About Jefferson Lab

- Jefferson Lab is a world leading facility for the study of Hadronic Physics through accelerator experiments funded by the US DOE Office of Science
- It is a user facility operating a 12-GeV continuous electron beam accelerator
- Experiments are carried out in 4 Experimental Halls (A,B,C & D)
- Experiments are supported by the Theory Center, and Scientific Computing
- Jefferson Lab is also a world leader in accelerator design and RF technology
- Jefferson Lab is one of the contender sites for a forthcoming Electron Ion Collider

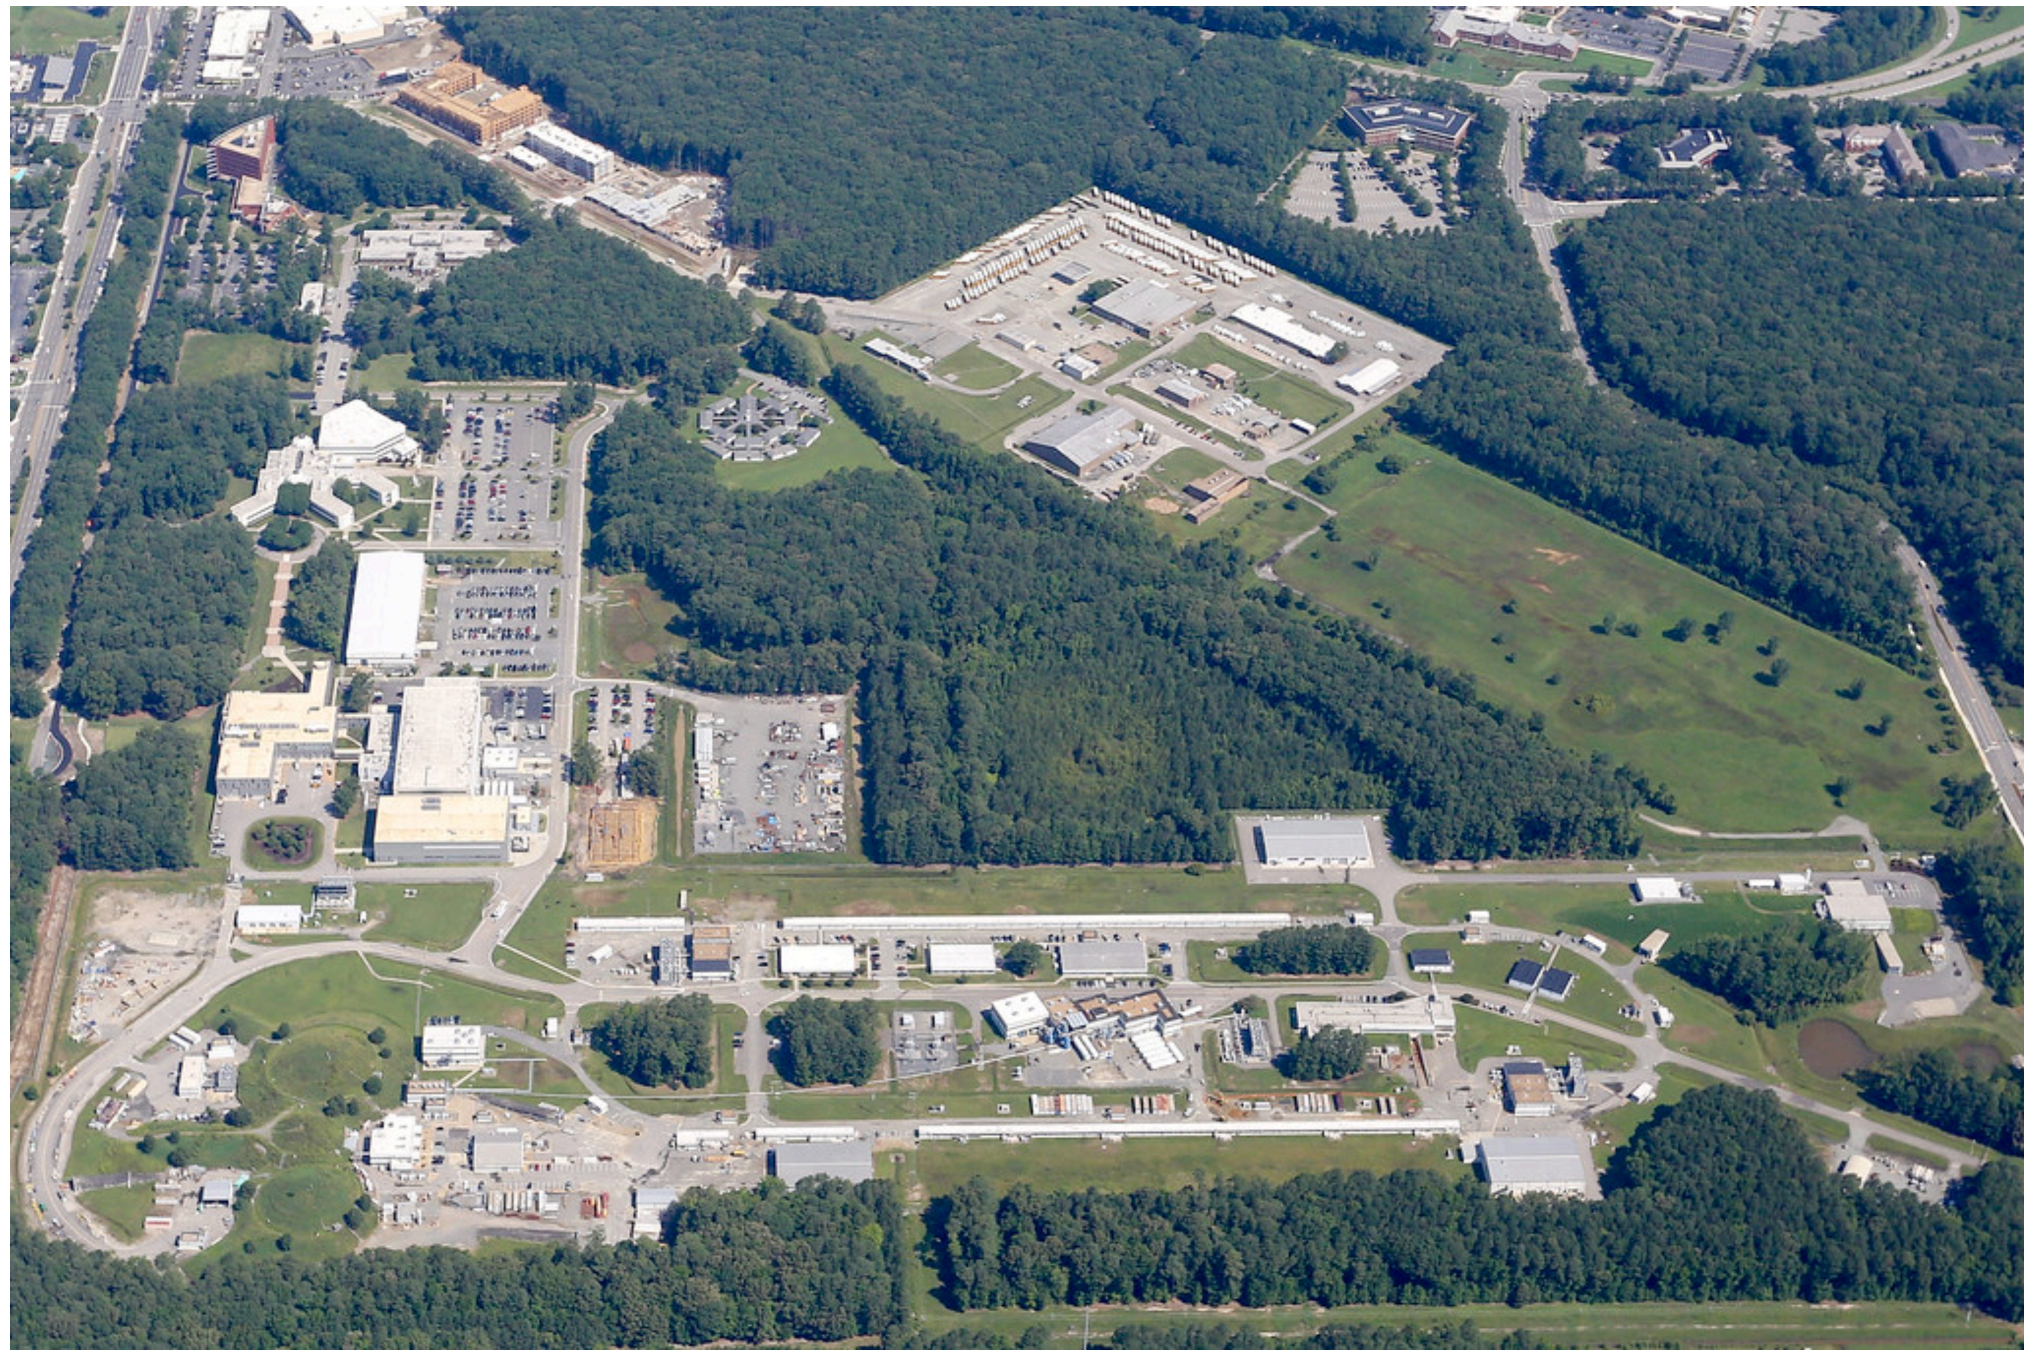

Experiments at JLab

- The CEBAF accelerator delivers high intensity electrons, up to an energy of 12 GeV, to four end stations, called *halls*.
- The halls contain detectors that are used, with differing targets and beam parameters to study properties and structure of atomic nuclei.
- The detectors detect particles from interactions of electrons from the beam with a nucleus in the target.
- Three basic types of data :
  - Charges - measured by ADCs.
  - Times - measured by TDCs.
  - Counts - measured by Scalars.
- The basic data unit is the “event”.
- Events are recorded by online systems.
- The bulk processing takes place offline.

