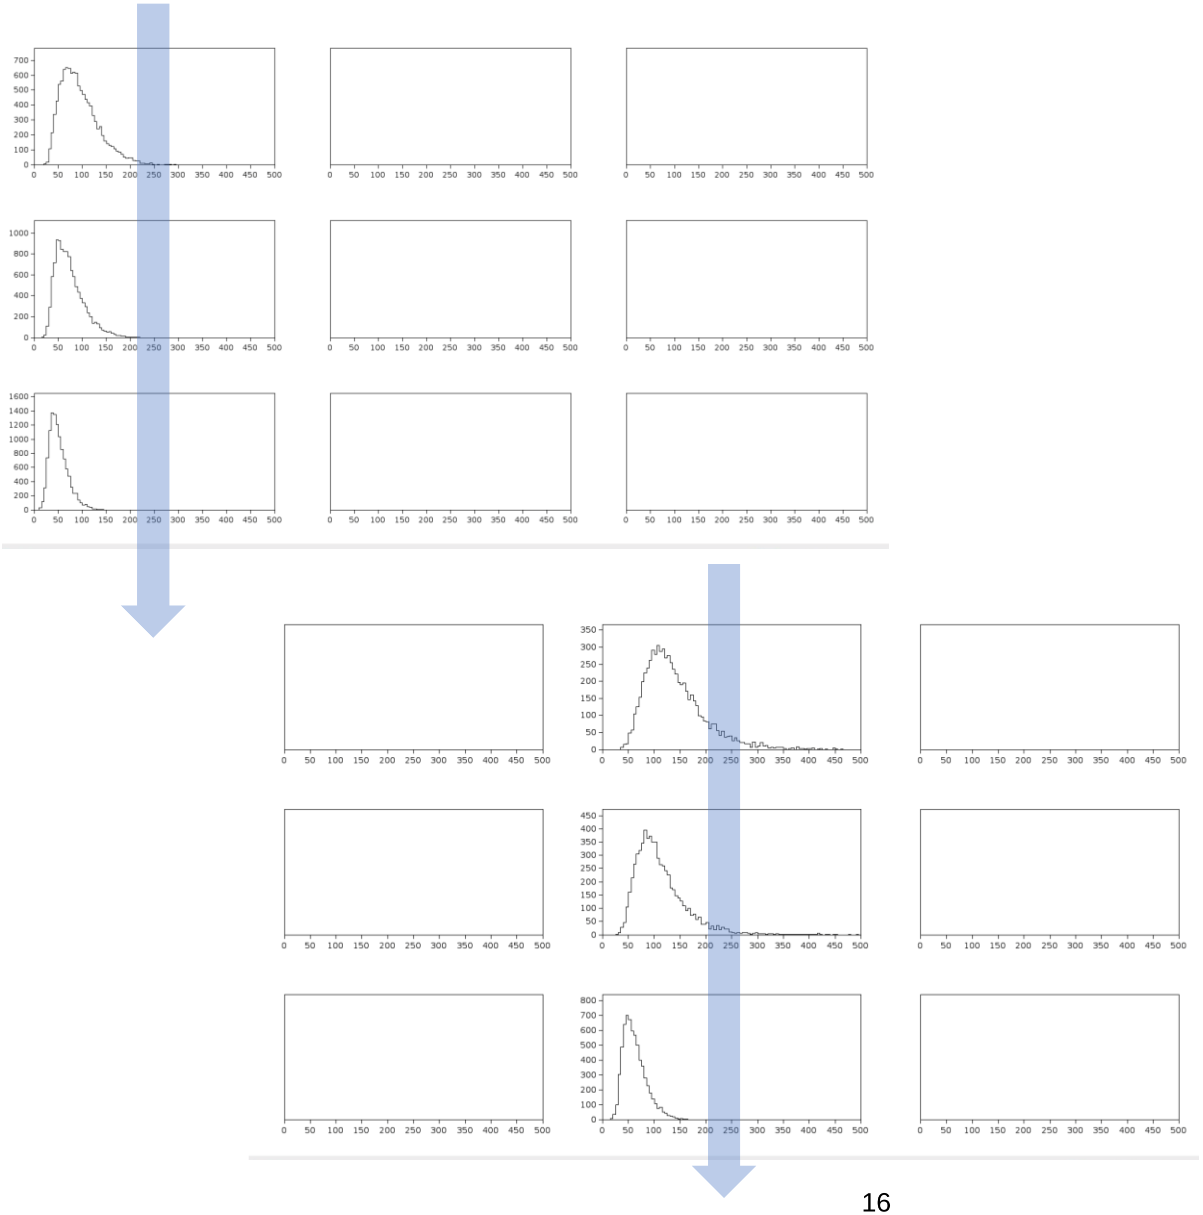

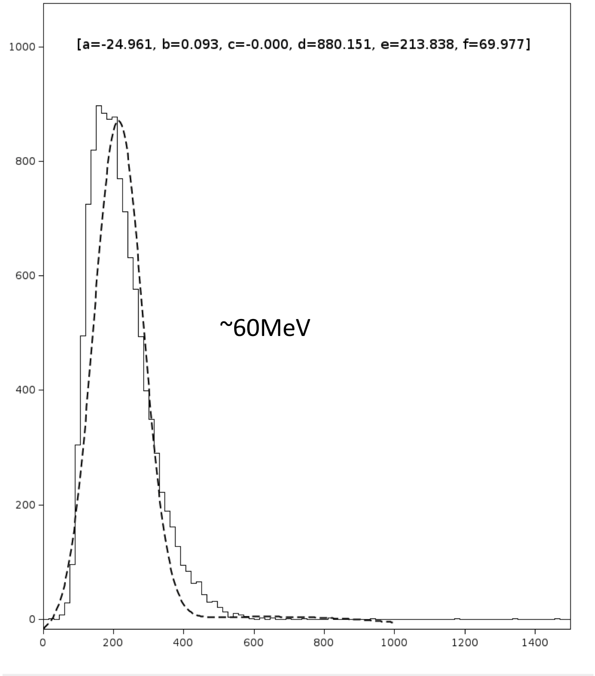

~60MeV

Hall-D prototype calorimeter ERSAP SRO pipeline

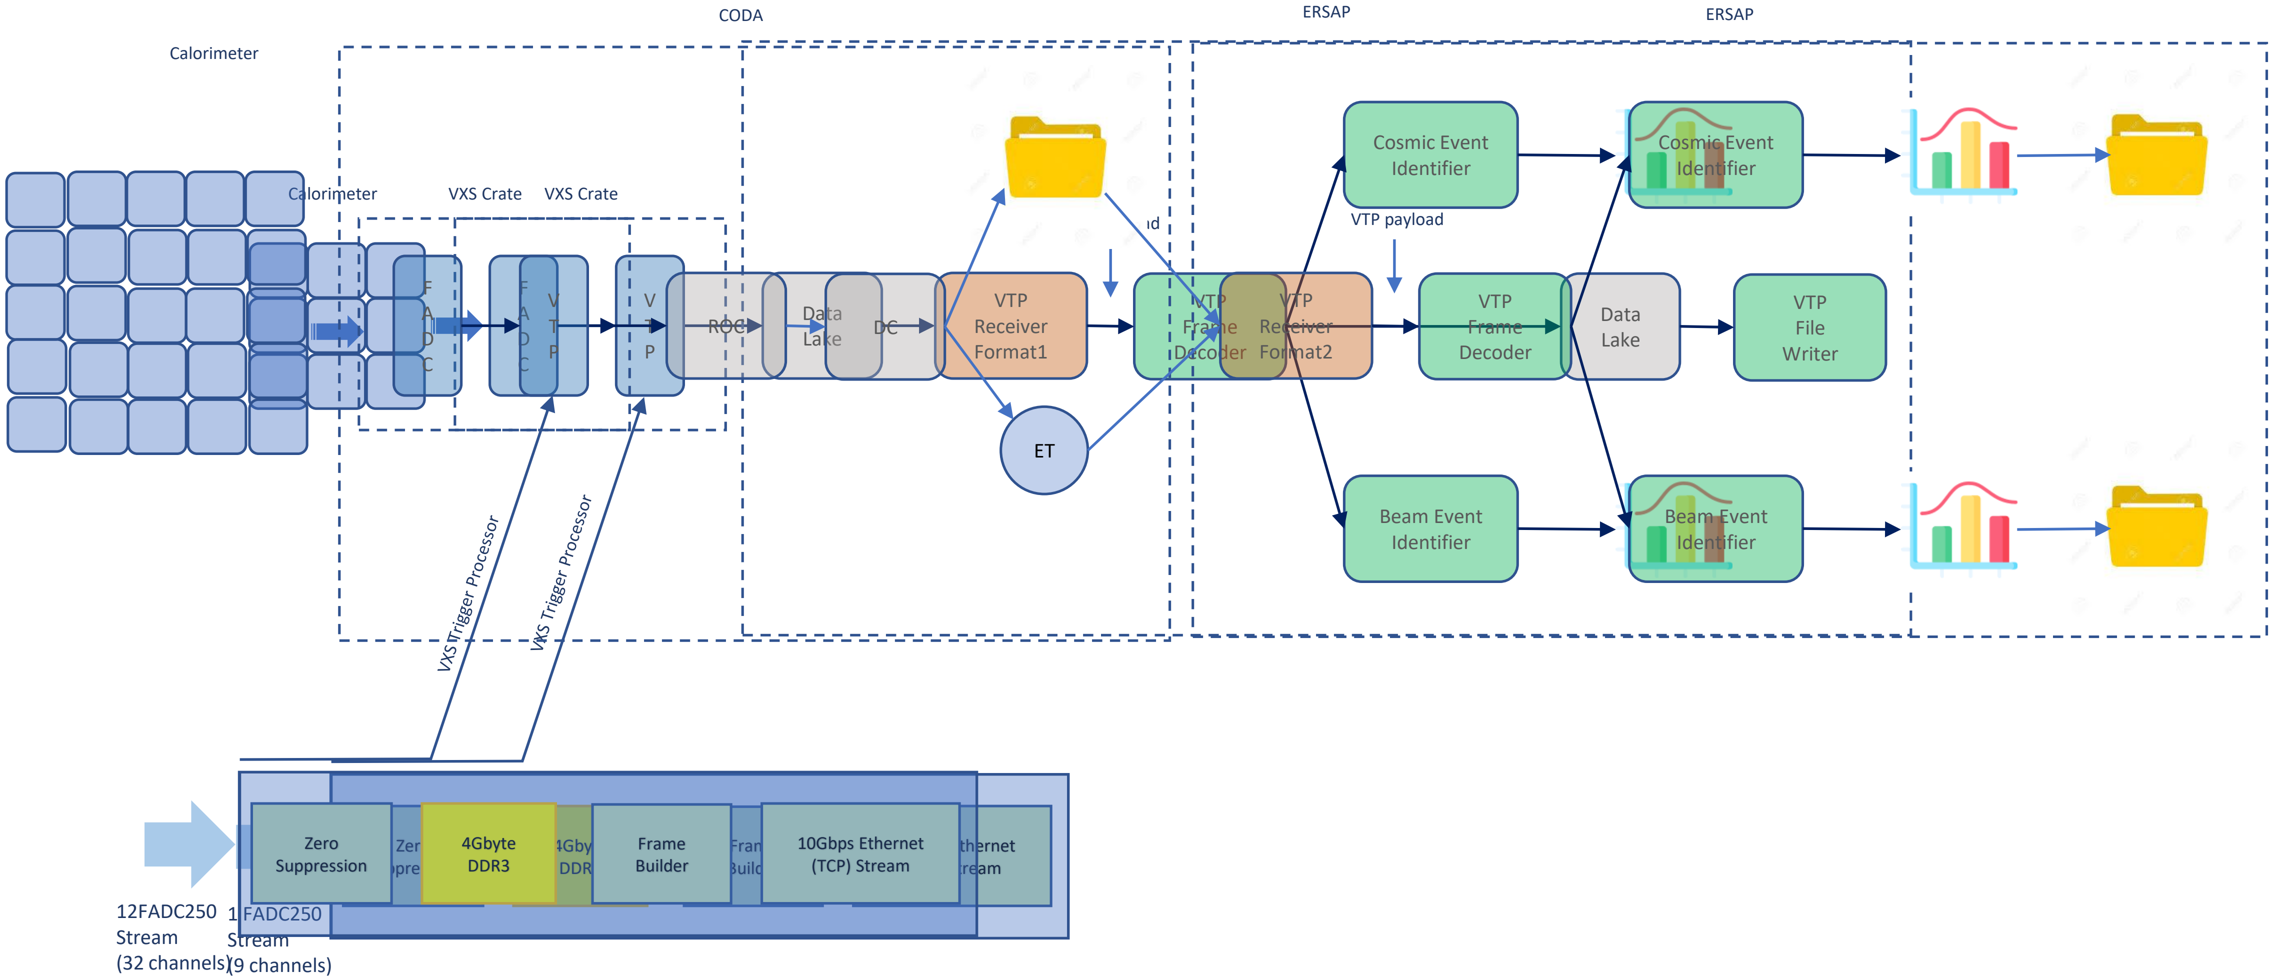

EIC prototype calorimeter SRO pipeline at DESY

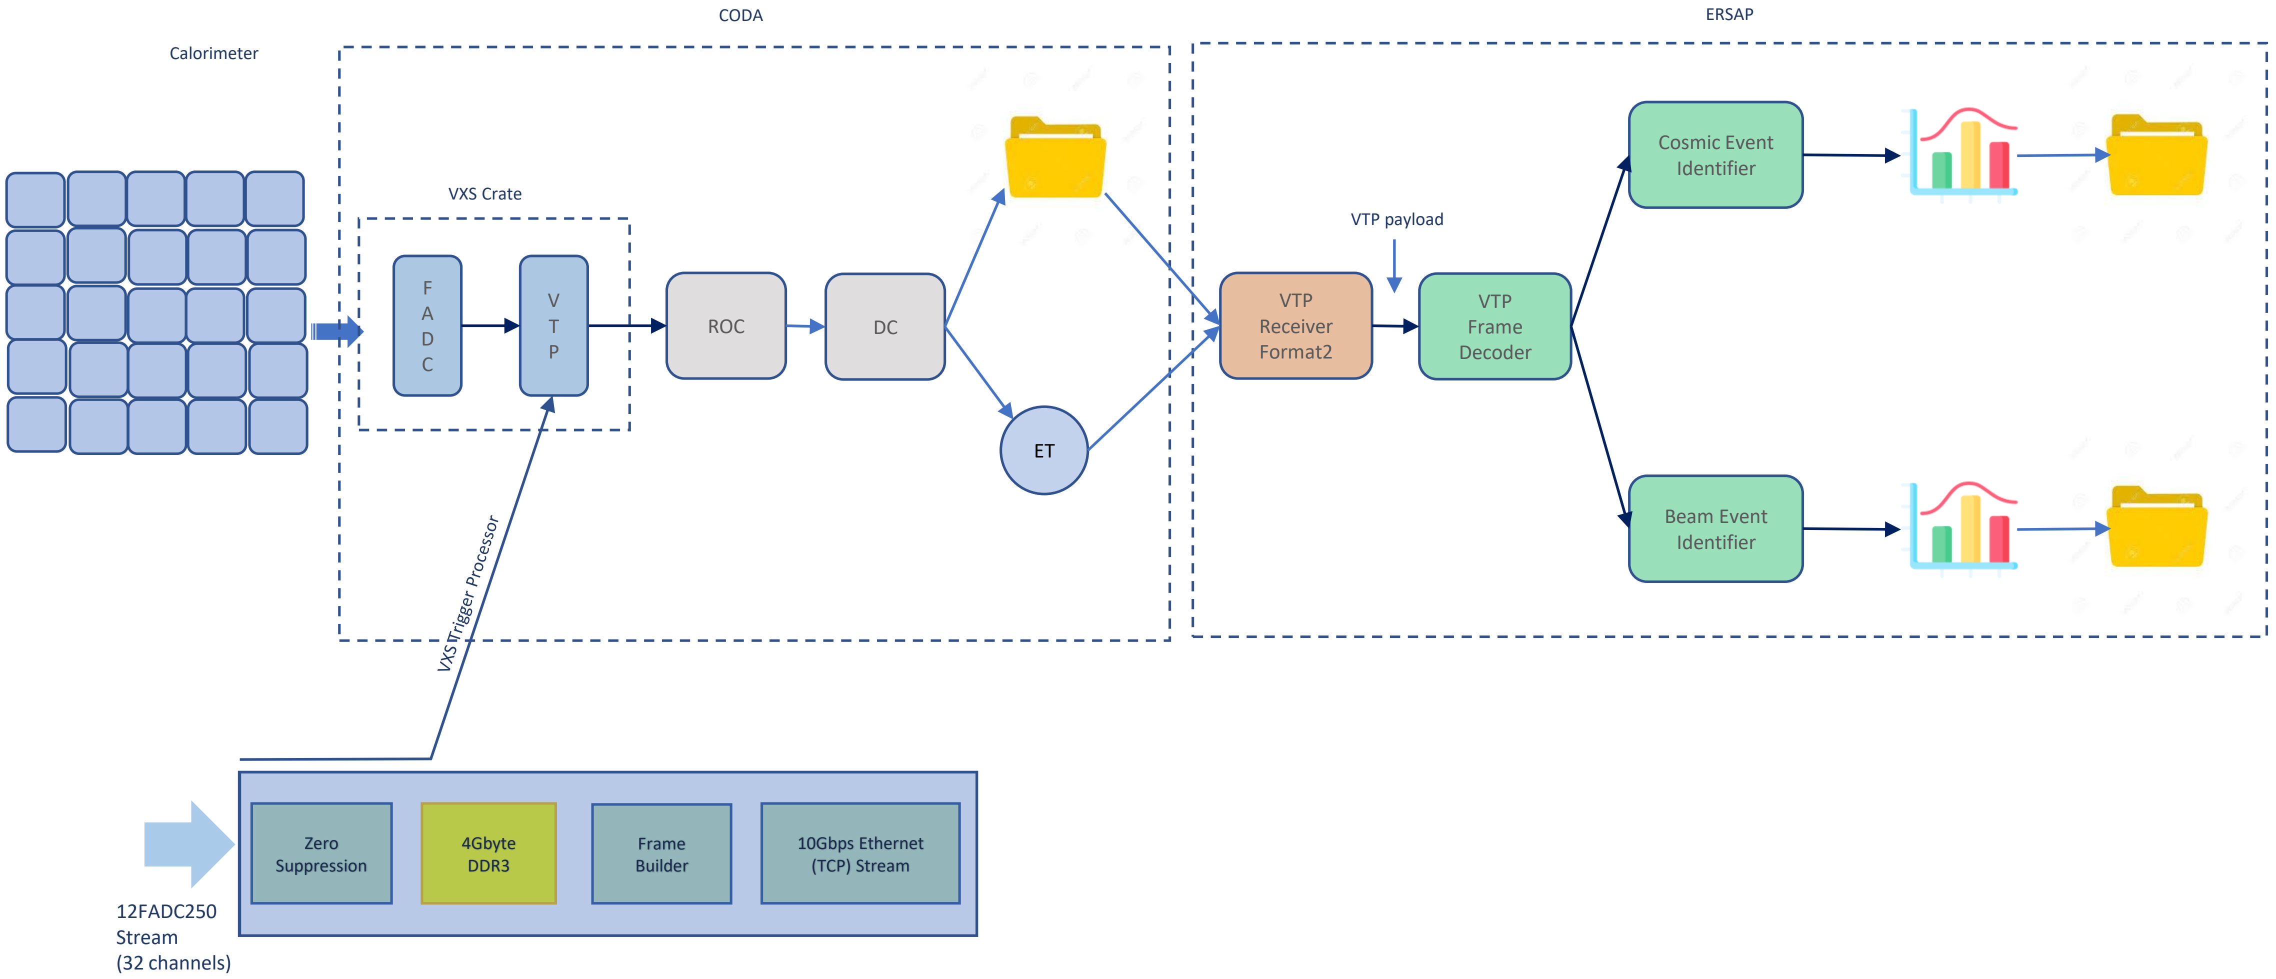

EIC prototype calorimeter SRO application design and configuration

```
---
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
```

Time distribution of hits

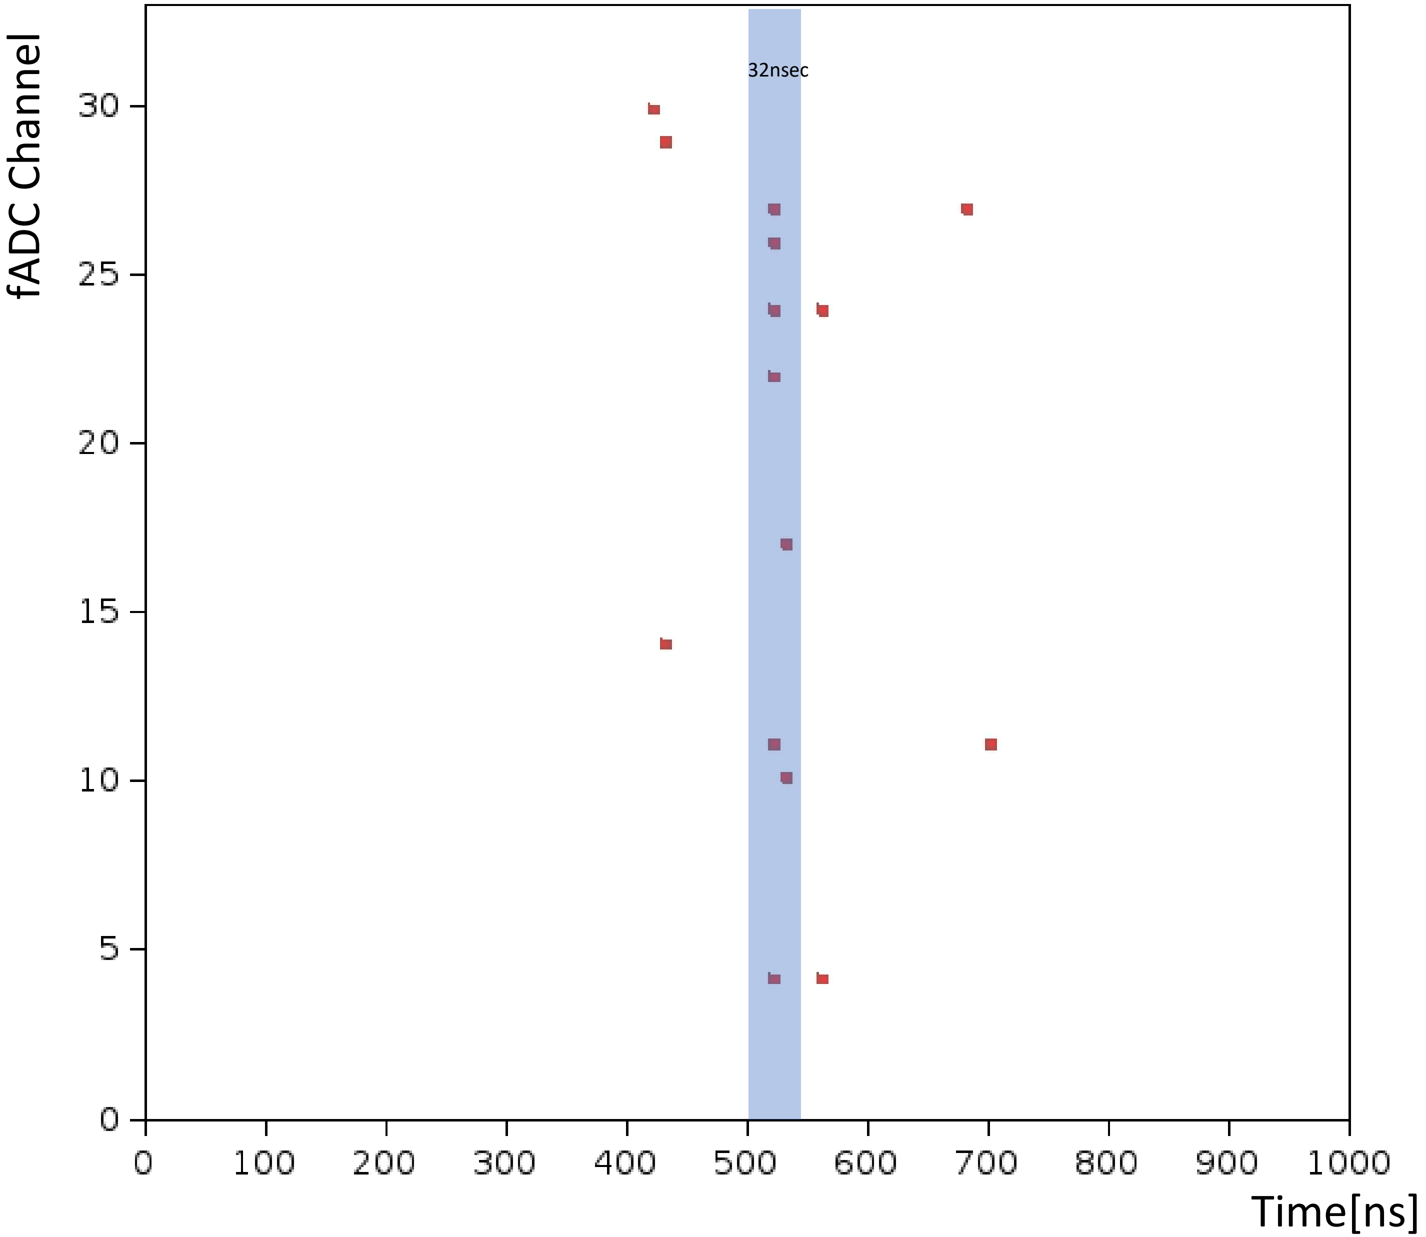

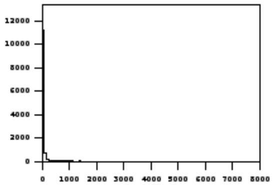

1-17-0

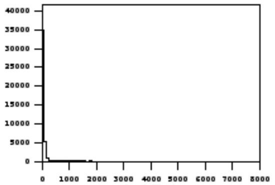

1-17-1

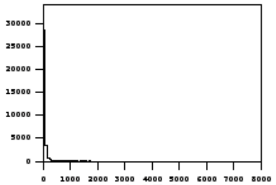

1-17-2

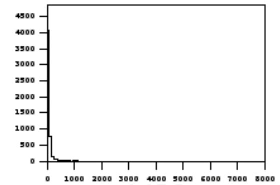

1-17-3

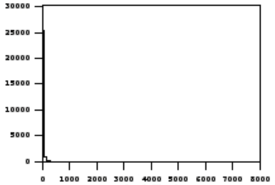

1-17-4

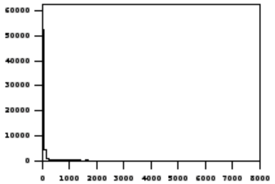

1-17-5

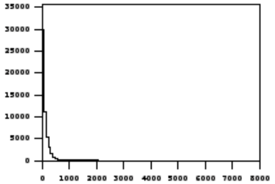

1-17-6

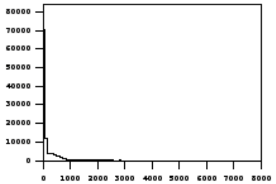

1-17-7

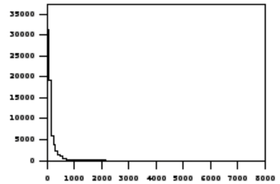

1-17-8

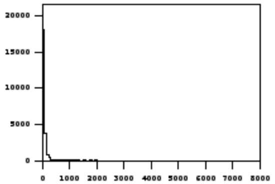

1-17-9

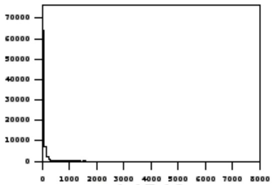

1-17-10

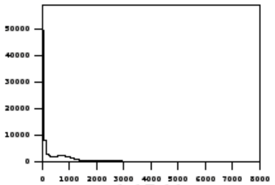

1-17-11

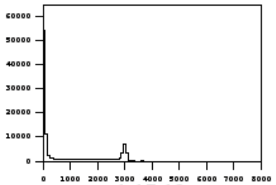

1-17-12

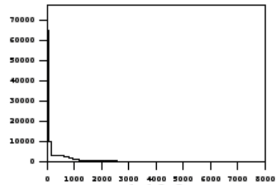

1-19-0

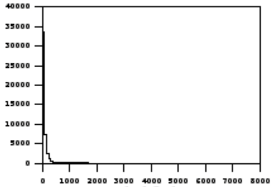

1-19-1

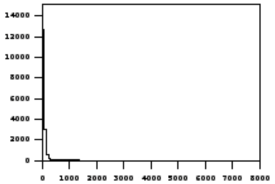

1-19-2

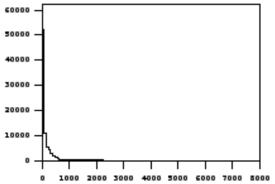

1-19-3

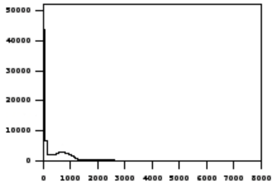

1-19-4

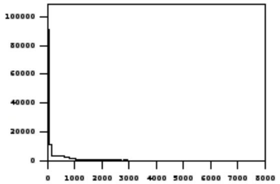

1-19-5

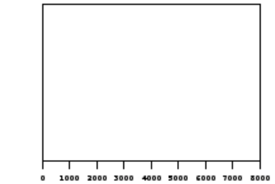

1-19-6

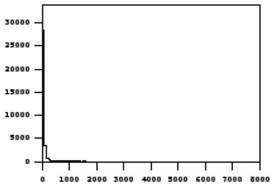

1-19-7

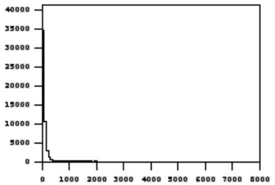

1-19-8

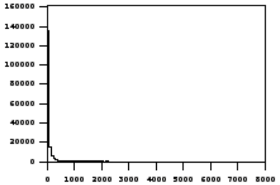

1-19-9

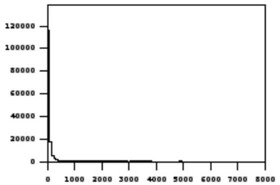

1-19-10

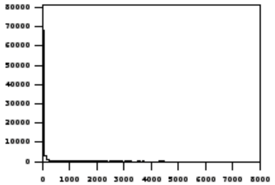

1-19-11

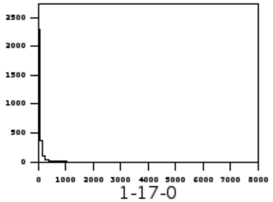

1-17-0

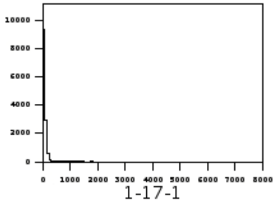

1-17-1

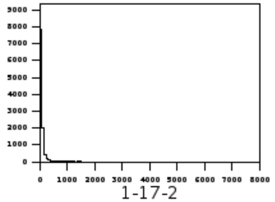

1-17-2

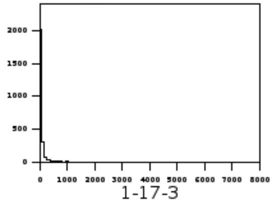

1-17-3

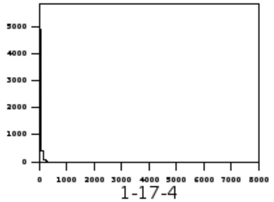

1-17-4

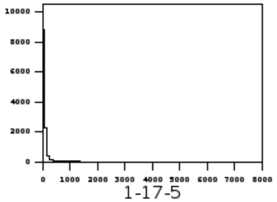

1-17-5

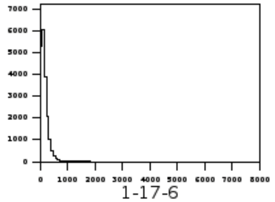

1-17-6