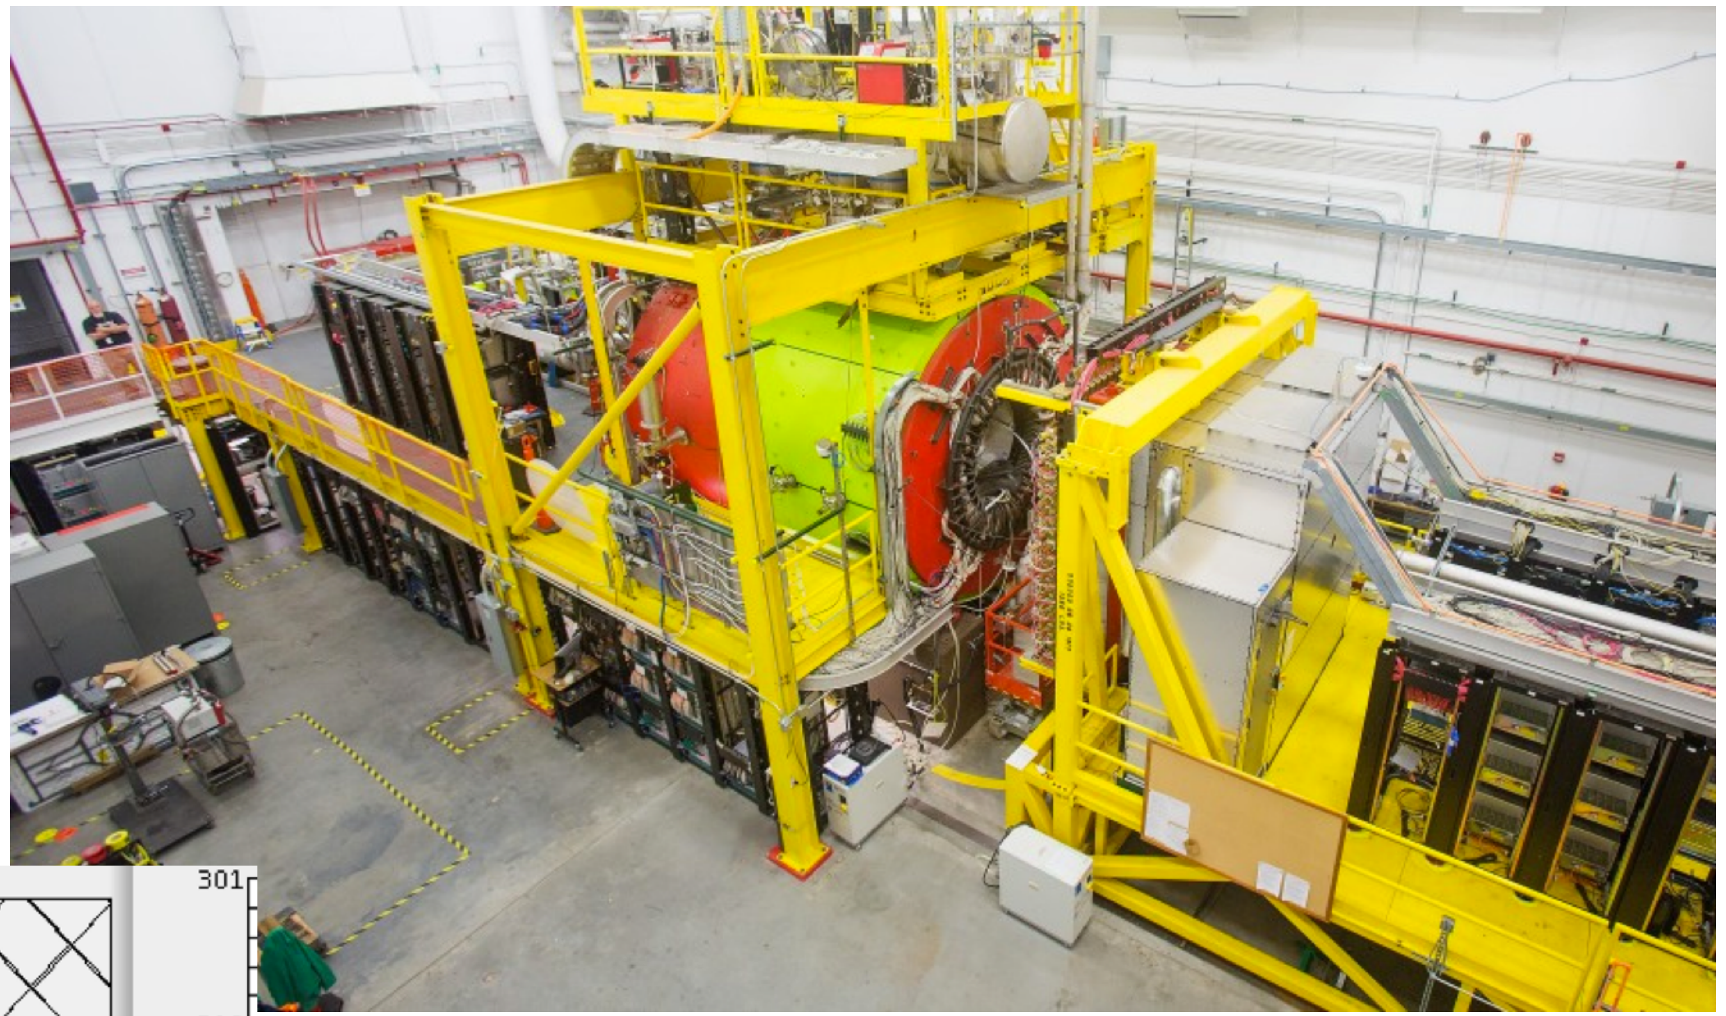

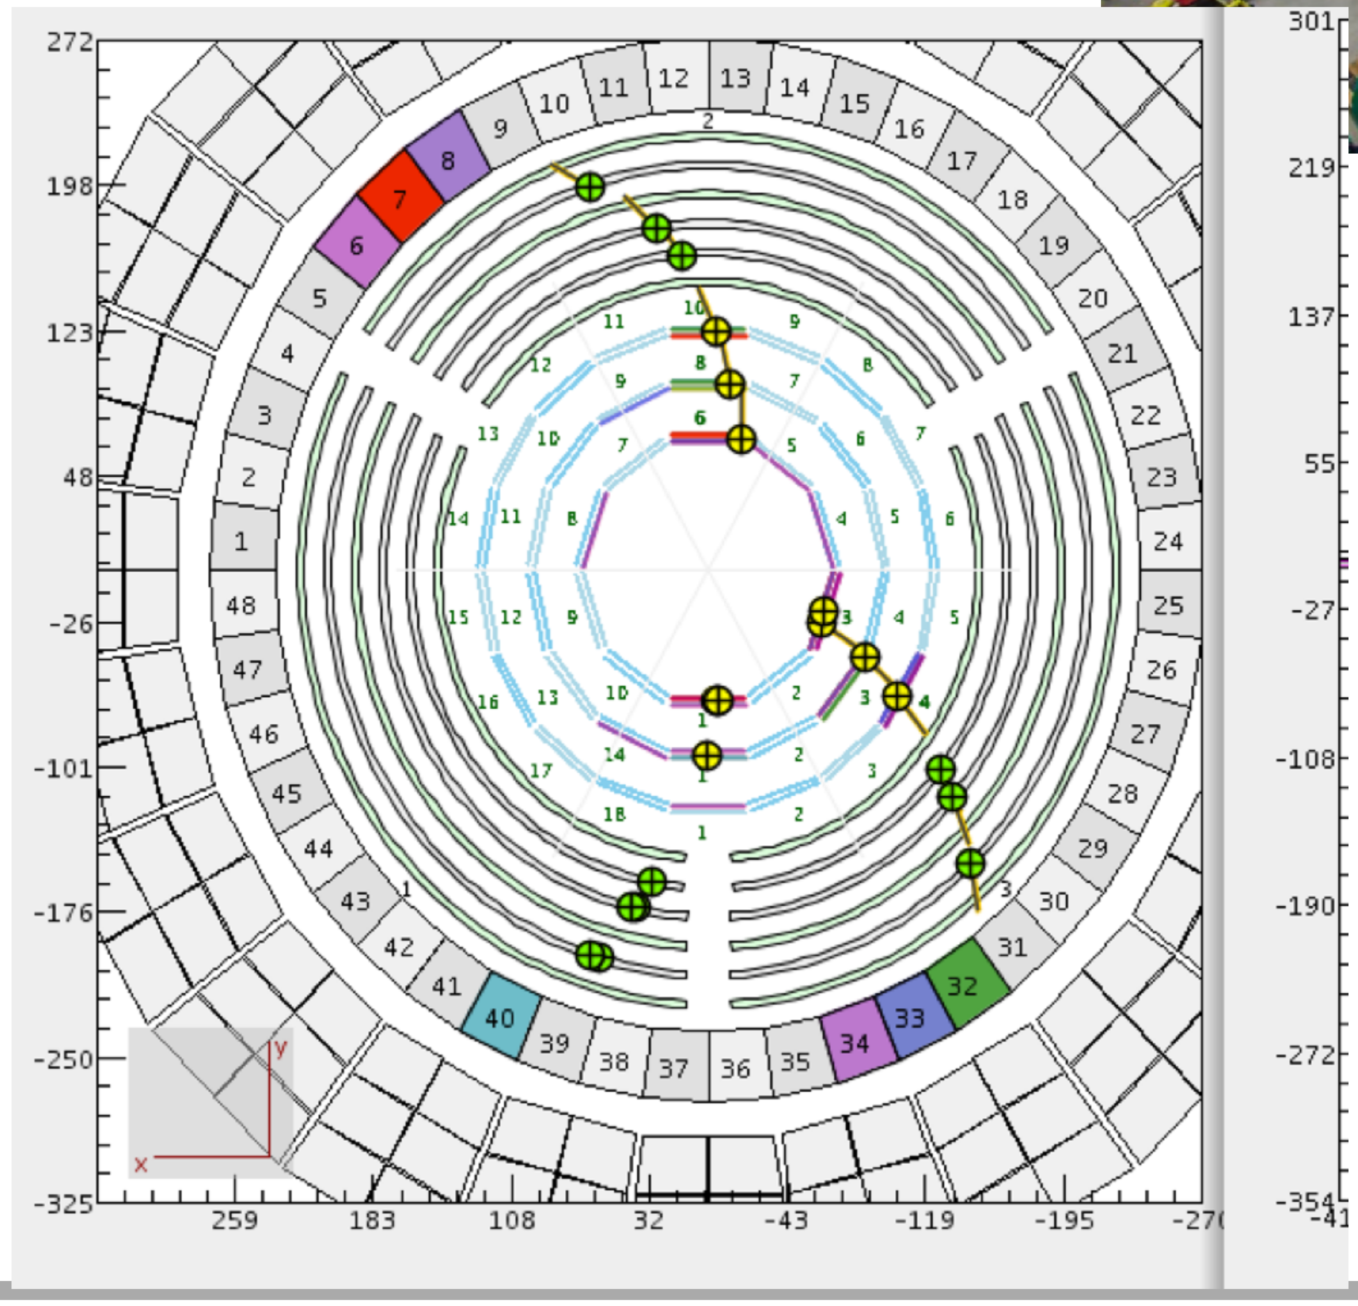

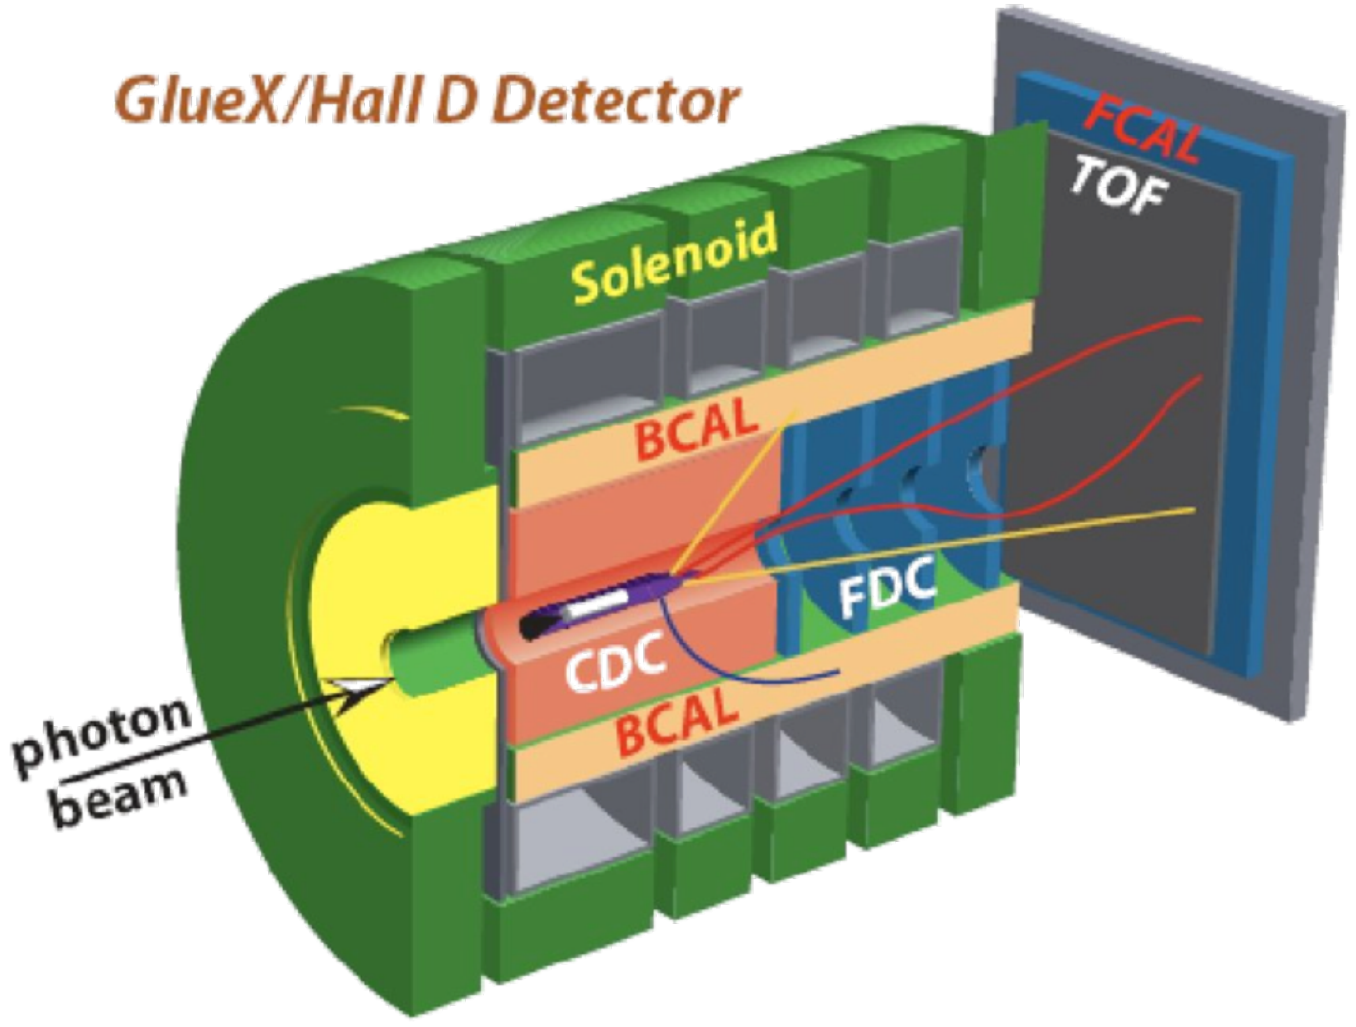

CLAS12 Detector in Hall-B

Forward Detector (FD)

- TORUS magnet
- High Threshold Cherenkov Counter (HTCC)
- Drift chamber system (DC)
- Low Threshold Cherenkov Counter (LTCC)
- RICH detector
- Forward Time-of-Flight System (FTOF)
- Pre-shower calorimeter (PCAL)
- E.M. calorimeter (EC)

Central Detector (CD)

- SOLENOID magnet
- Silicon Vertex Tracker (SVT)
- MicroMegas (MM)
- Central Time-of-Flight (CTOF)
- Central Neutron Detector (CND)