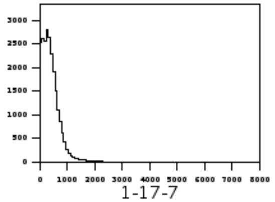

1-17-7

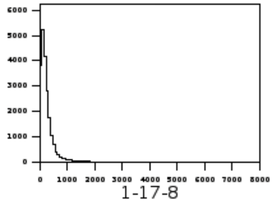

1-17-8

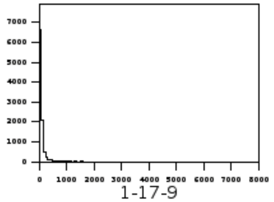

1-17-9

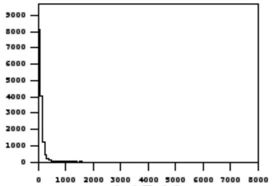

1-17-10

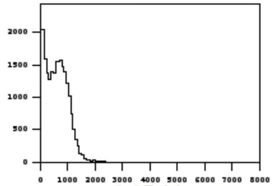

1-17-11

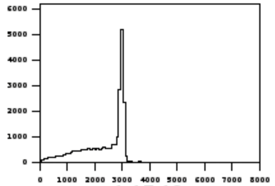

1-17-12

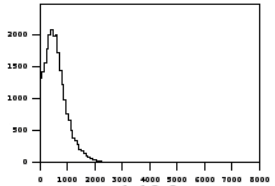

1-19-0

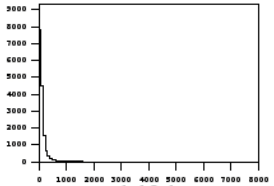

1-19-1

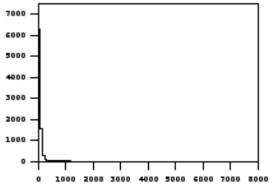

1-19-2

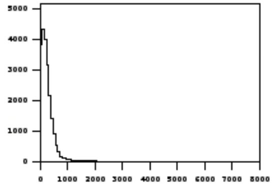

1-19-3

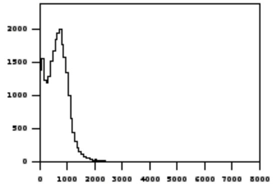

1-19-4

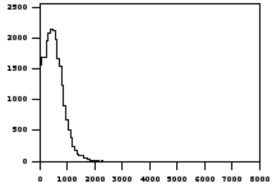

1-19-5

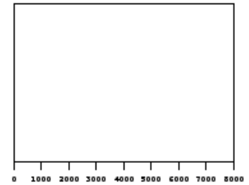

1-19-6

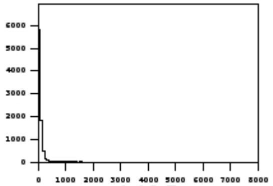

1-19-7

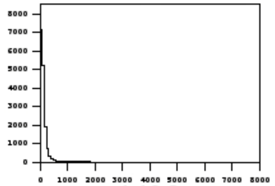

1-19-8

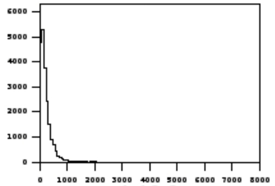

1-19-9

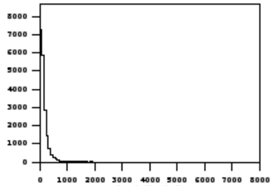

1-19-10

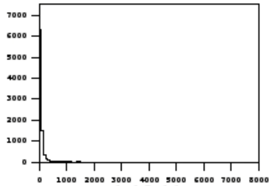

1-19-11

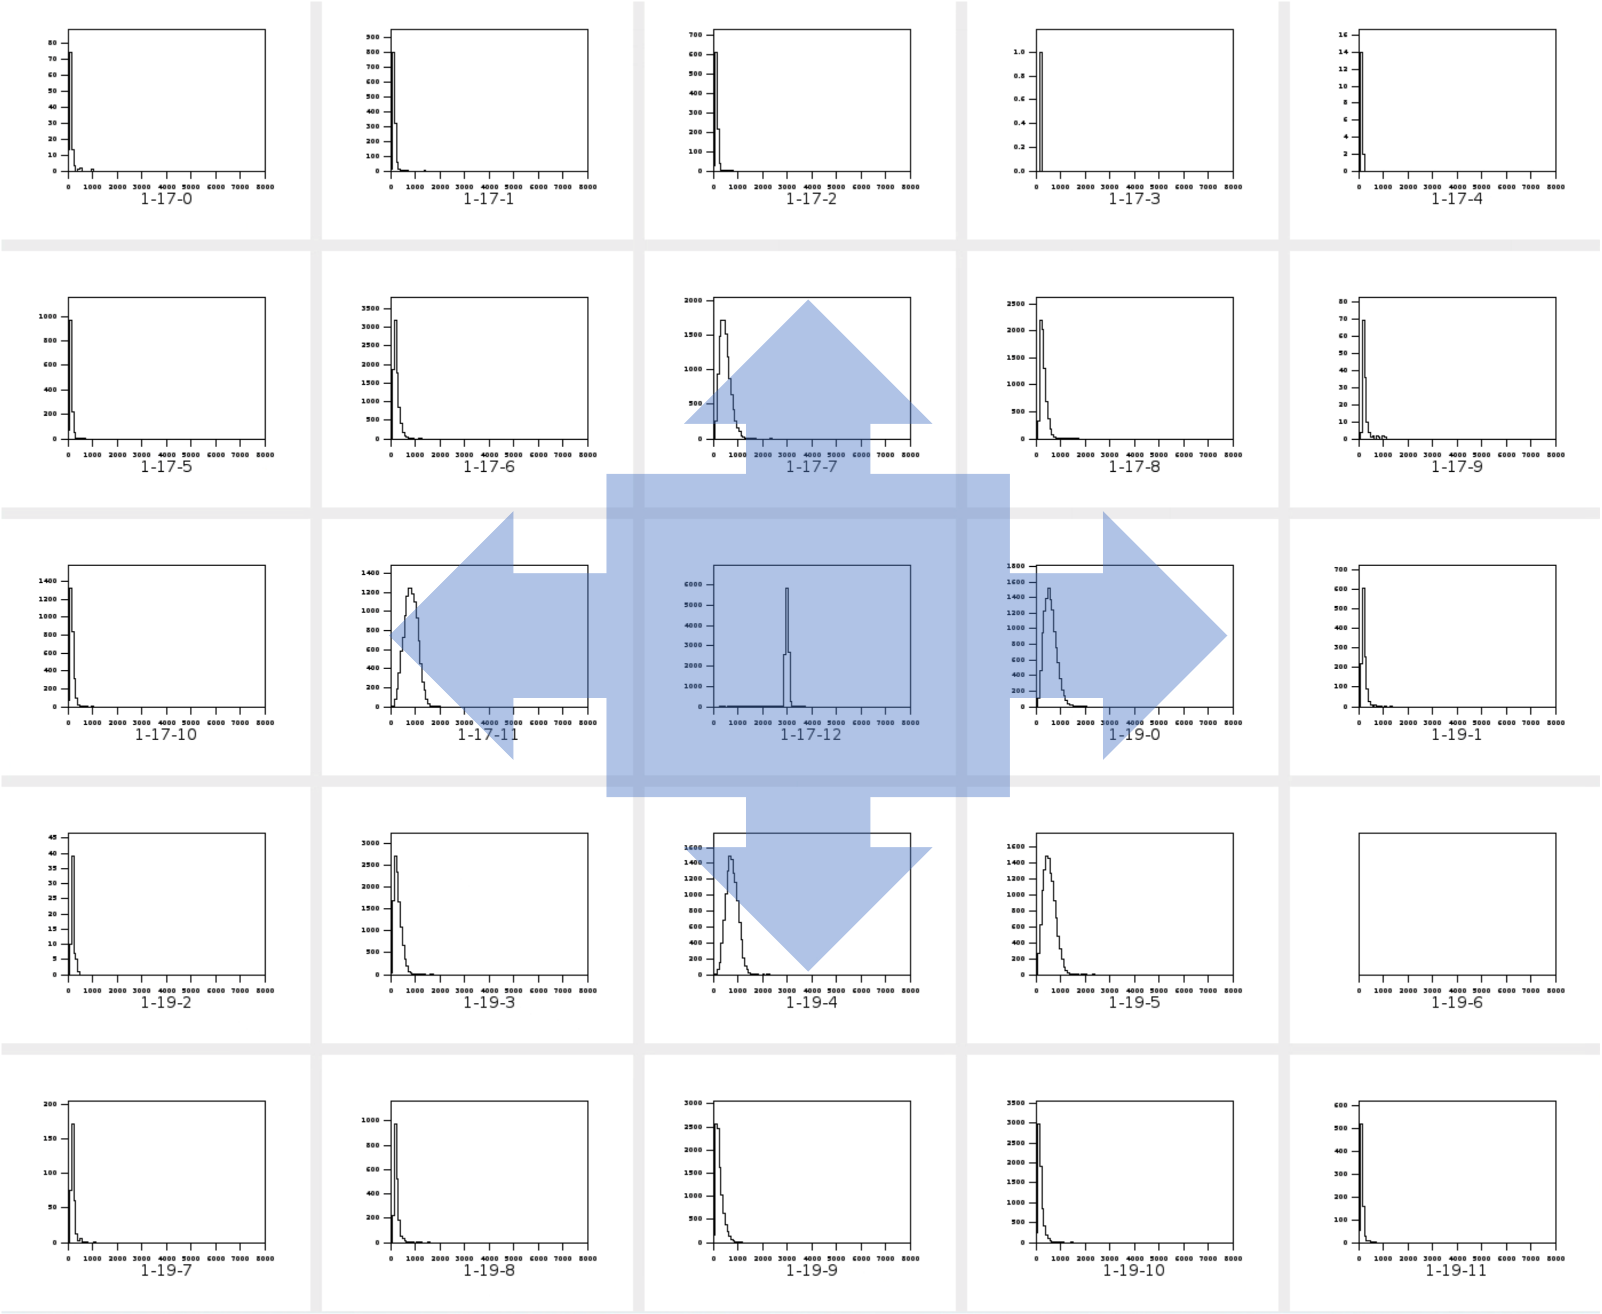
1-17-0
1-17-1
1-17-2
1-17-3
1-17-4
1-17-5
1-17-6
1-17-7
1-17-8
1-17-9
1-17-10
1-17-11
1-17-12
1-19-0
1-19-1
1-19-2
1-19-3
1-19-4
1-19-5
1-19-6
1-19-7
1-19-8
1-19-9
1-19-10
1-19-11

Charge cut effect

No energy/charge cut in the soft trigger

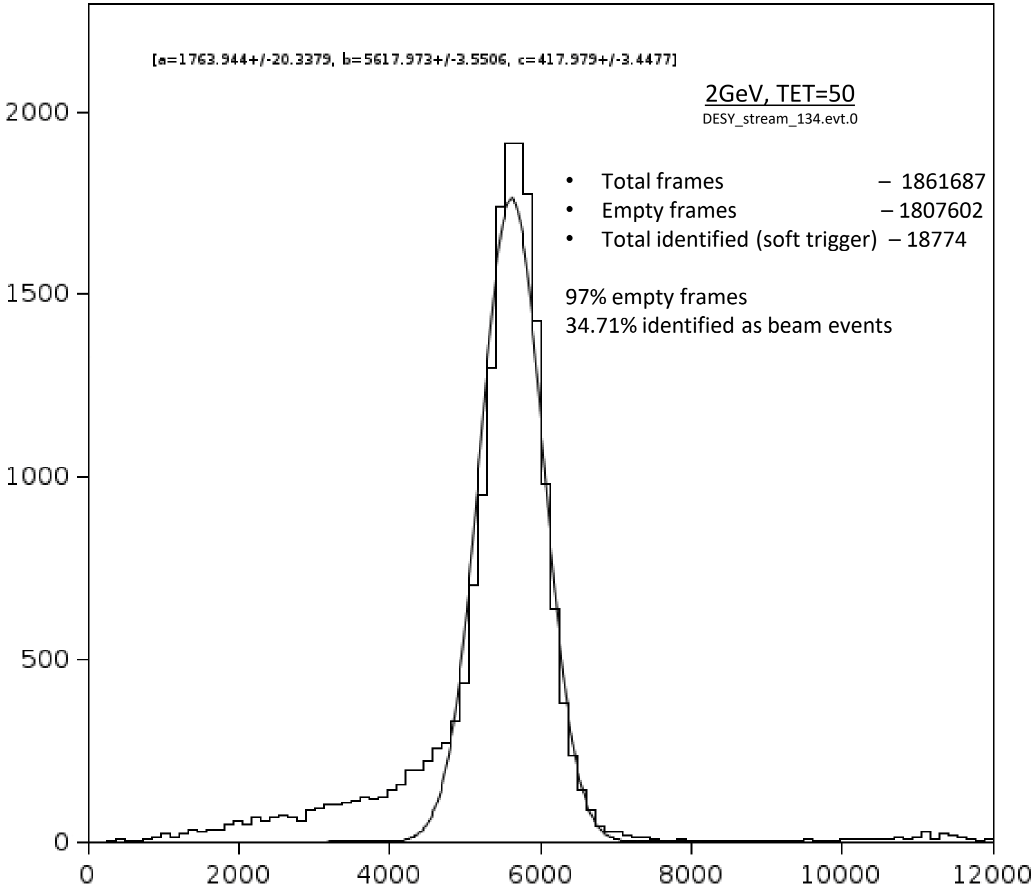

2GeV, TET=50

DESY_stream_134.evt.0

- Total frames – 1861687
- Empty frames – 1807602
- Total identified (soft trigger) – 18774

97% empty frames

34.71% identified as beam events

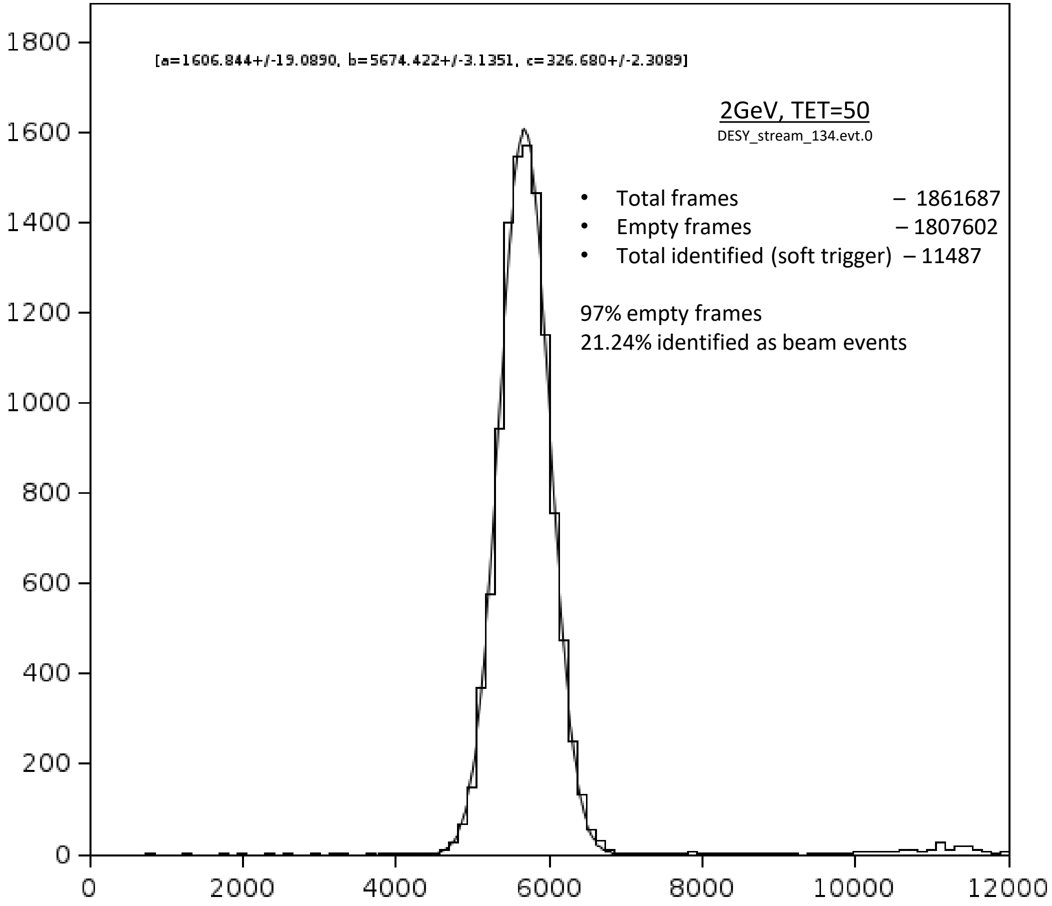

2GeV, TET=50

DESY_stream_134.evt.0

- Total frames – 1861687
- Empty frames – 1807602
- Total identified (soft trigger) – 11487

97% empty frames

21.24% identified as beam events

fADC threshold effects (2GeV)

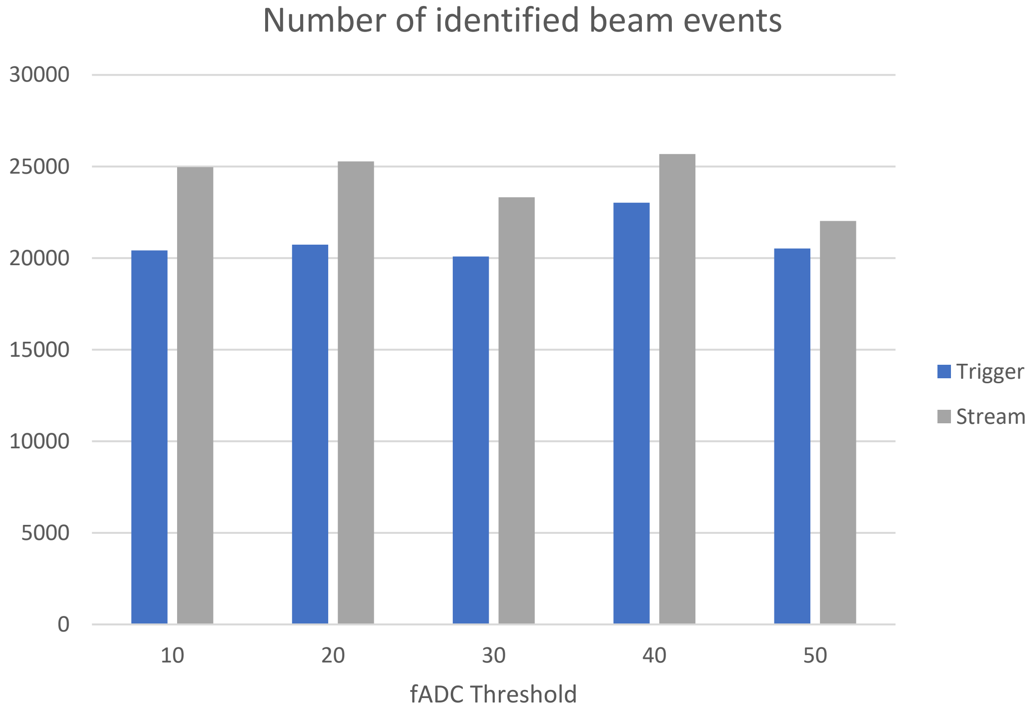
Number of identified beam events

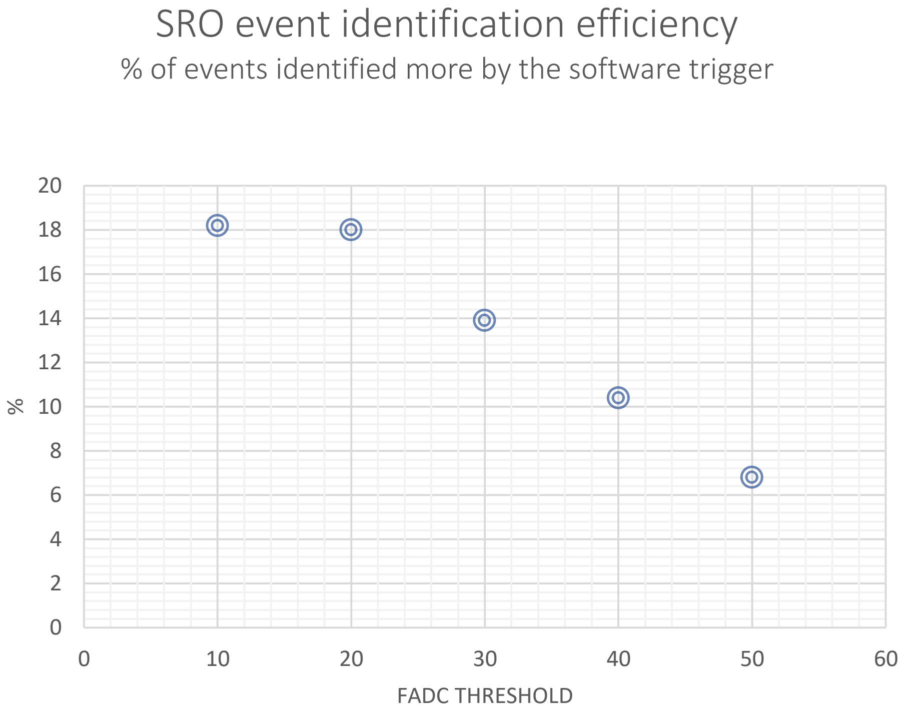
SRO event identification efficiency

% of events identified more by the software trigger

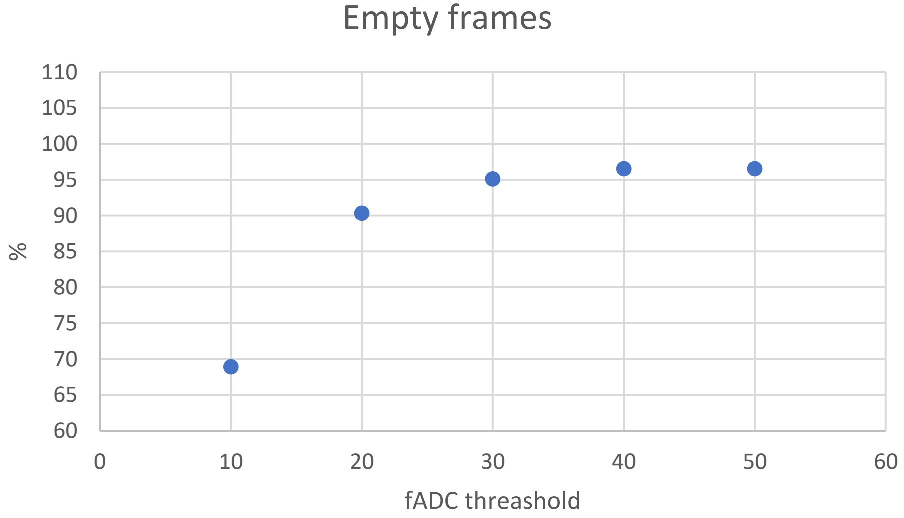
Empty frames

Calorimeter response linearity

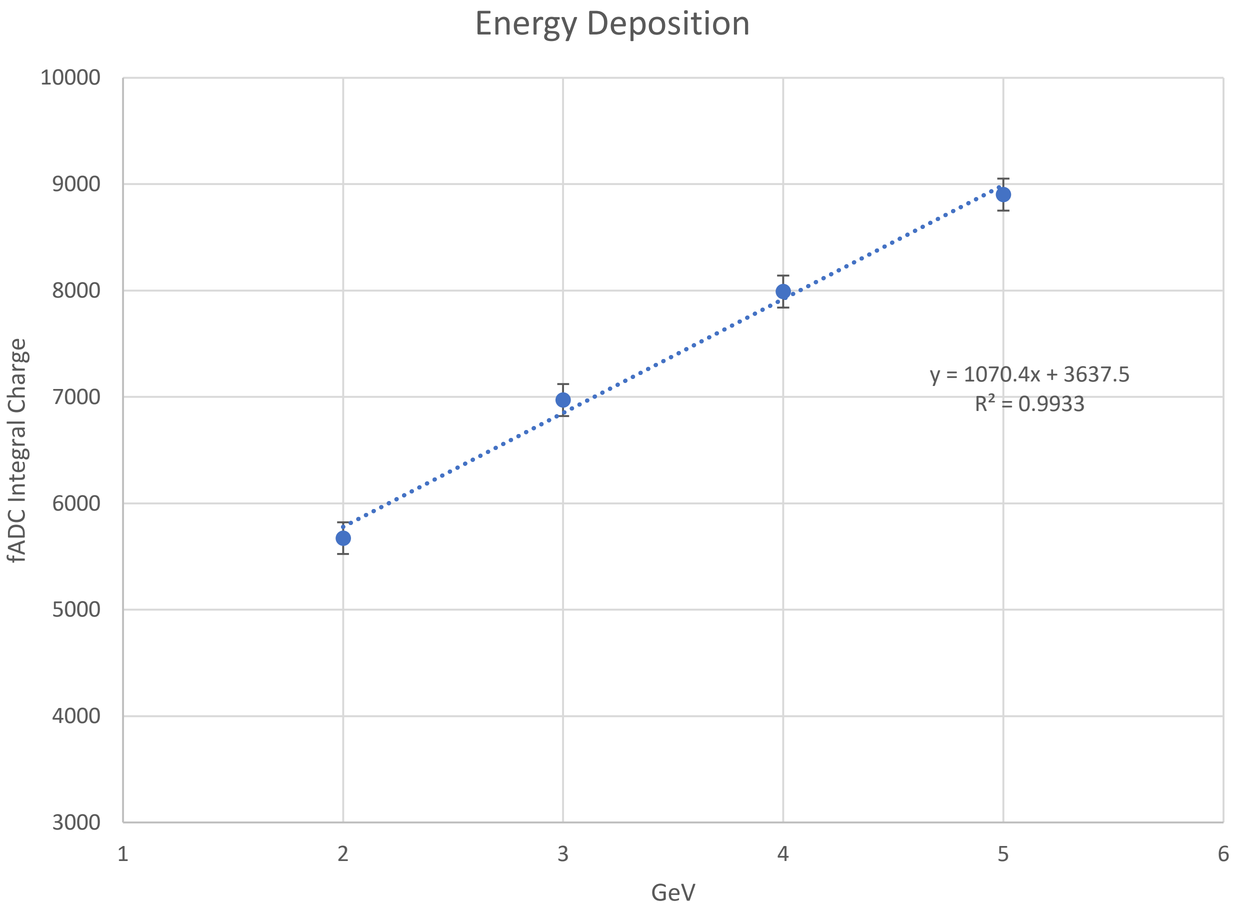

Summary

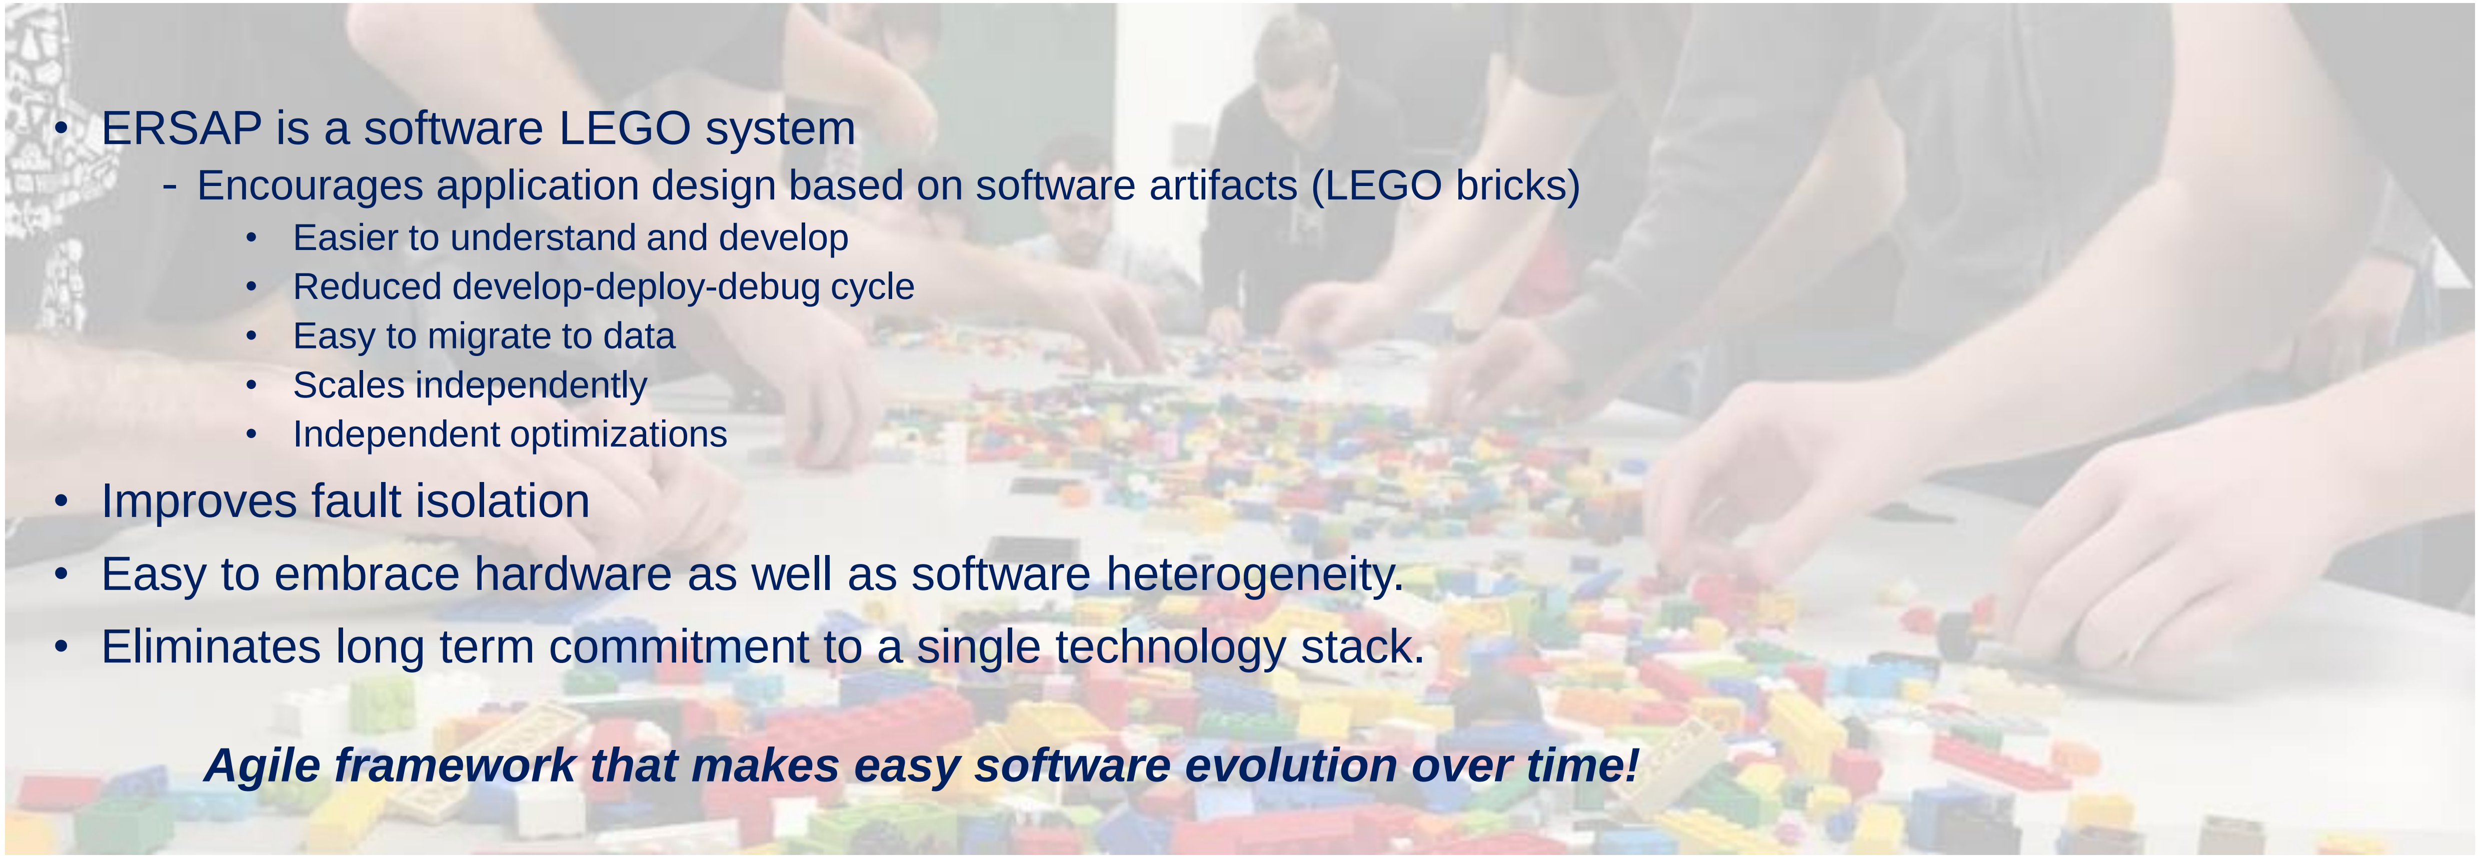

- ERSAP is a software LEGO system
  - Encourages application design based on software artifacts (LEGO bricks)
    - Easier to understand and develop
    - Reduced develop-deploy-debug cycle
    - Easy to migrate to data
    - Scales independently
    - Independent optimizations
- Improves fault isolation
- Easy to embrace hardware as well as software heterogeneity.
- Eliminates long term commitment to a single technology stack.

Agile framework that makes easy software evolution over time!

Current status and future plans

- ERSAP is a reactive actor/micro-service based data-stream processing framework. https://wiki.jlab.org/epsciwiki/index.php/ERSAP
- Combines decade-long experience: CODA, AFECS and CLARA
  - ERSAP Java binding, betta release: https://github.com/JeffersonLab/ersap-java.git
  - ERSAP C++ binding development in progress: https://github.com/JeffersonLab/ersap-cpp.git
  - ERSAP Python binding in the design stage
  - Plans to design ERSAP Julia binding
- Many ERSAP engine development projects are in progress
  - CODA engines: https://github.com/JeffersonLab/ersap-coda.git
  - JANA2 based engines: https://github.com/JeffersonLab/ersap-jana.git
  - TriDAS engines: https://github.com/JeffersonLab/ersap-tridas.git
  - CLAS12 AI reconstruction engines https://github.com/JeffersonLab/ersap-vtp.git
  - INDRA ASTRA project ML engines
- Collaborative effort between JLAB Physics and CST divisions.