Others

- Forward Tagger (FT)
- Beamline
- Cryo Target
- Moller polarimeter

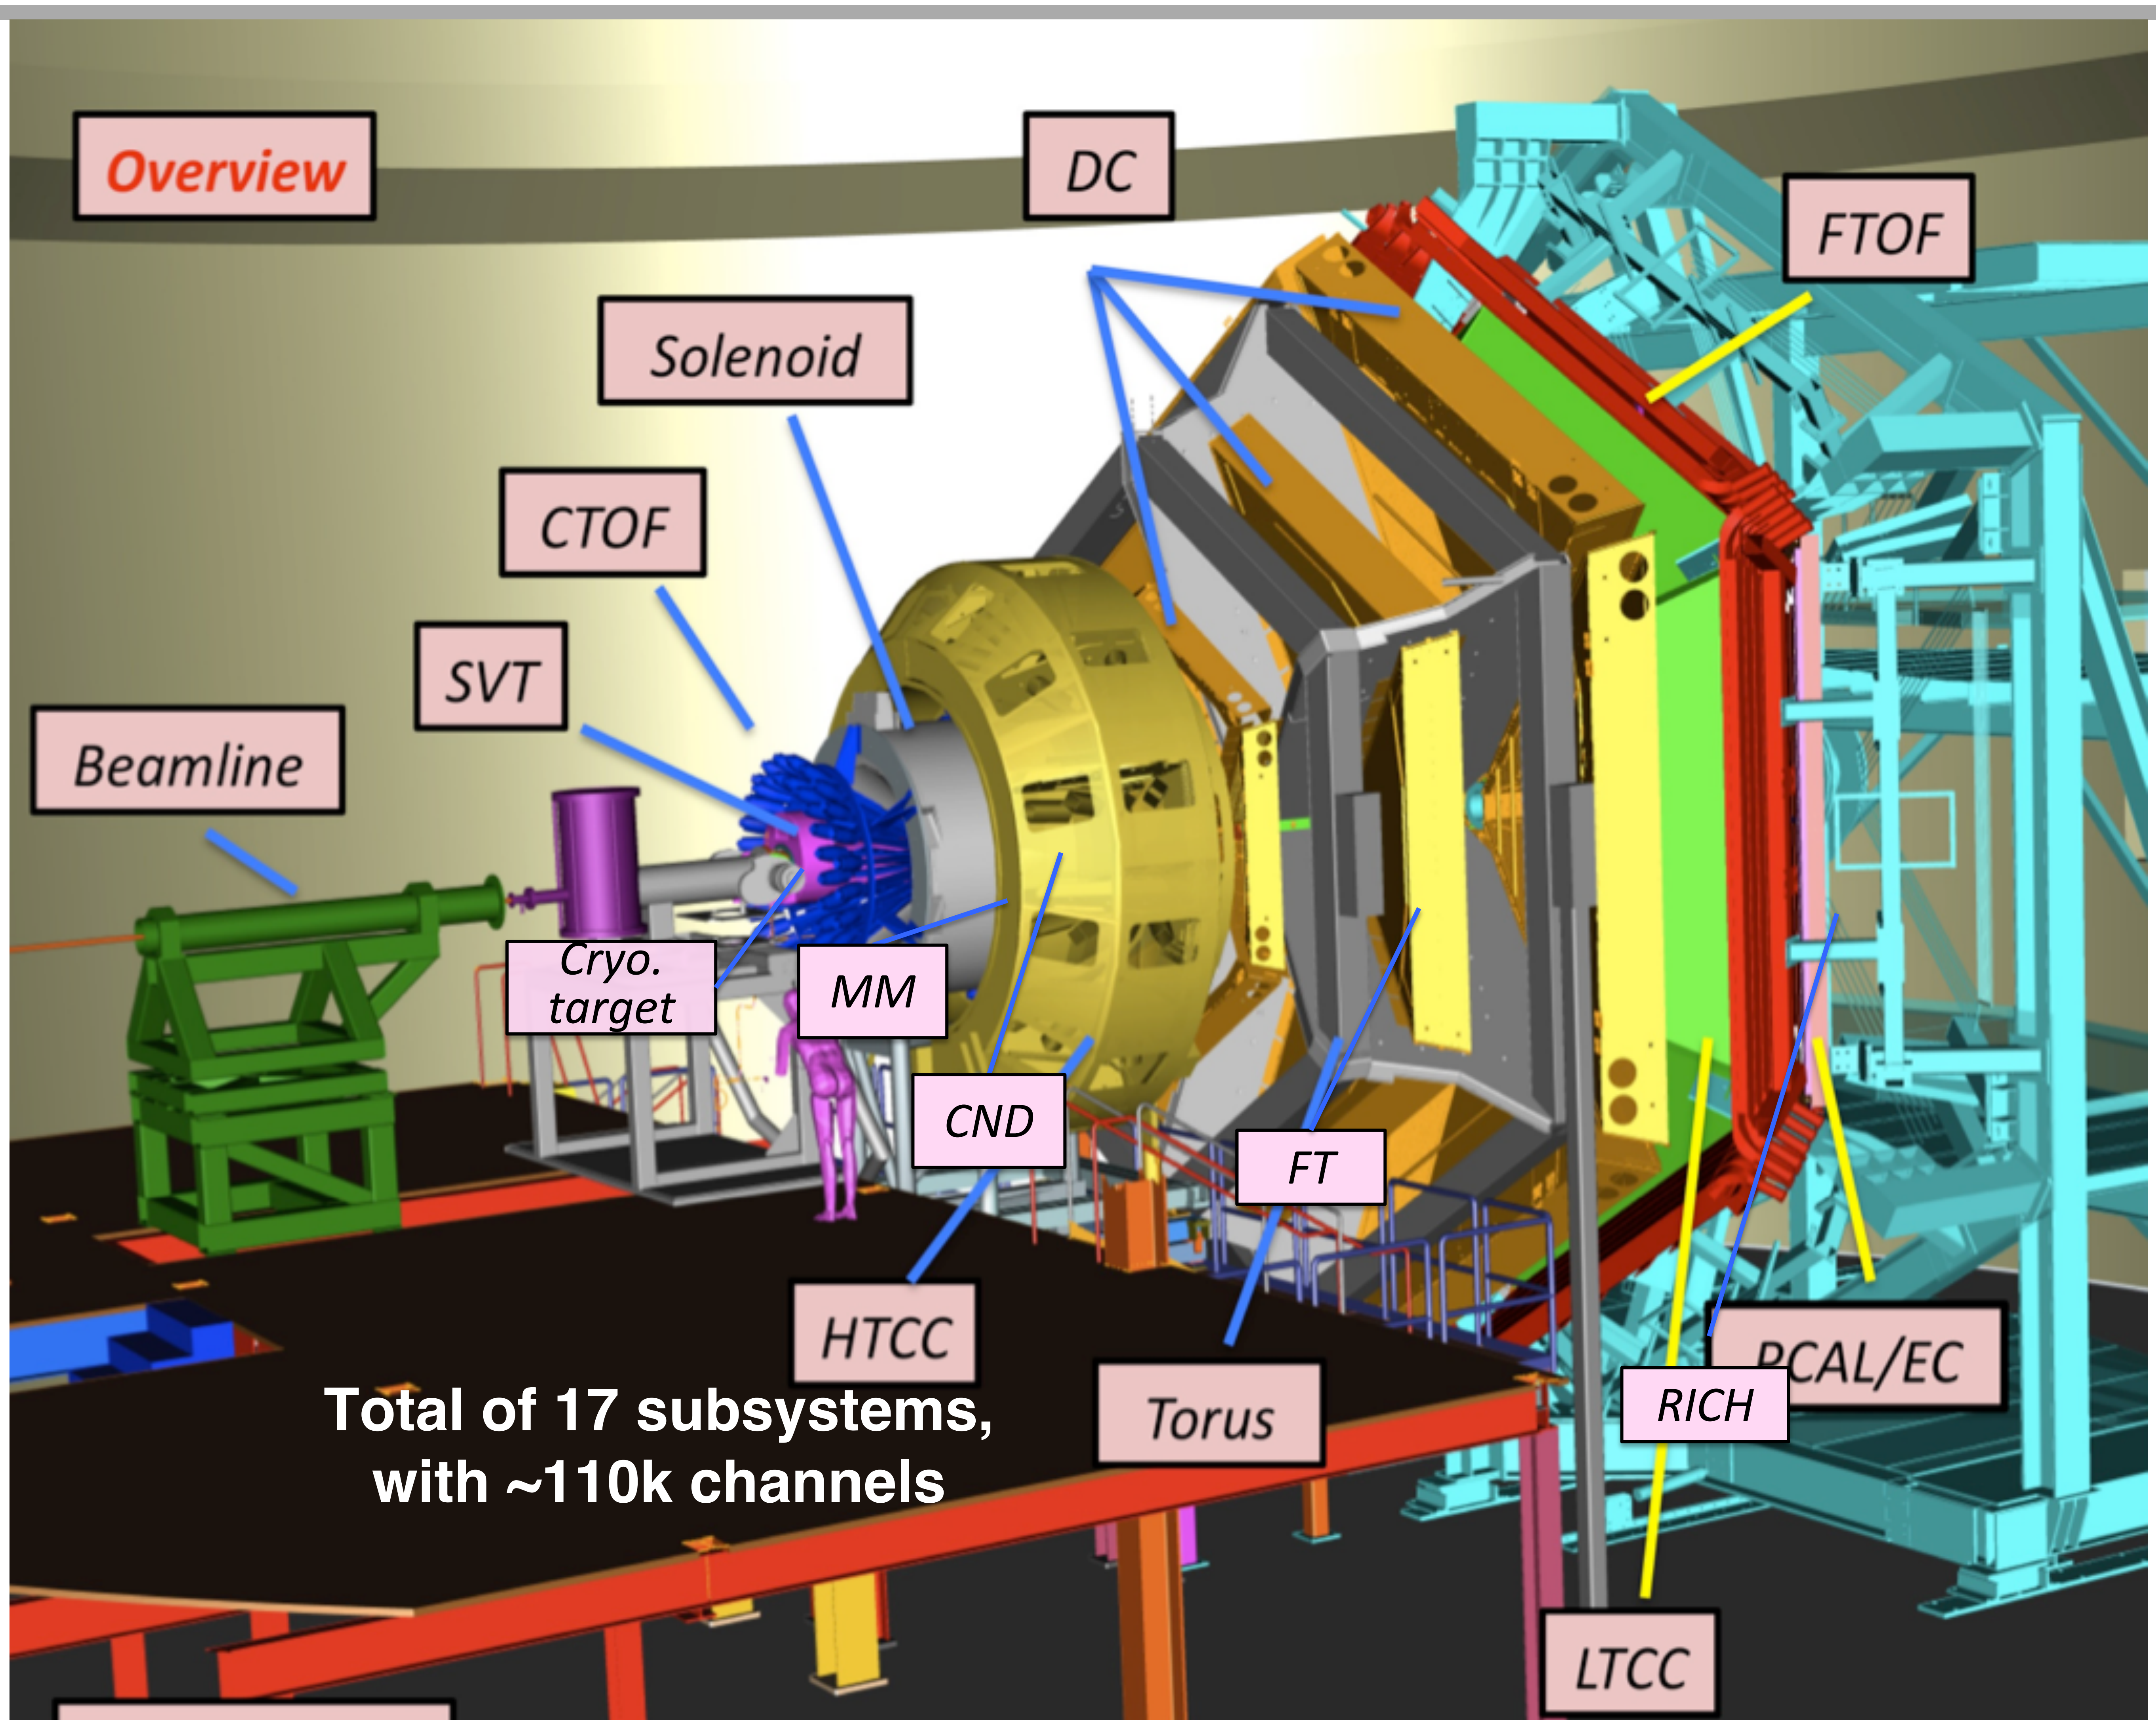

The event as a unit of data

- Data from one event has no history. It doesn't depend upon events that went before and doesn't influence later events.
- Events occur with random timing.
  - Hardware may not be ready for new data.
    - Dead time when data is lost.
  - Events may overlap in time, event pileup.
  - Peak event rate can be more than the average.
- Event size depends upon the physics.
  - Accidental hits unconnected with event.
  - Electronic noise.
  - Distribution of event sizes.
  - Some very large events.