Thank You

Streaming CODA and ERSAP to achieve data stream acquisition and processing

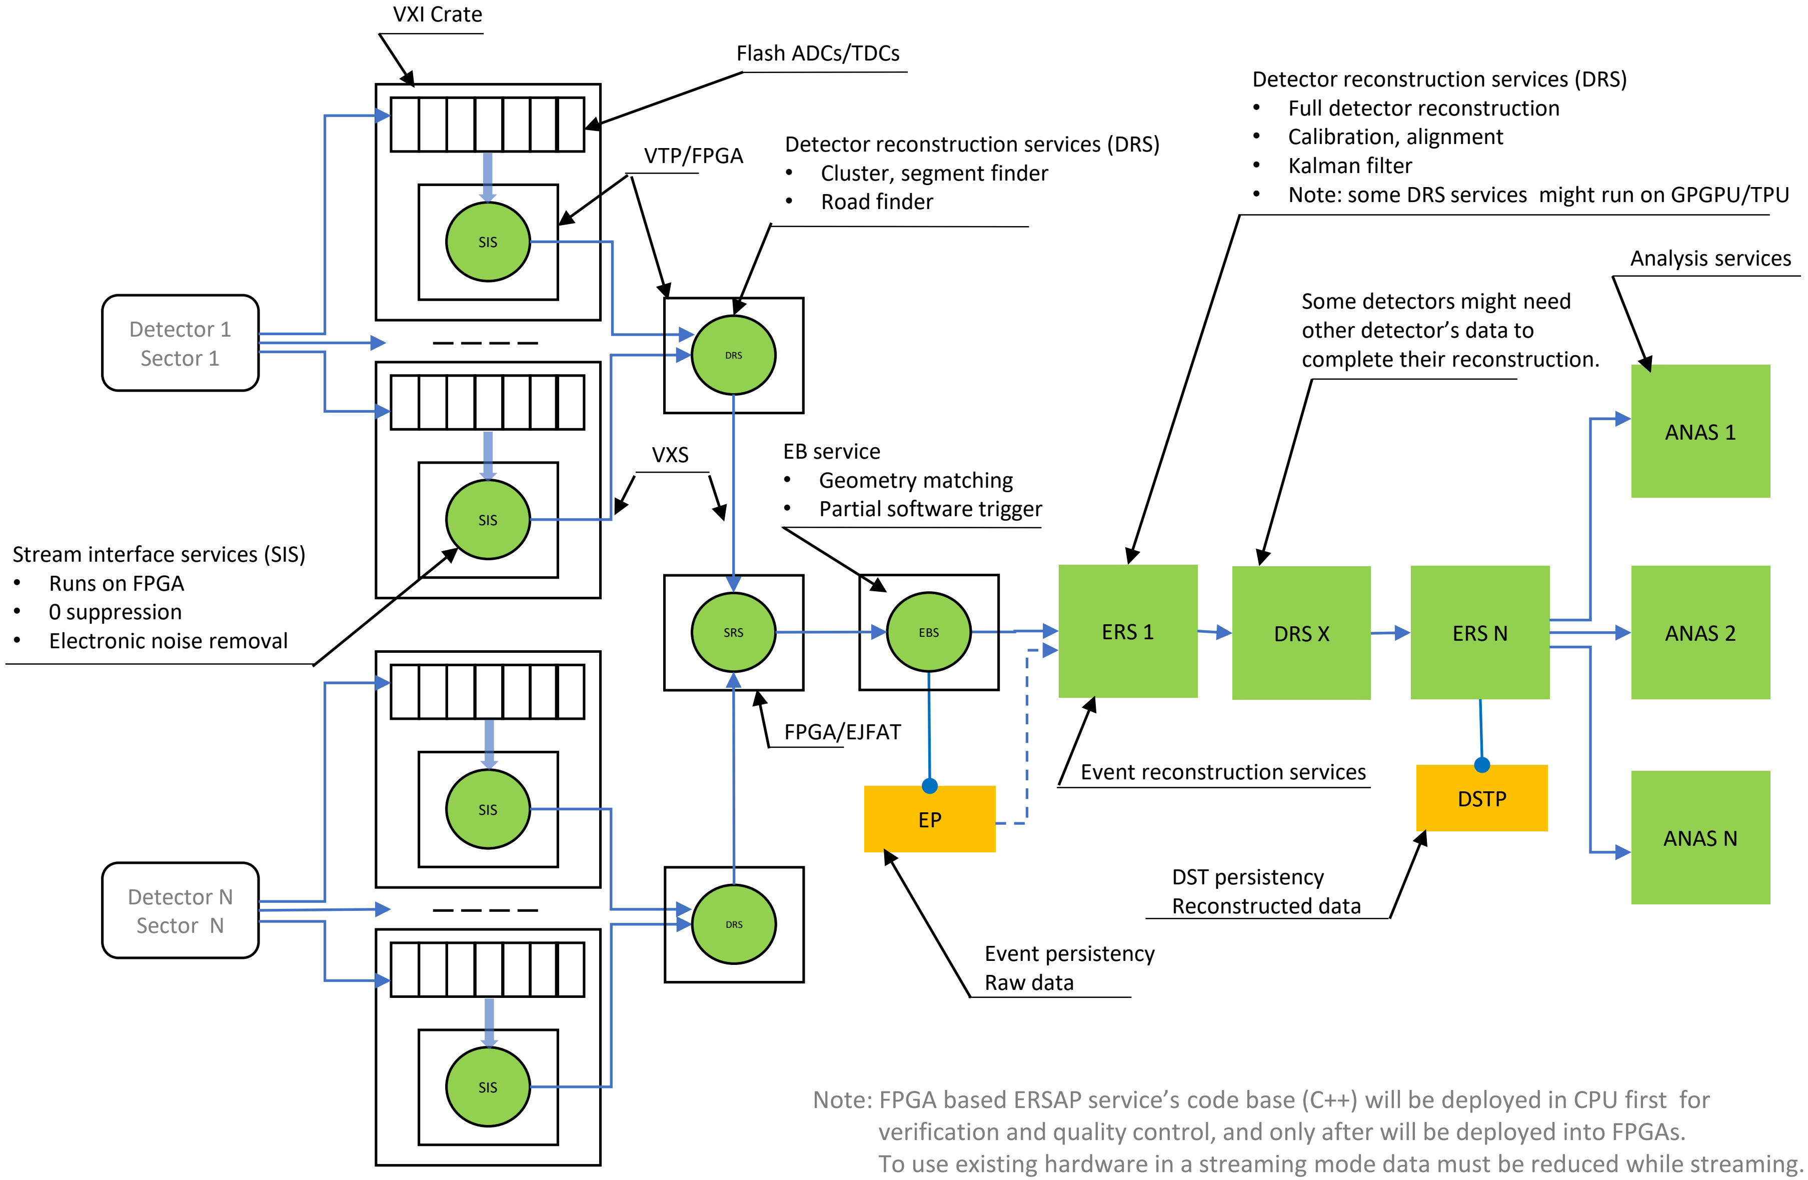

Detector reconstruction services (DRS)
- Cluster, segment finder
- Road finder

Detector reconstruction services (DRS)
- Full detector reconstruction
- Calibration, alignment
- Kalman filter
- Note: some DRS services might run on GPGPU/TPU

Some detectors might need other detector's data to complete their reconstruction.

EB service
- Geometry matching
- Partial software trigger

Stream interface services (SIS)
- Runs on FPGA
- 0 suppression
- Electronic noise removal

Event reconstruction services

DST persistency
Reconstructed data

Event persistency
Raw data

Note: FPGA based ERSAP service's code base (C++) will be deployed in CPU first for verification and quality control, and only after will be deployed into FPGAs.
To use existing hardware in a streaming mode data must be reduced while streaming.

Hardware vs. software event identification

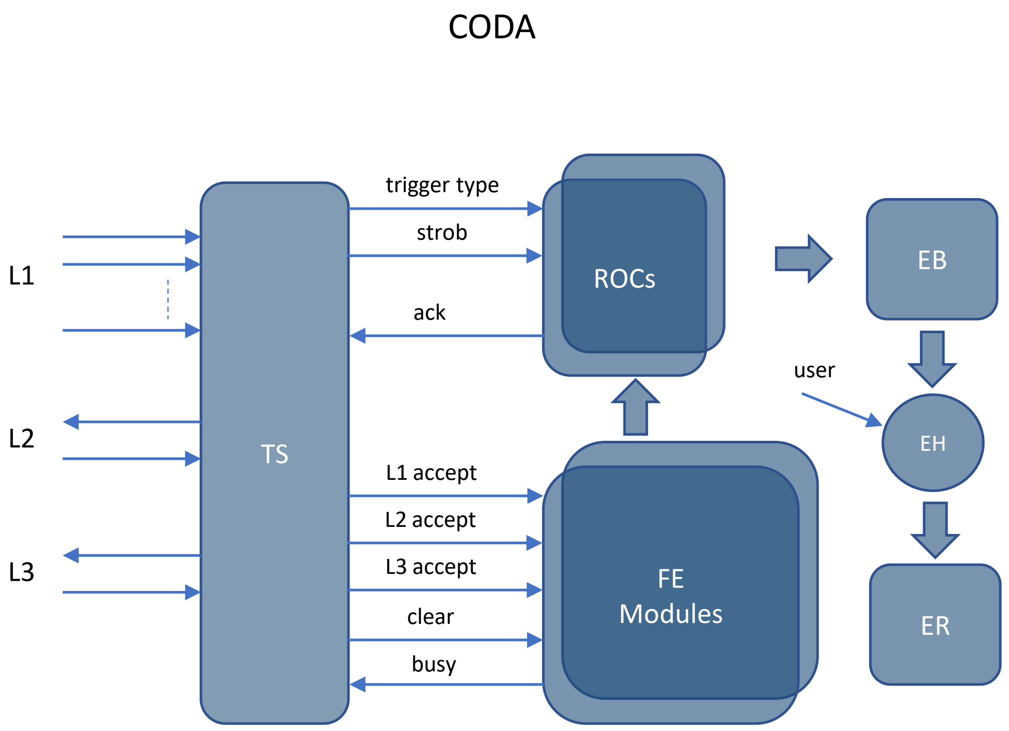

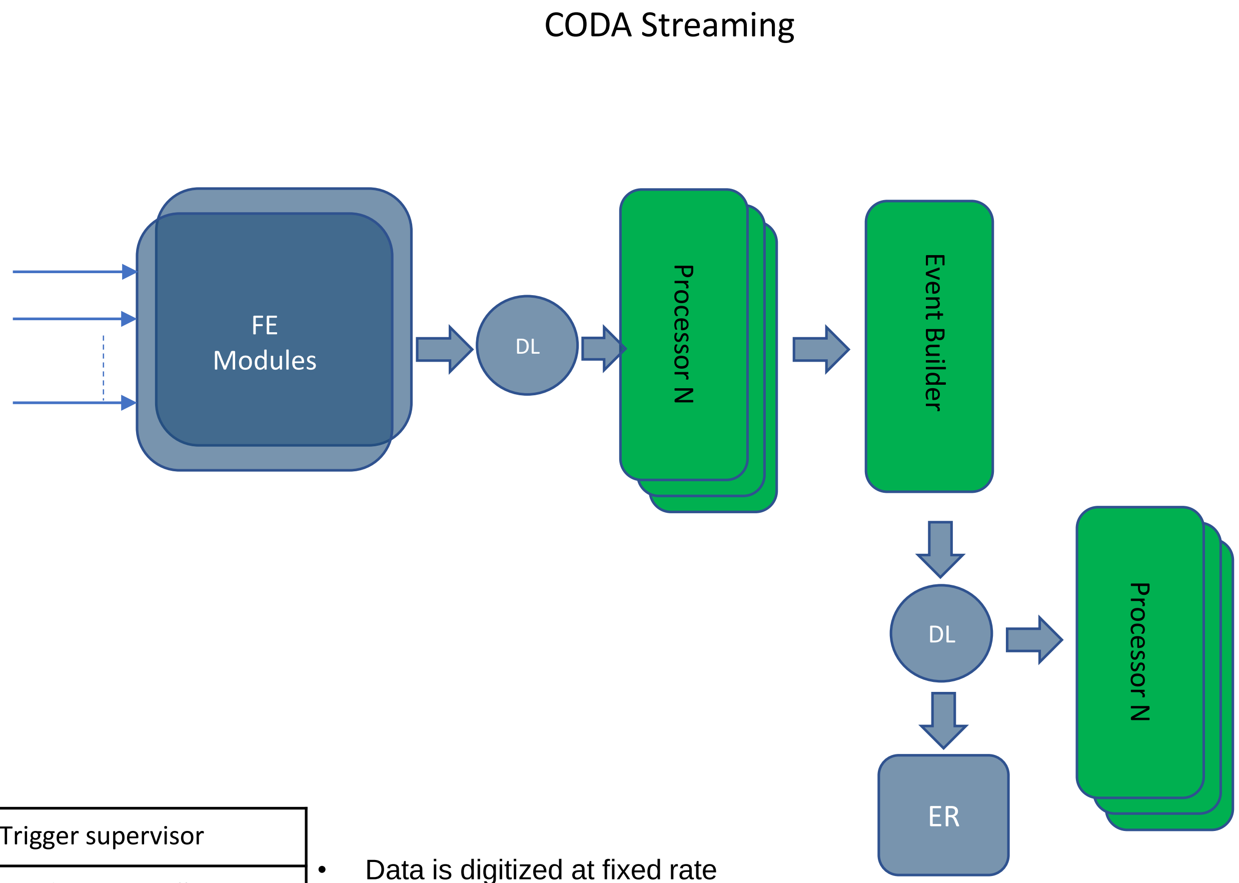

- Challenging to deal with event-pileups
- Not ideal to read general purpose detectors
- Cary bias to low-energy particles

TS	Trigger supervisor
ROC	Readout controller
EB	Event Builder
EH	Event Hub
ER	Event Recorder
DL	Data Lake/tiered storage

- Data is digitized at fixed rate
- Data is read out continuous parallel streams
- Data is cooled down at the tiered storages

fADC threshold effect

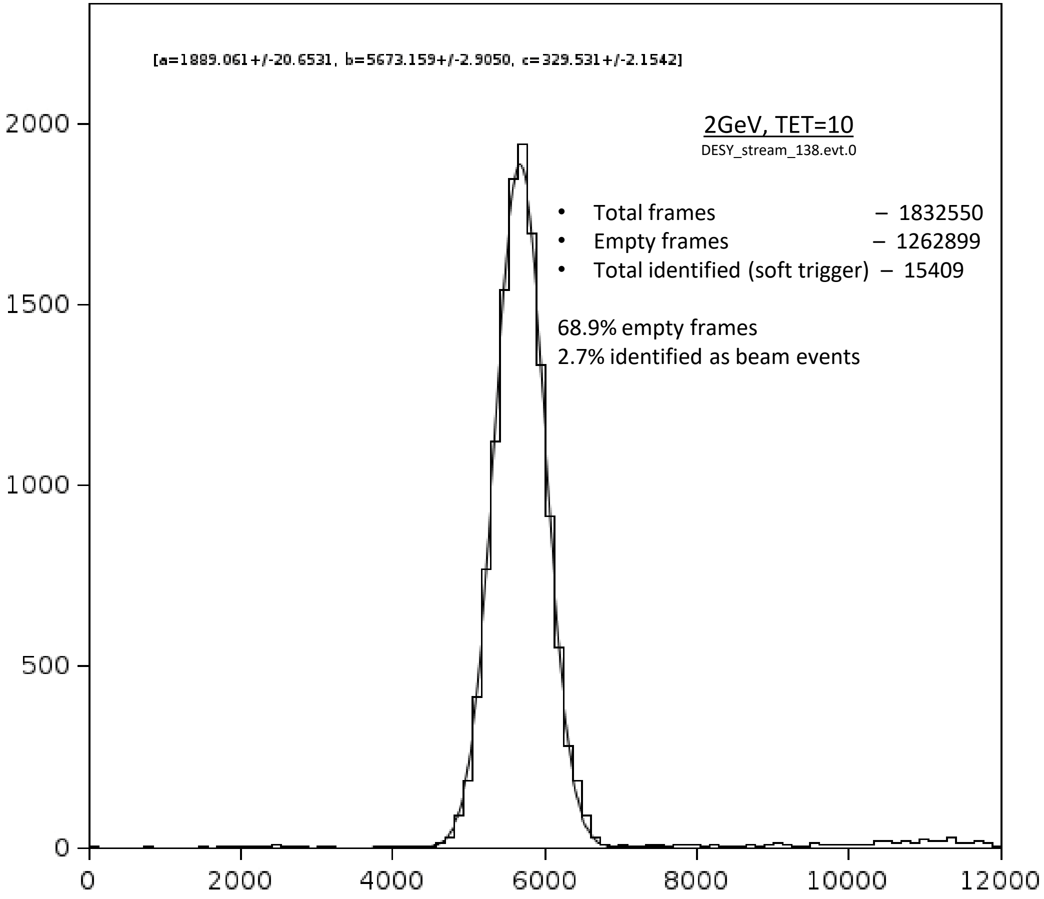

2GeV, TET=10

DESY_stream_138.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 15409

68.9% empty frames
2.7% identified as beam events

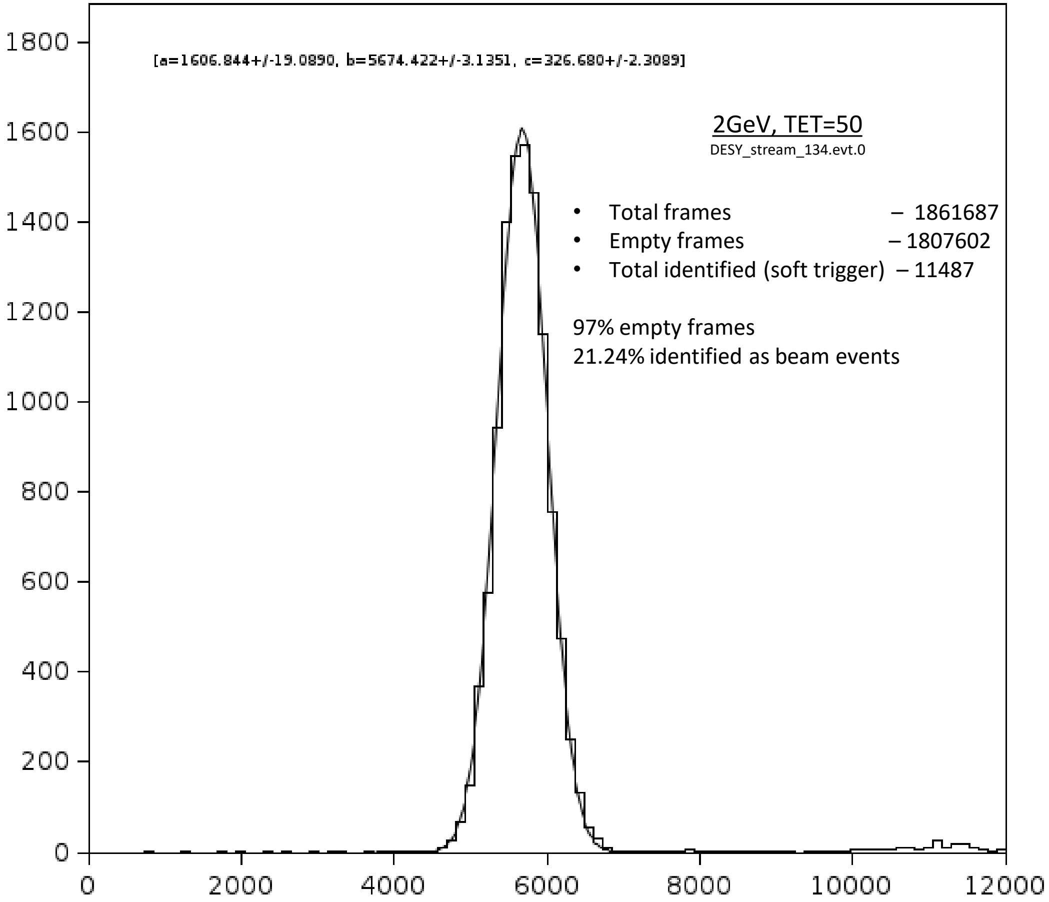

2GeV, TET=50

DESY_stream_134.evt.0

- Total frames – 1861687
- Empty frames – 1807602
- Total identified (soft trigger) – 11487

97% empty frames
21.24% identified as beam events

No charge cut in the soft trigger

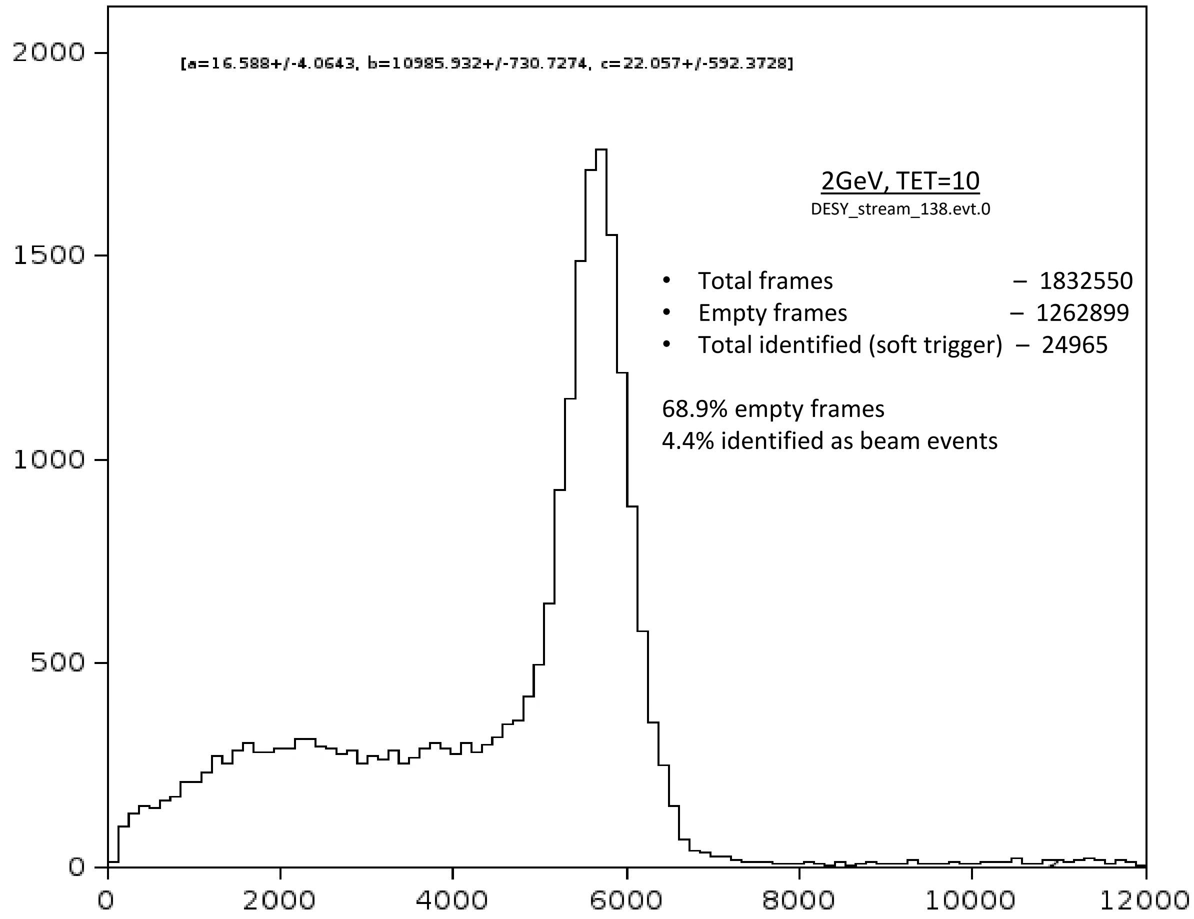

2GeV, TET=10

DESY_stream_138.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 24965

68.9% empty frames

4.4% identified as beam events

No energy/charge cut in the soft trigger

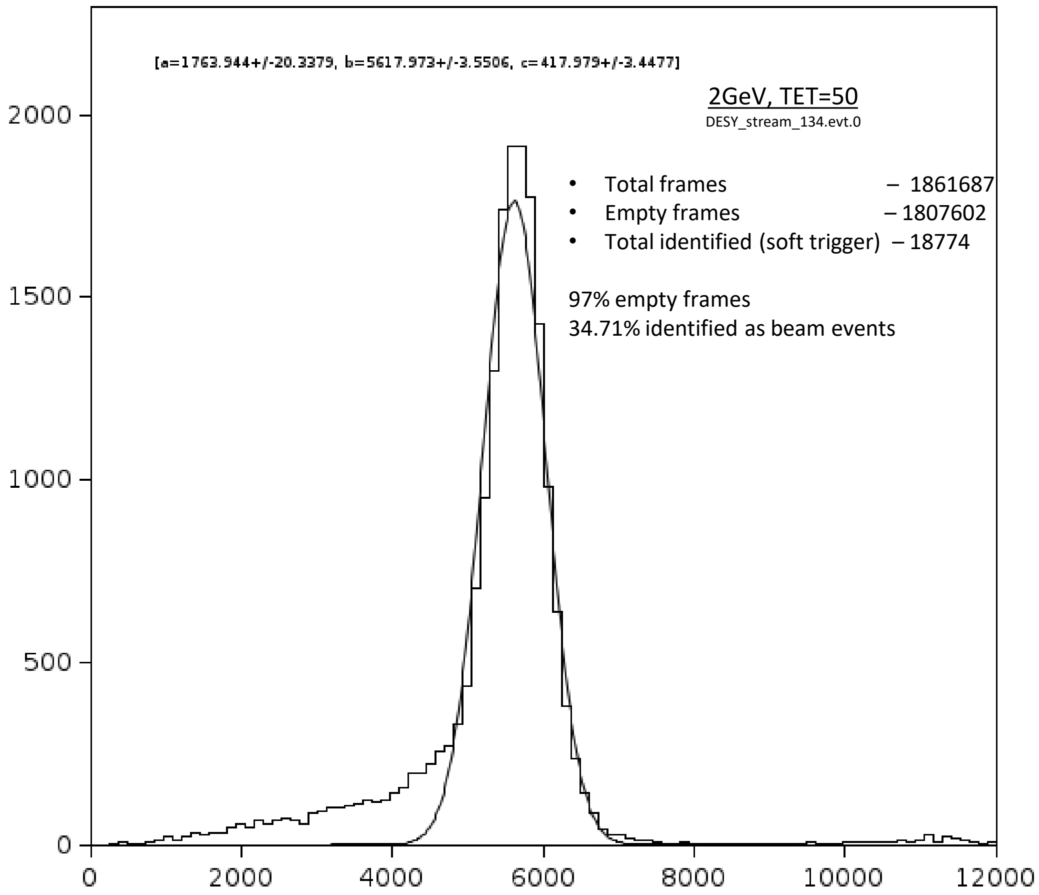

2GeV, TET=50

DESY_stream_134.evt.0

- Total frames – 1861687
- Empty frames – 1807602
- Total identified (soft trigger) – 18774