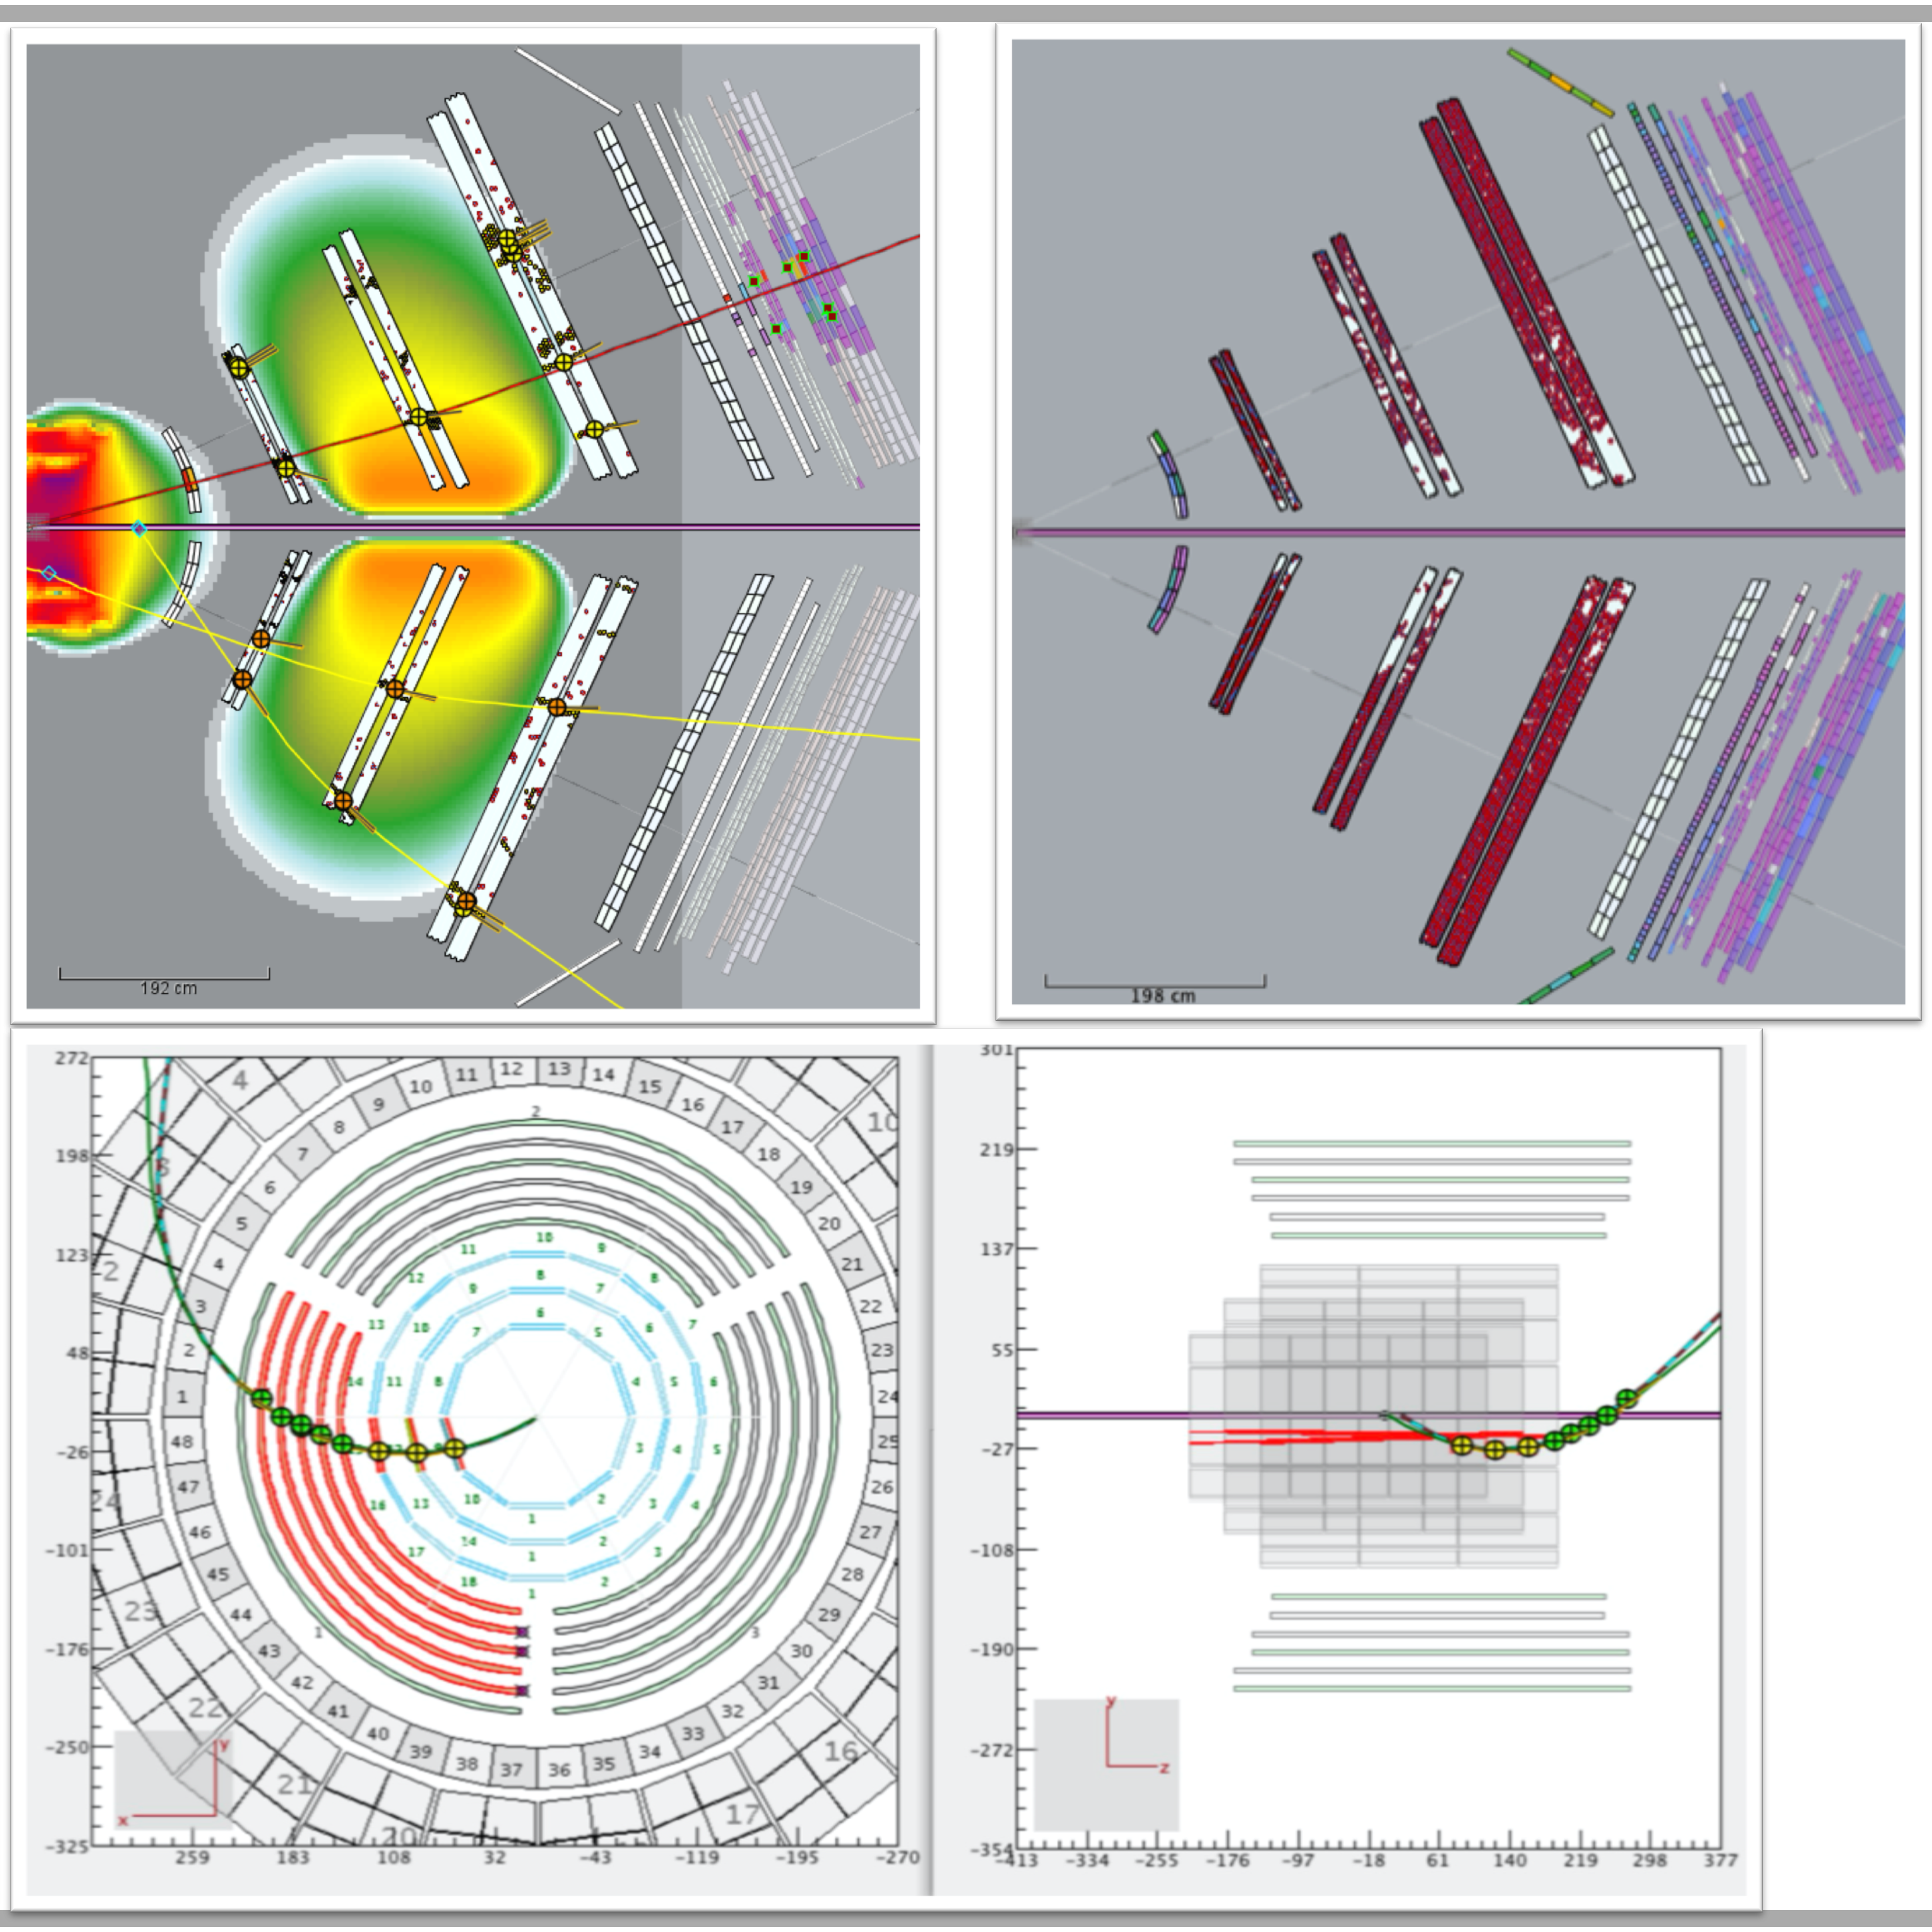

Some numbers

- The halls vary in the scale and longevity of experiments.
  - Halls A and C
    - High turnaround of short-lived individual experiments, weeks or months.
    - Typically high luminosity but simple, and therefore small, events.
    - Data rates < 100 MB/s at 5kB/event.
  - Hall B
    - General purpose detector – group experiments in “run groups” that share a common dataset, typically few months to a year each.
    - Data rates ~300 MB/s (13 kHz event rate, 22 kB/event).
  - Hall D
    - GLUEX represents a single experiment.
    - Low and high intensity phases running for years each.
    - Up to 1.5 GB/s data rate (90 kHz event rate, 17 kB/event) in high intensity phase.

EVIO Data format, a real example

- EVIO is a self describing hierarchical format consisting of nested banks.
  - Bank - container for other data.
- Each bank starts with a header.
  - Length.
  - Description of content.
- The code 0x10 in the outer header tells us this bank contains other banks that are 32-bits wide.
- The first bank is a list of trigger information for all the events in the block. To save space this info is 16-bits.
  - The code 0x20 tells us that this is a bank of 16-bit wide segments (mini banks).
- The following “payload banks” contain blocks of raw data read from ADCs or TDCs.

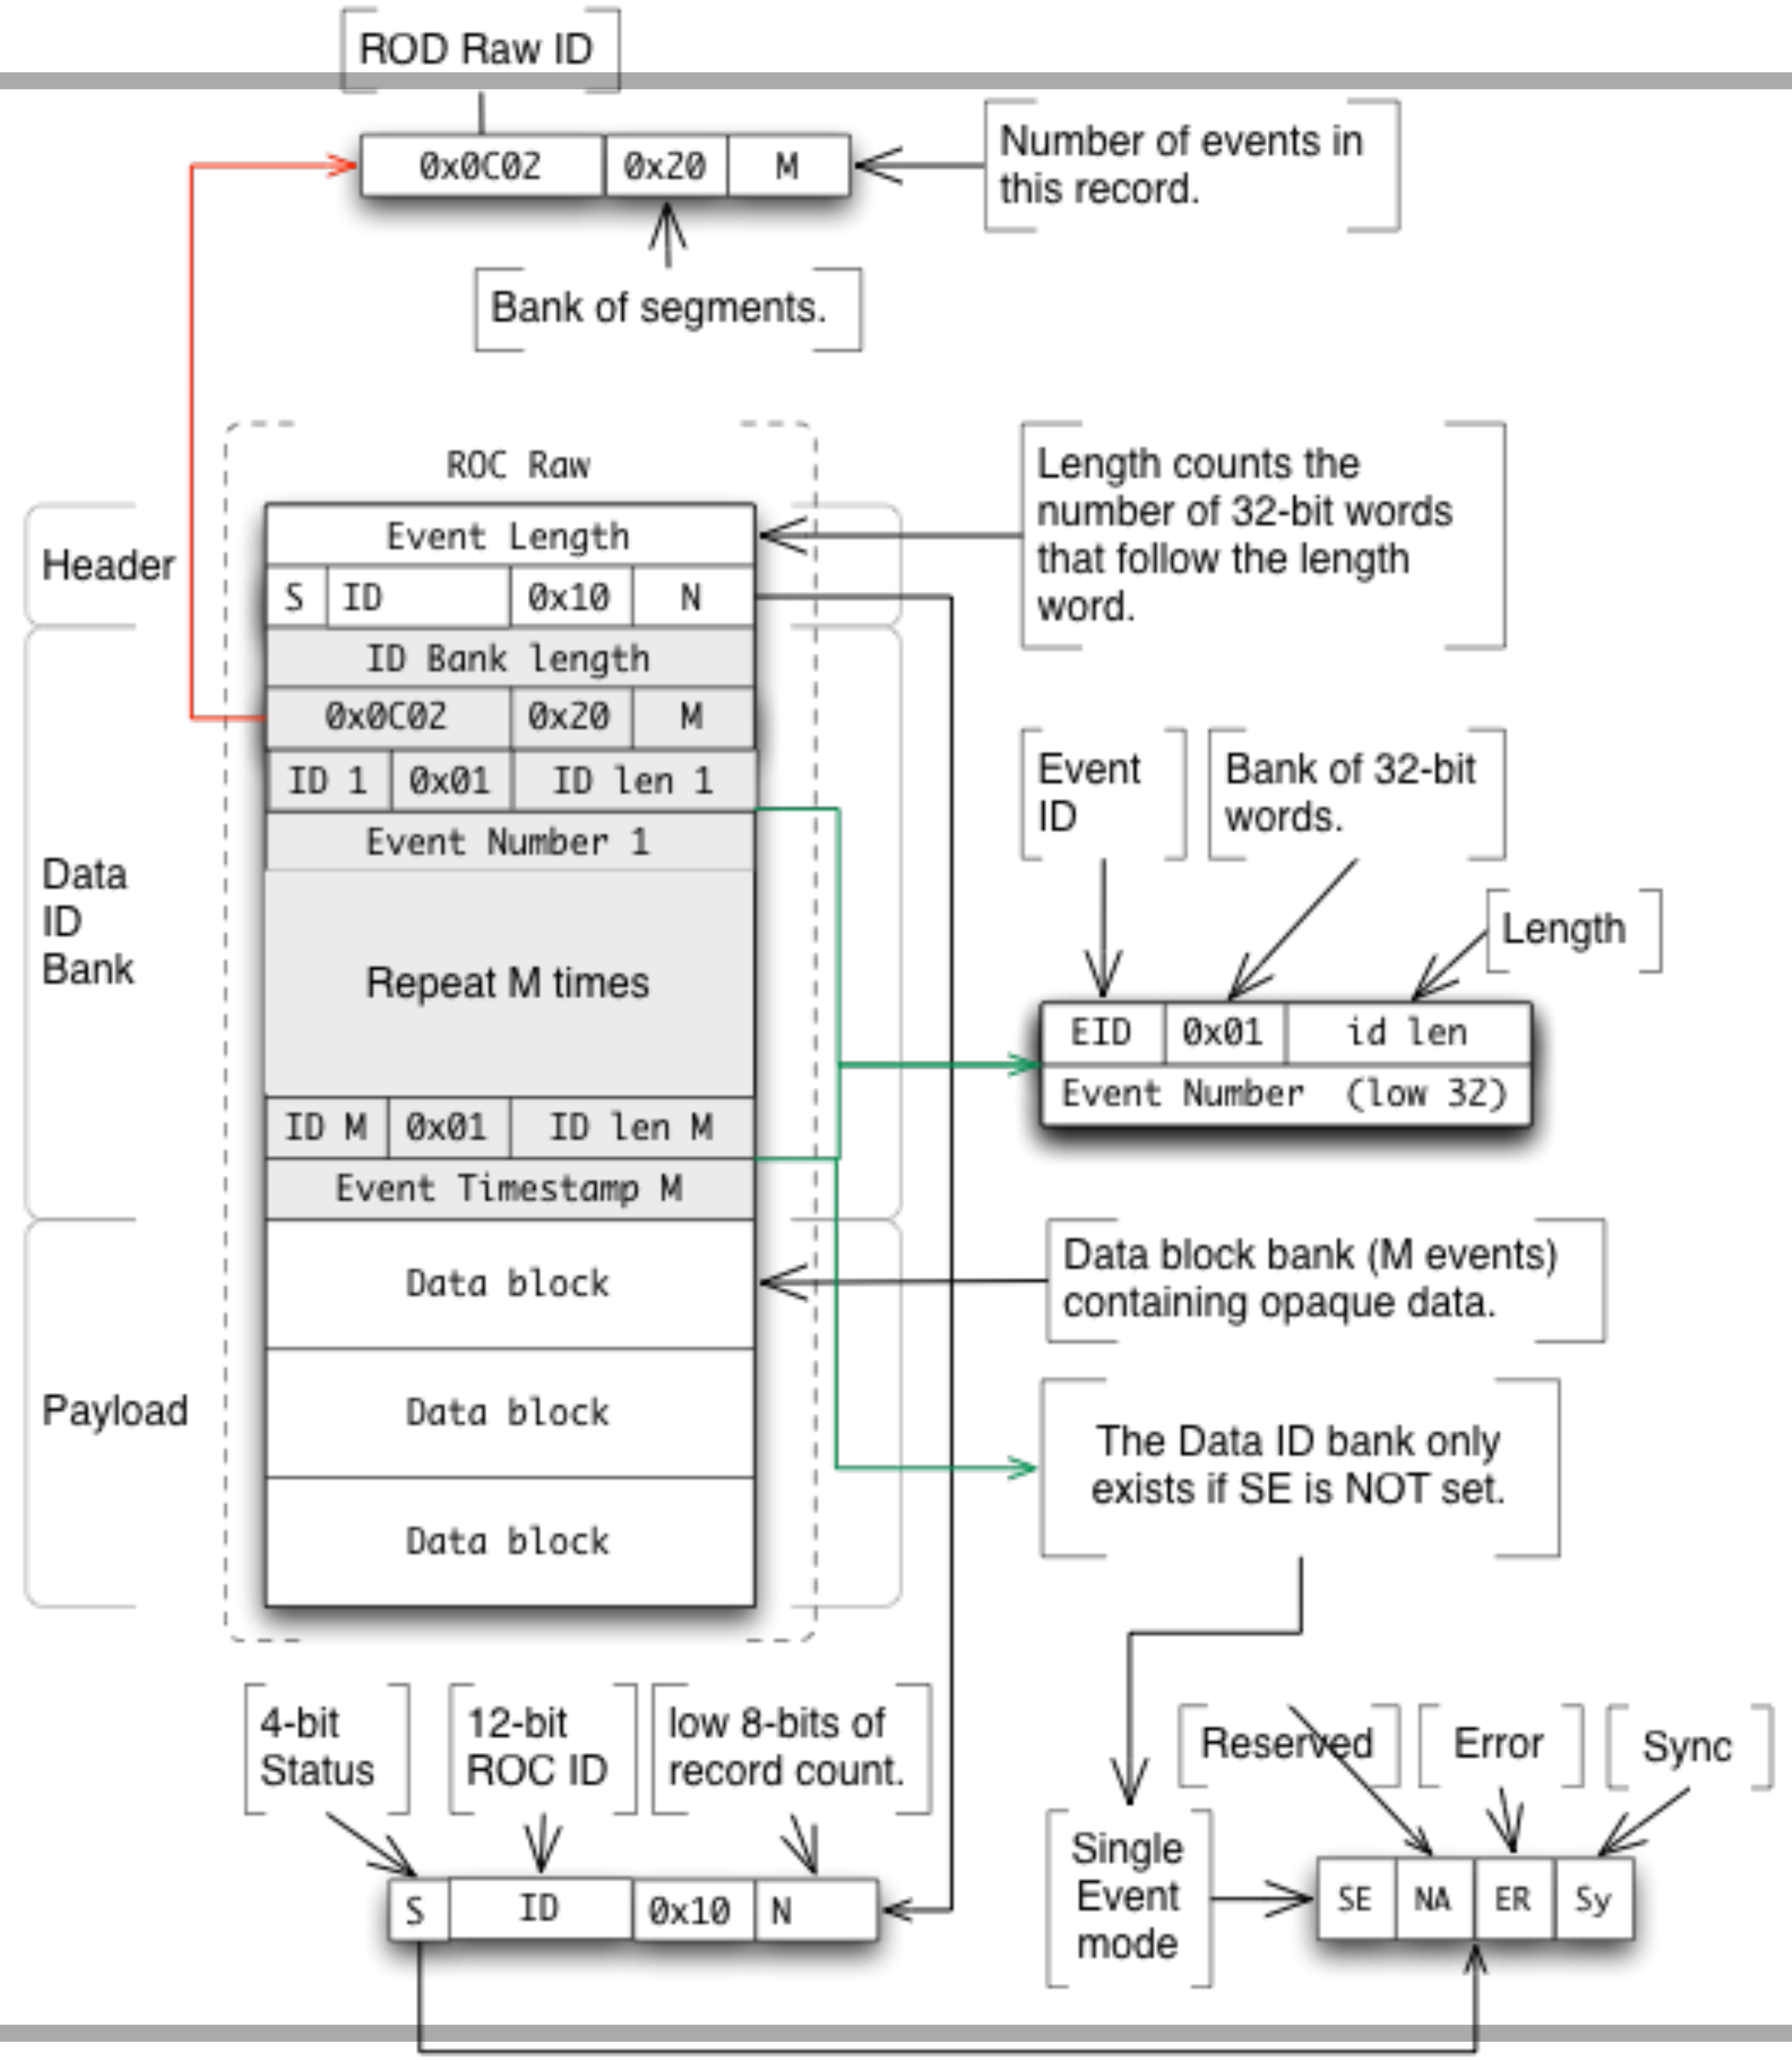

Current data/work flow

- Custom and commercial electronics in VME format digitizes signals from the detectors.
- Most of the signals are noise or uninteresting events.
  - Algorithms implemented in FPGA firmware use prompt data to form a ‘trigger’.
  - Trigger data = blue lines, trigger = red, main data path = black.
  - Analogy: camera shutter - event = photo
- Online Event Builder - Merges event fragments from different detectors.
- Event Recorder writes to files on disk - mass storage system archives to tape.
- Each file contains a linear sequence of events each tagged with an event number.
  - Analogy: event = individual frame on rolls of cine film.
- Groups of files are processed offline by “jobs” on a batch system fronted by a workflow management tool.
  - Onsite cluster, OSG and NERSC in use.
- Algorithms validated using simulated data based on theory.