97% empty frames

34.71% identified as beam events

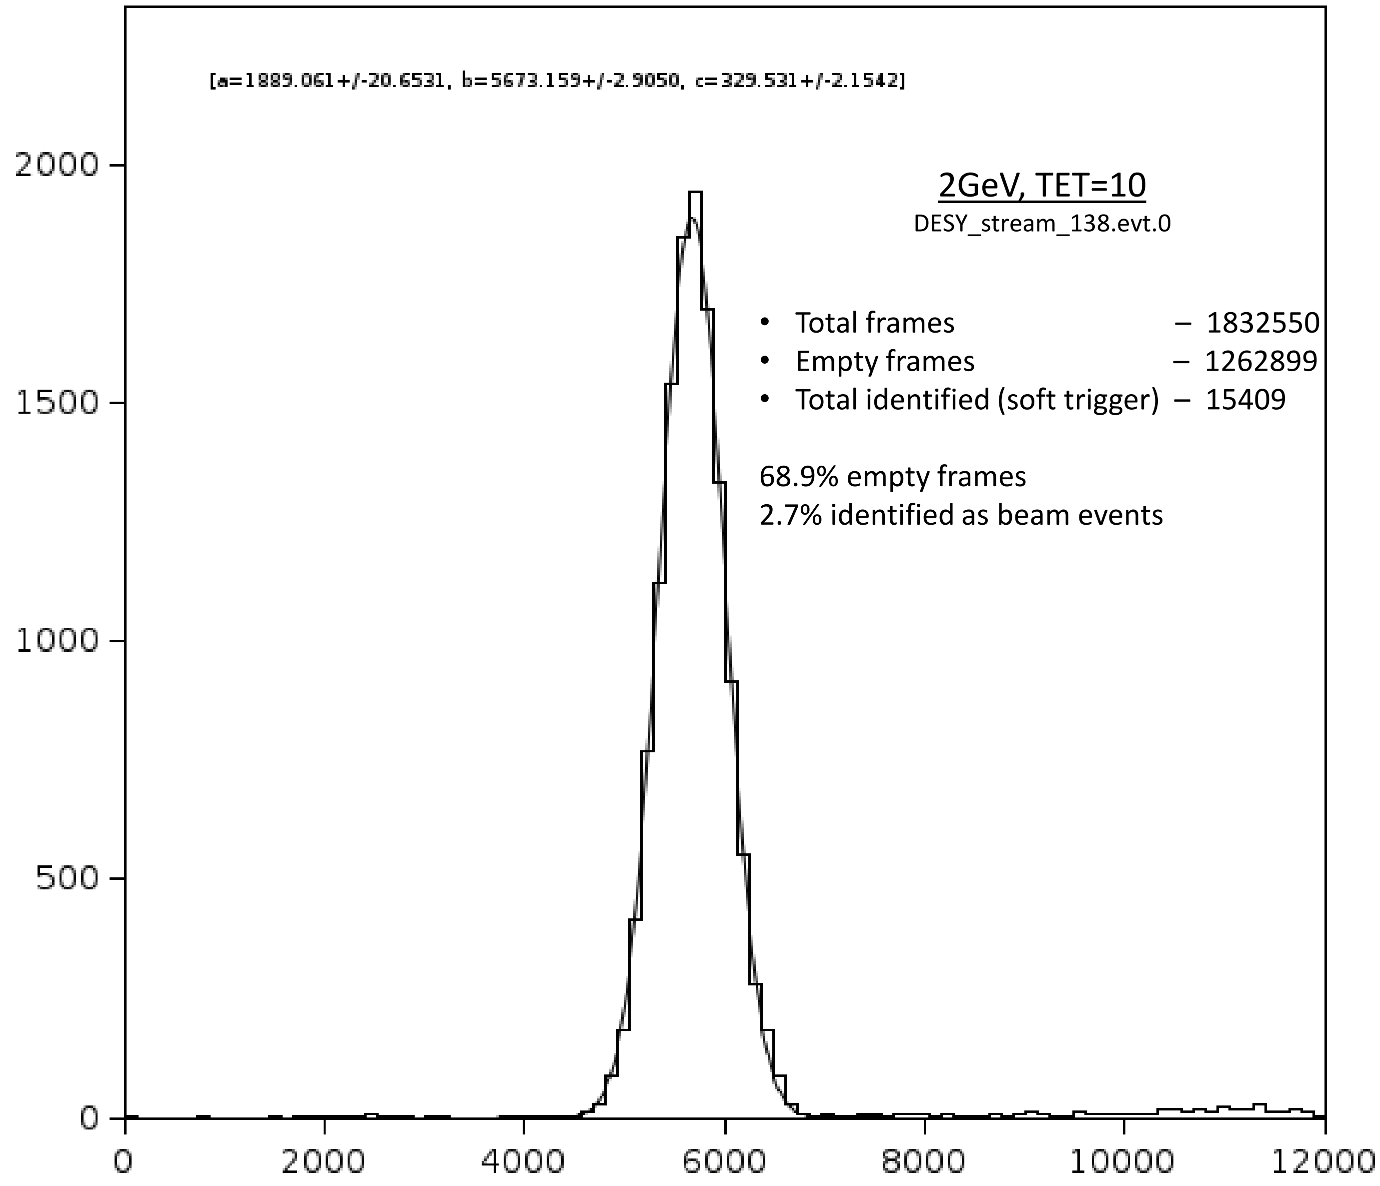

2GeV, TET=10

DESY_stream_138.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 15409

68.9% empty frames
2.7% identified as beam events

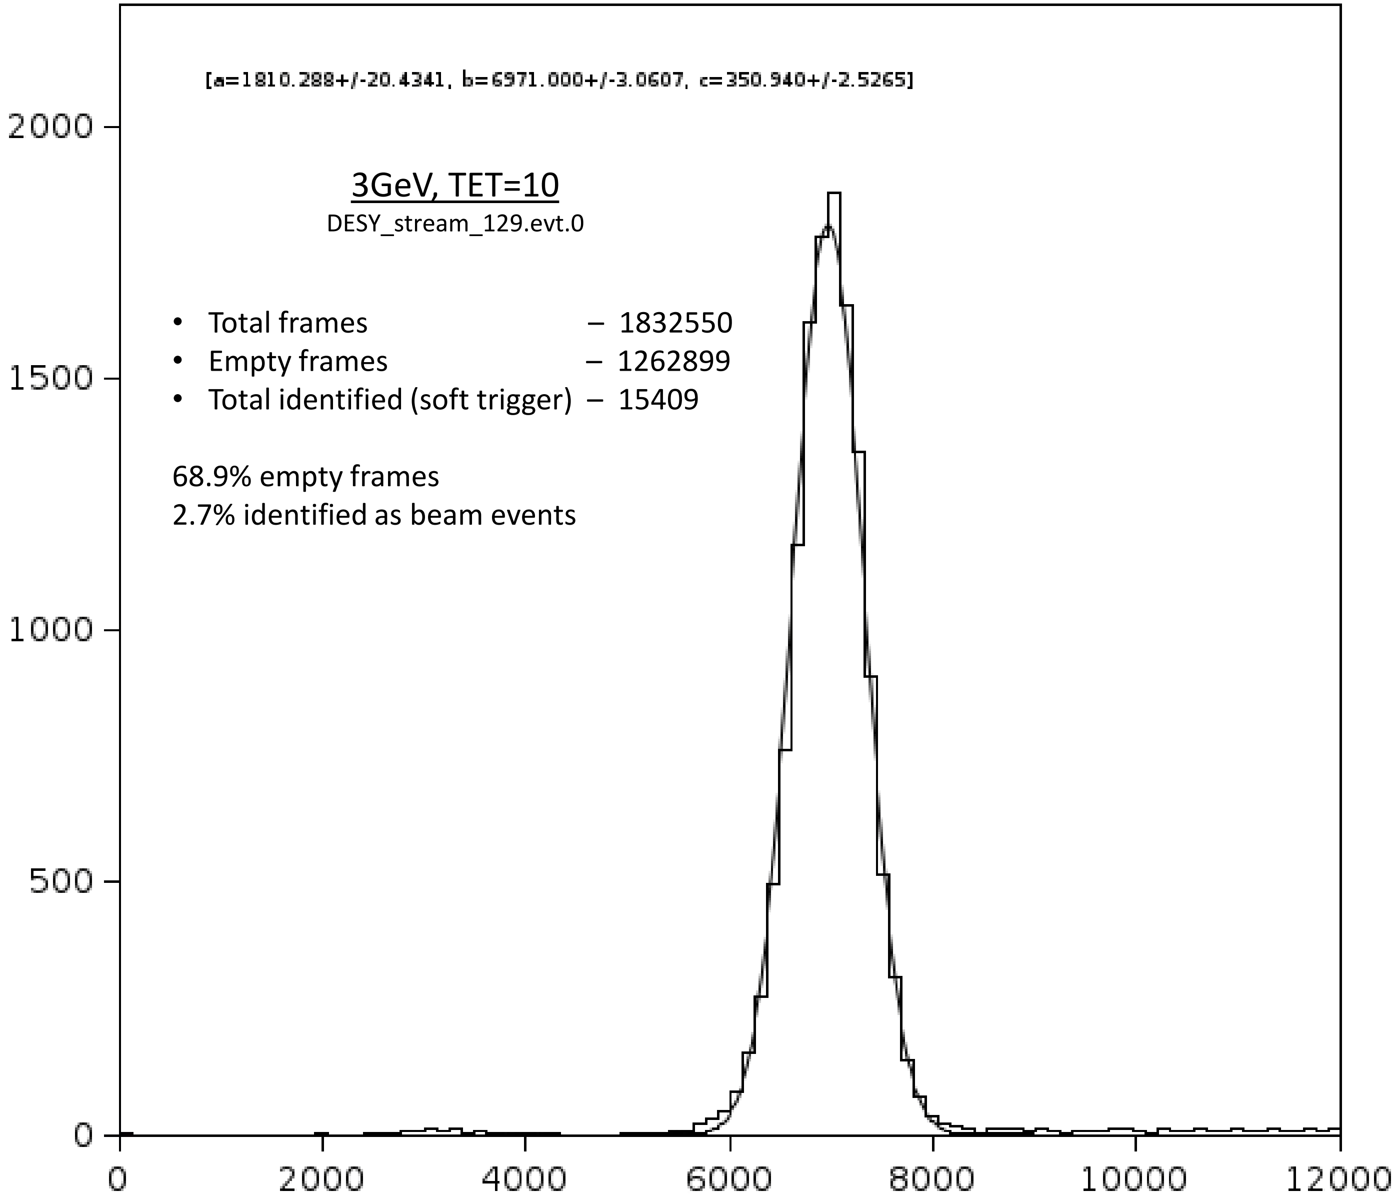

[a=1810.288+/-20.4341, b=6971.000+/-3.0607, c=350.940+/-2.5265]

3GeV, TET=10

DESY_stream_129.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 15409

68.9% empty frames
2.7% identified as beam events

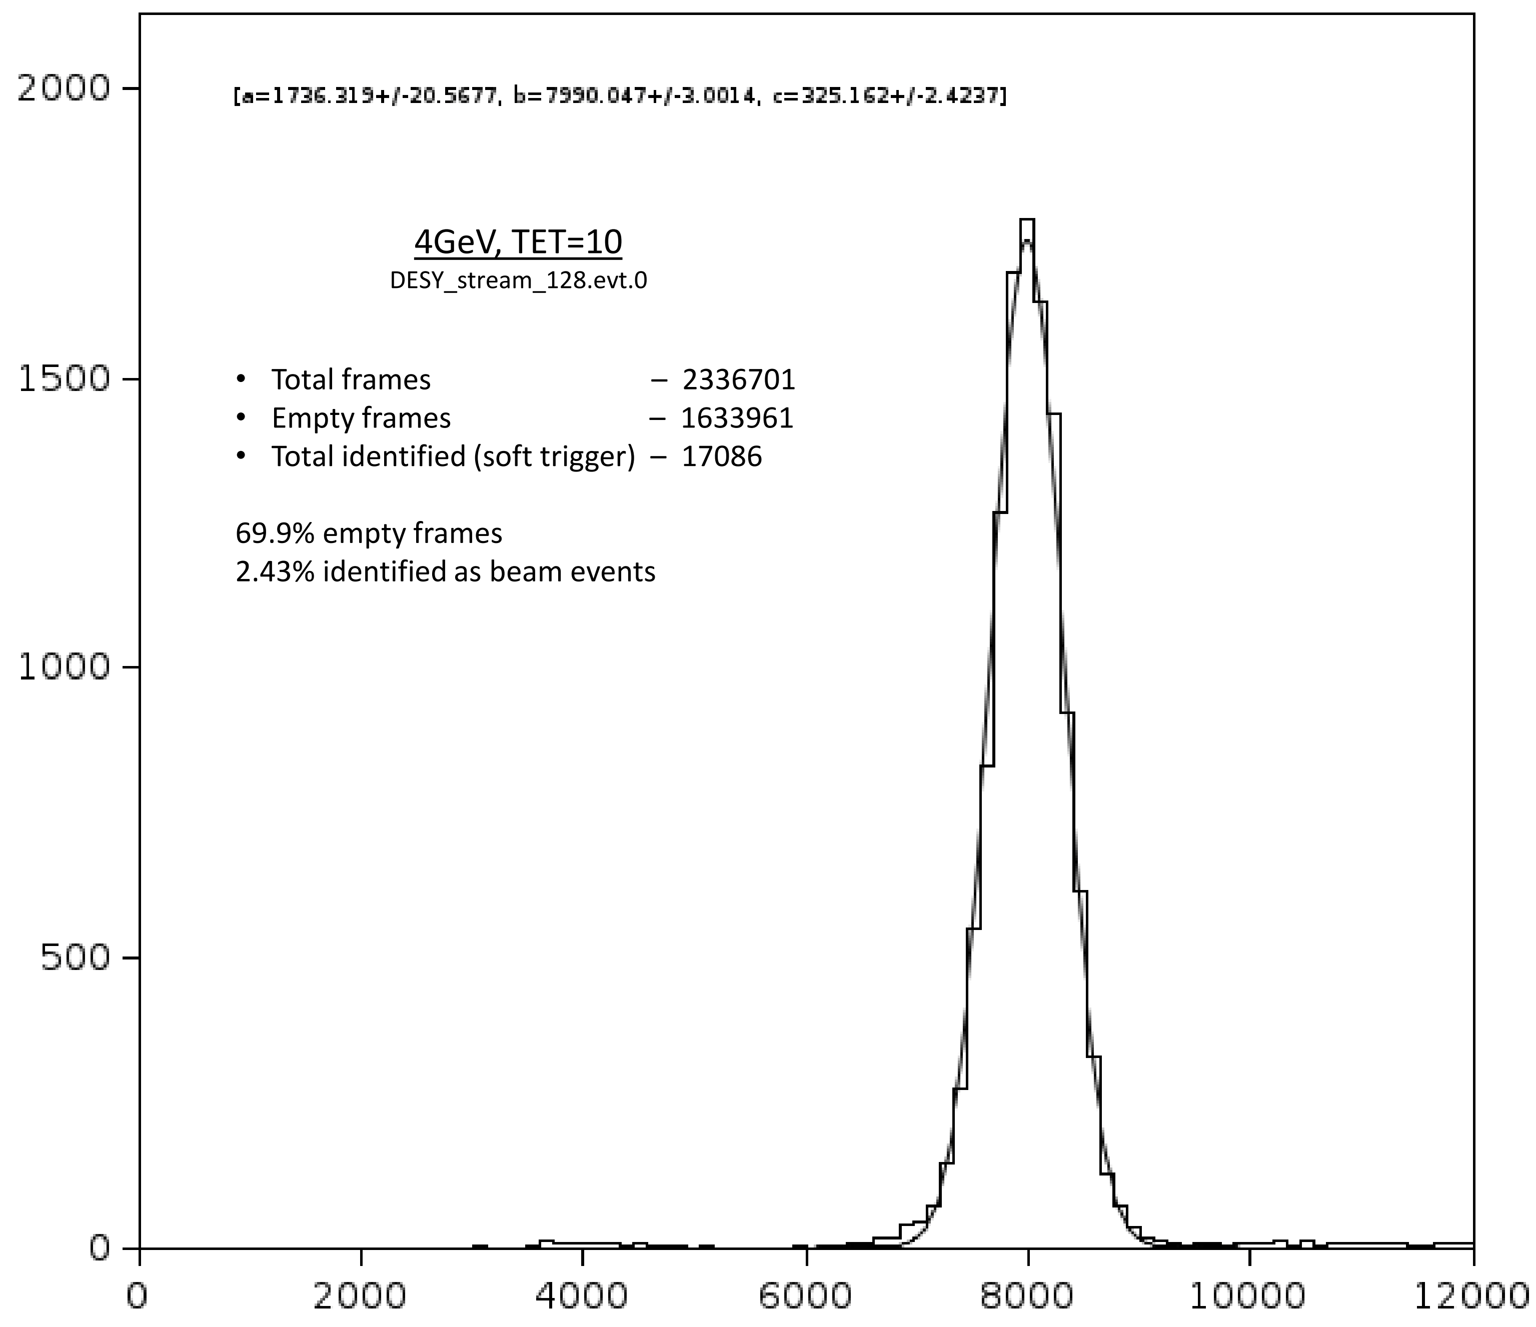

4GeV, TET=10

DESY_stream_128.evt.0

- Total frames – 2336701
- Empty frames – 1633961
- Total identified (soft trigger) – 17086

69.9% empty frames
2.43% identified as beam events

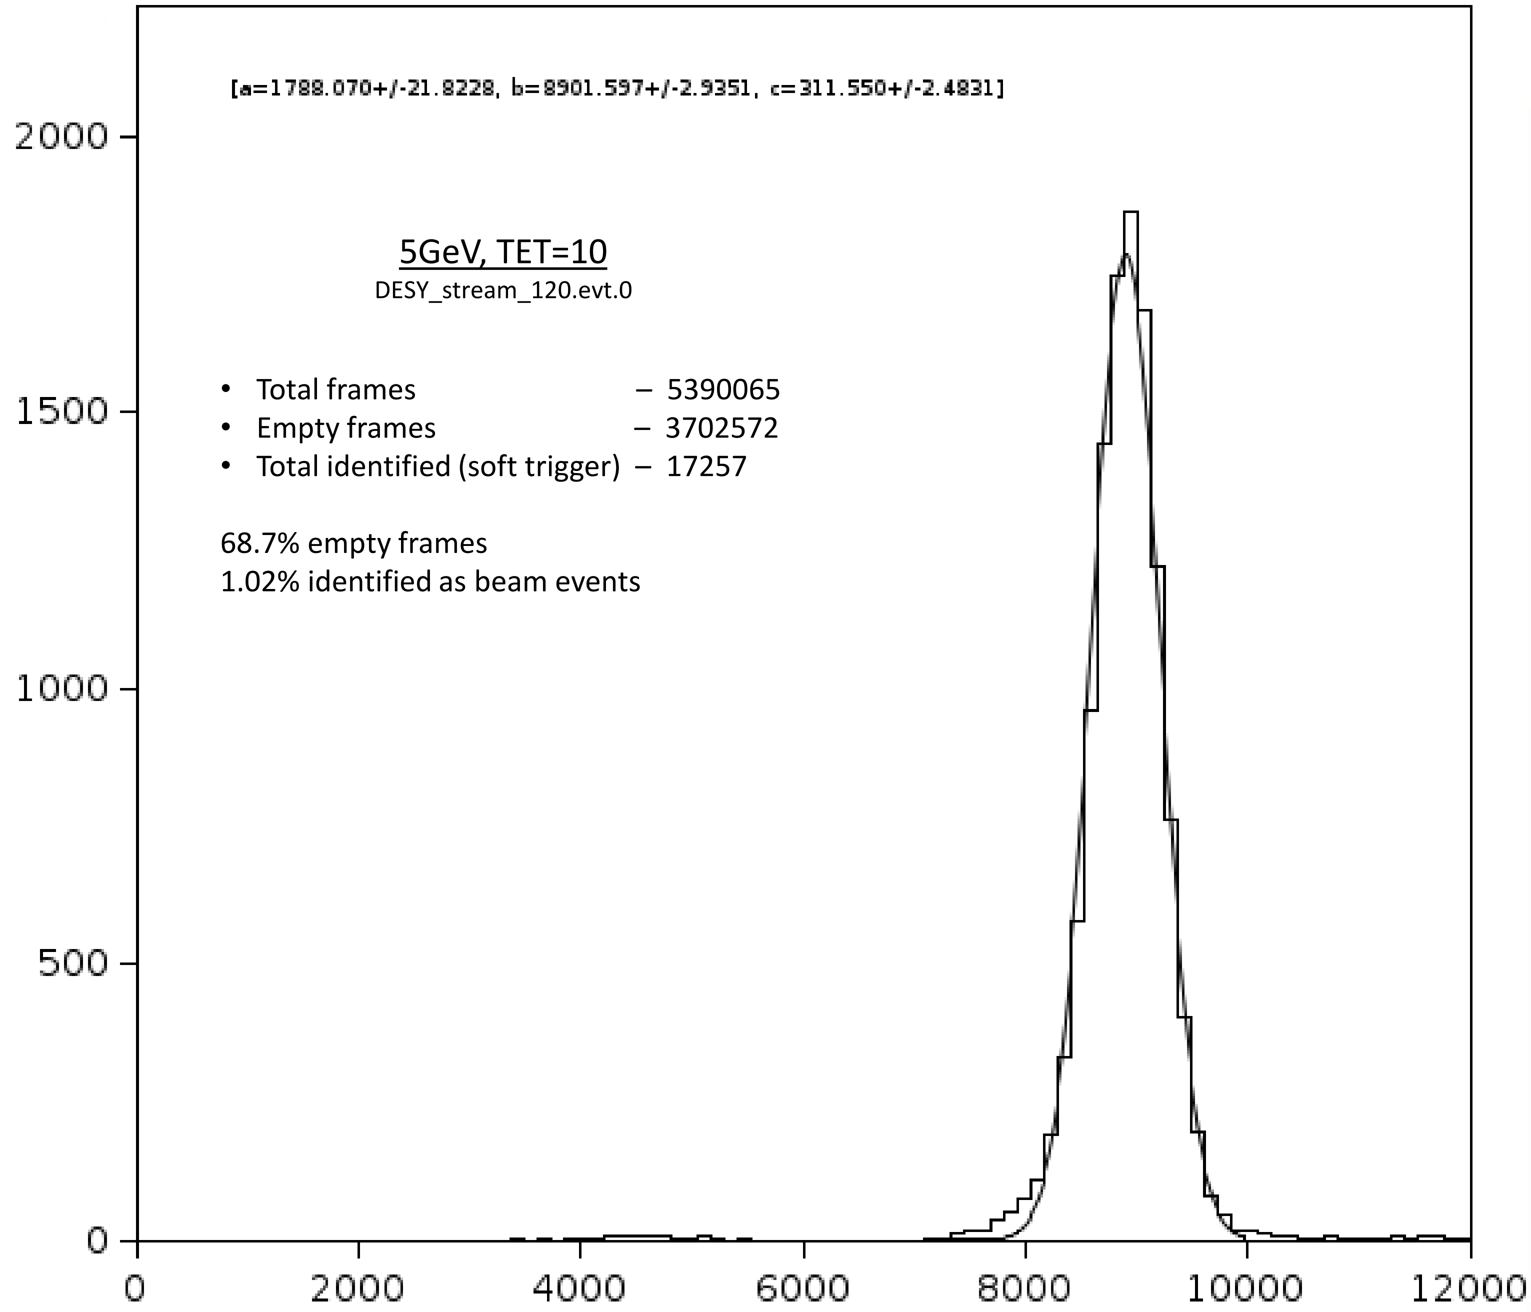

5GeV, TET=10

DESY_stream_120.evt.0

- Total frames – 5390065
- Empty frames – 3702572
- Total identified (soft trigger) – 17257

68.7% empty frames
1.02% identified as beam events

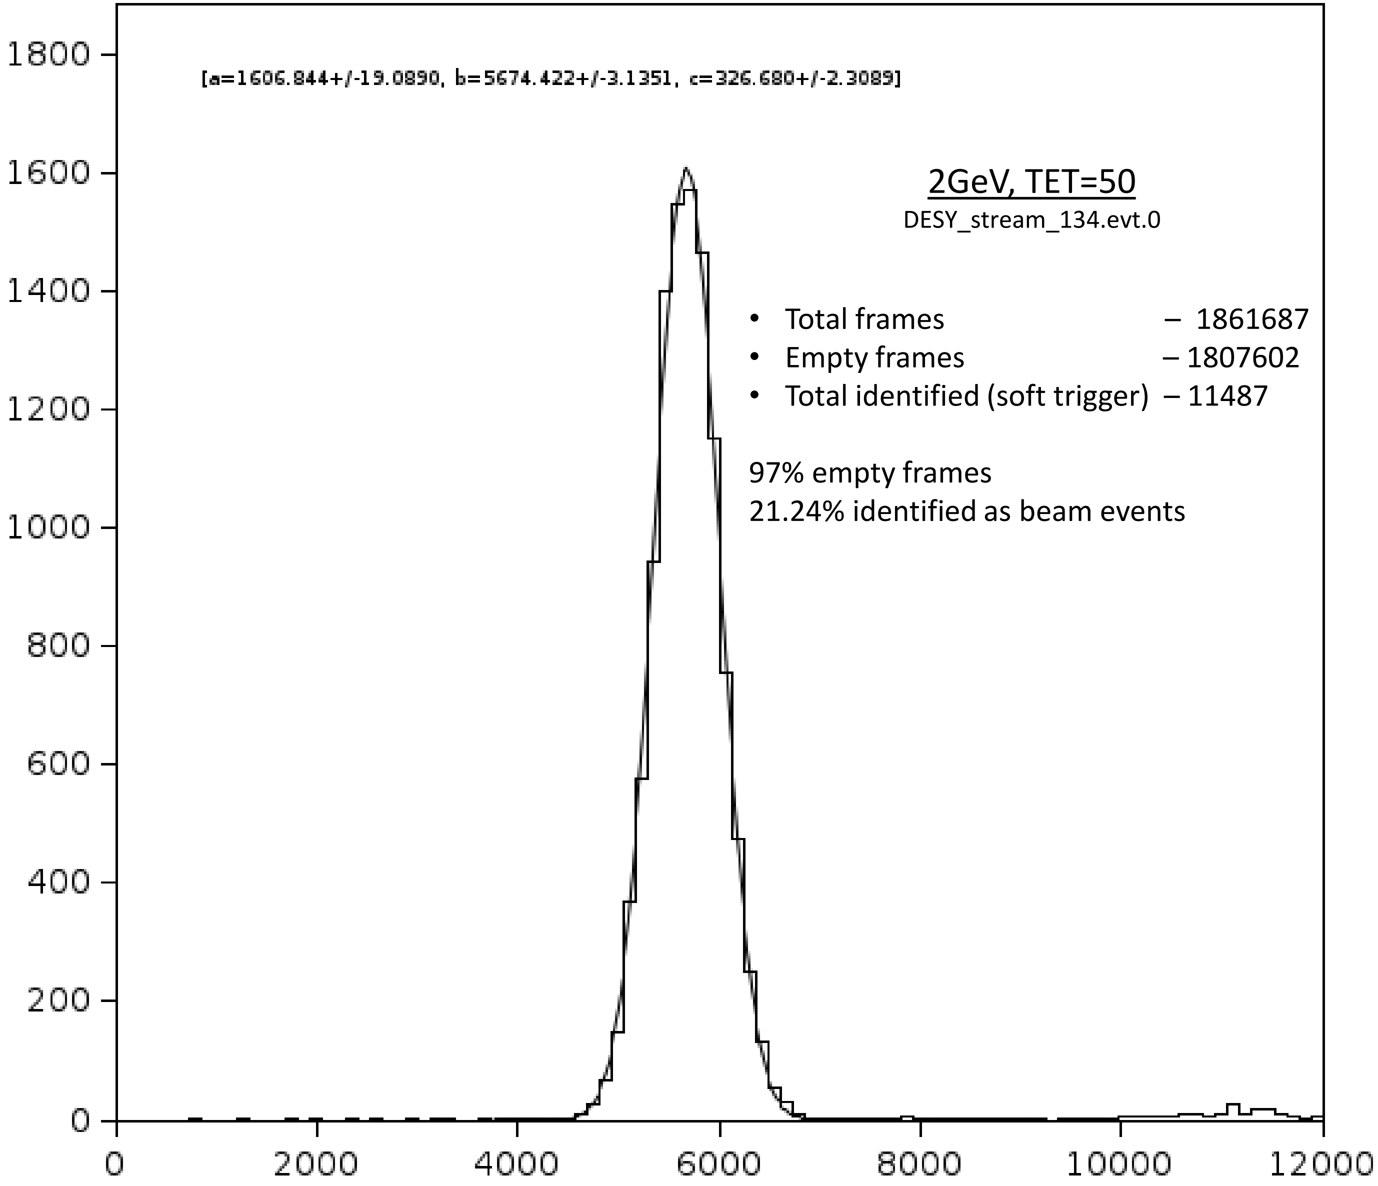

2GeV, TET=50

DESY_stream_134.evt.0

- Total frames – 1861687
- Empty frames – 1807602
- Total identified (soft trigger) – 11487

97% empty frames
21.24% identified as beam events

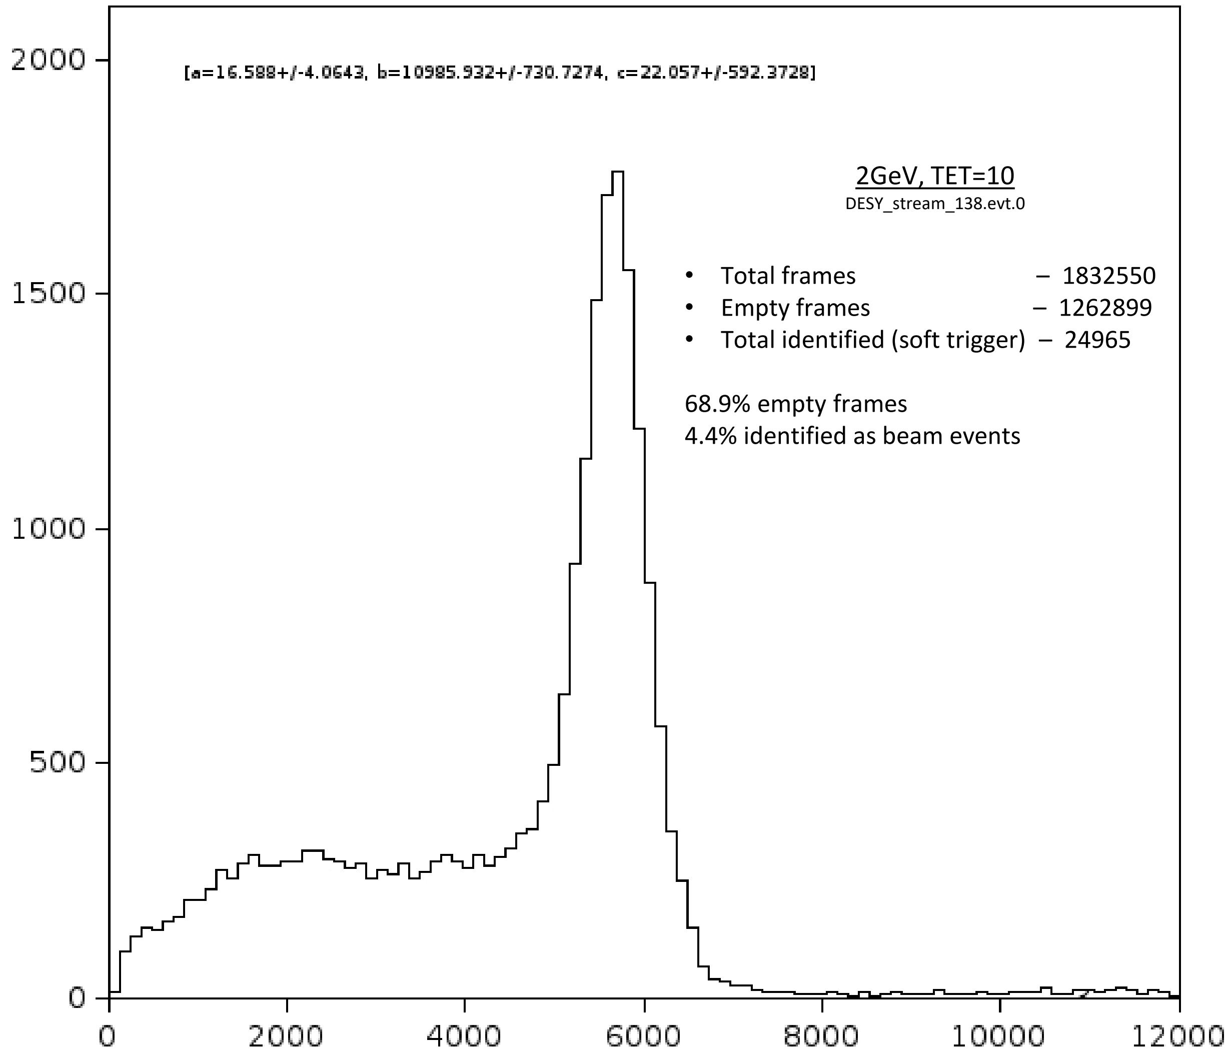

[a=16.588+/-4.0643, b=10985.932+/-730.7274, c=22.057+/-592.3728]

2GeV, TET=10

DESY_stream_138.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 24965

68.9% empty frames
4.4% identified as beam events

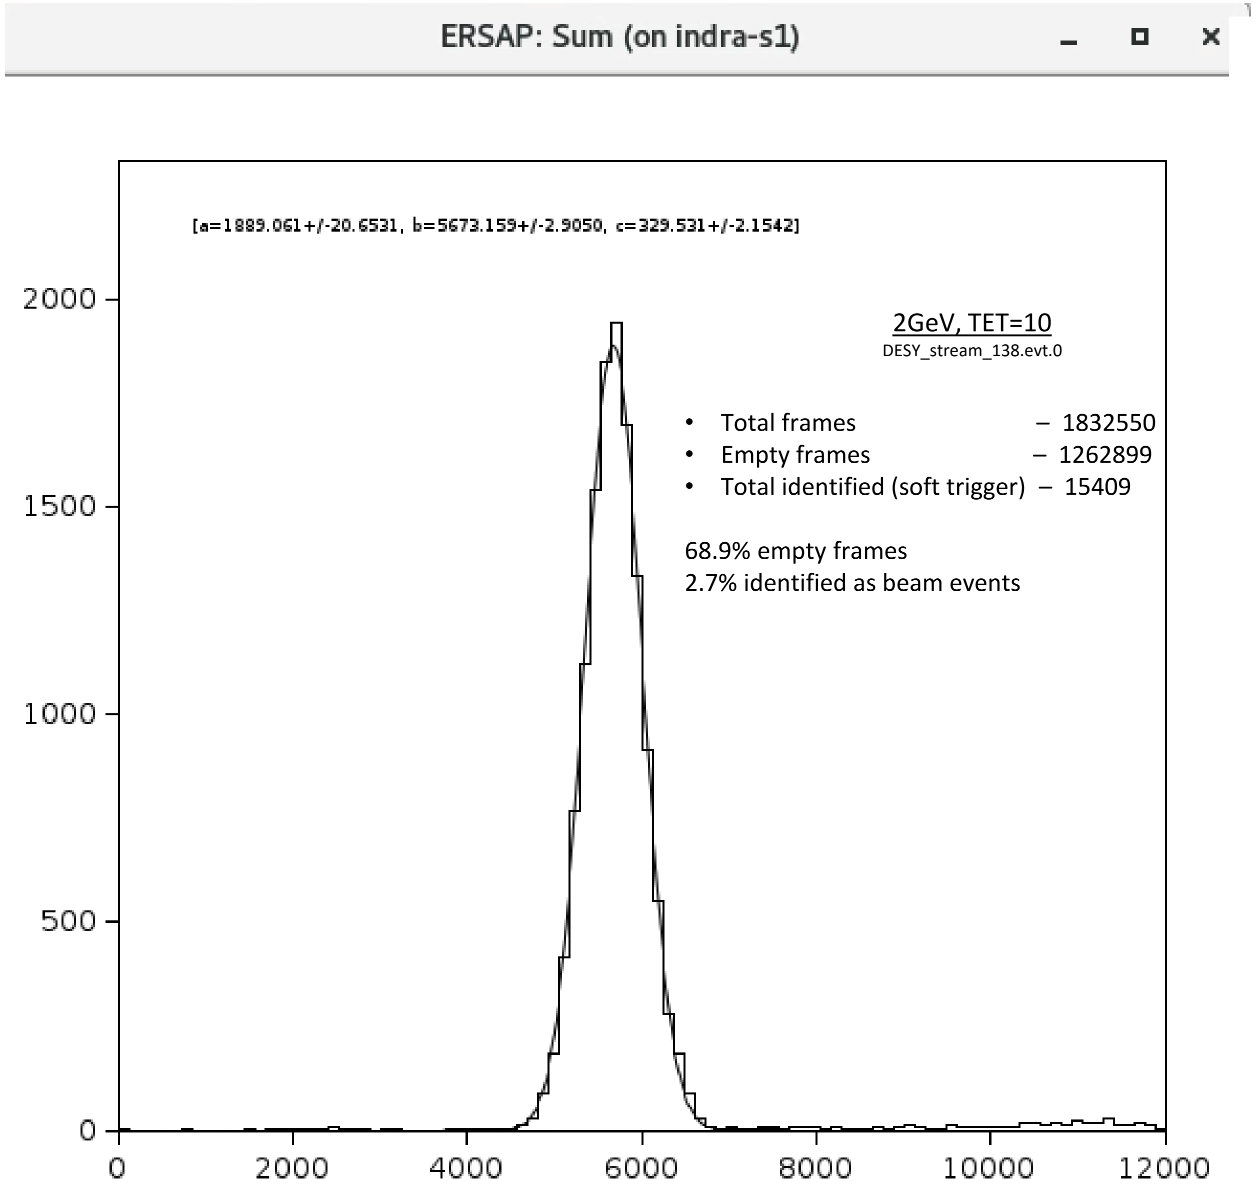

[a=1889.061+/-20.6531, b=5673.159+/-2.9050, c=329.531+/-2.1542]

2GeV, TET=10

DESY_stream_138.evt.0

- Total frames – 1832550
- Empty frames – 1262899
- Total identified (soft trigger) – 15409

68.9% empty frames
2.7% identified as beam events

$PbWO_4$ Christal based 3x3 Calorimeter

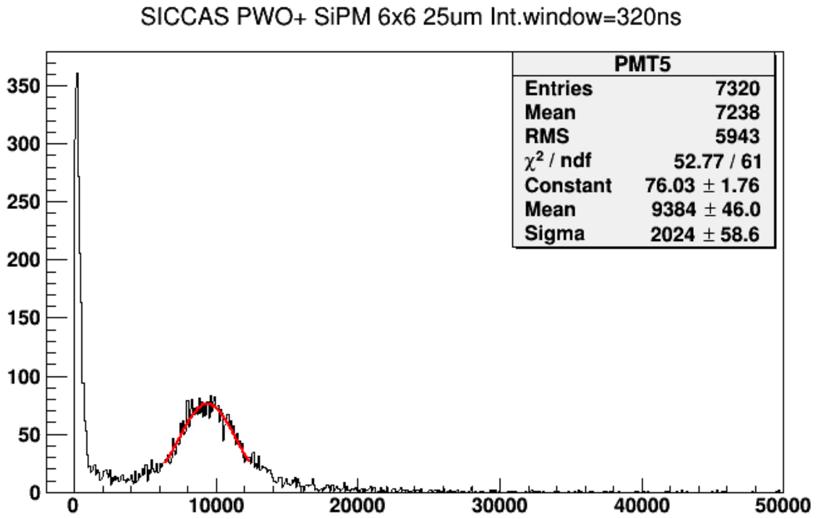
SICCAS PWO+ SiPM 6x6 25um Int.window=320ns
PMT5
Entries 7320
Mean 7238
RMS 5943
χ² / ndf 52.77 / 61
Constant 76.03 ± 1.76
Mean 9384 ± 46.0
Sigma 2024 ± 58.6