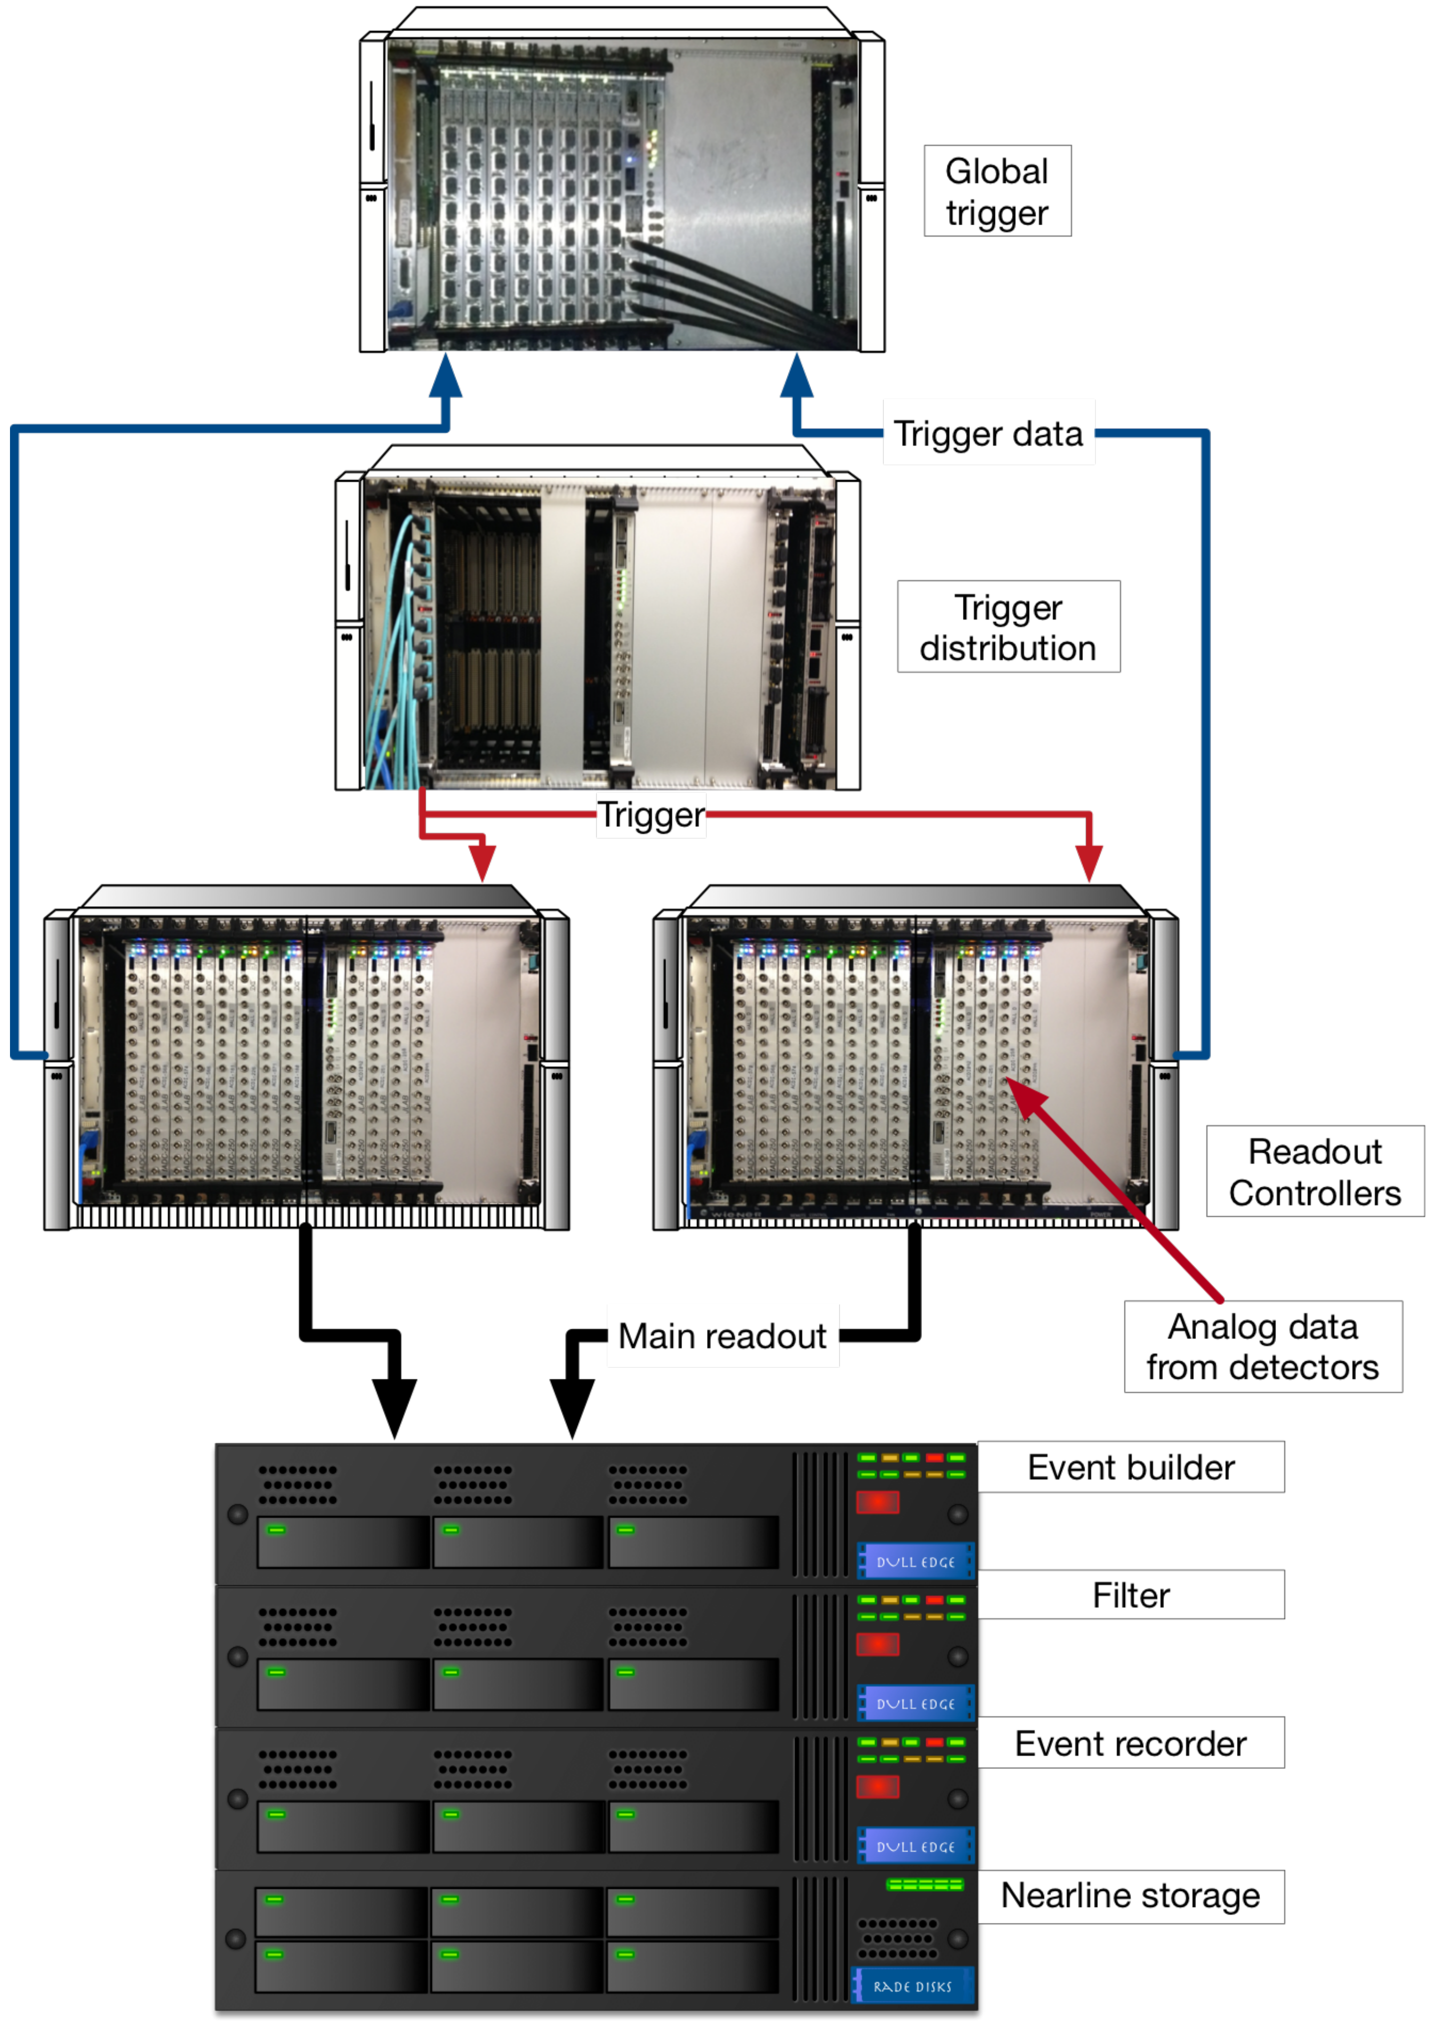

New technology driving change

- Existing DAQ designs at JLab are based on several assumptions:
  - Experiments generate data at a bandwidth that is too high for an affordable system to acquire.
  - Even if the data could be acquired it could not be stored.
  - At these rates the data could not be processed by software in real, or near-real, time.
- In recent years it has become clear to several groups, both at JLab and outside, that these assumptions are no longer true.
- It is now possible to acquire data with minimal filtering in parallel streams to short term storage and process in near real time to reduce it to a volume that can be permanently archived.
- This approach is known as streaming readout.
  - Much of what was formerly done online in custom electronics, firmware and embedded software is moved near/off-line.
  - Looks a lot like the diagram on the previous slide except that the expensive and complex hardware trigger is no longer needed.

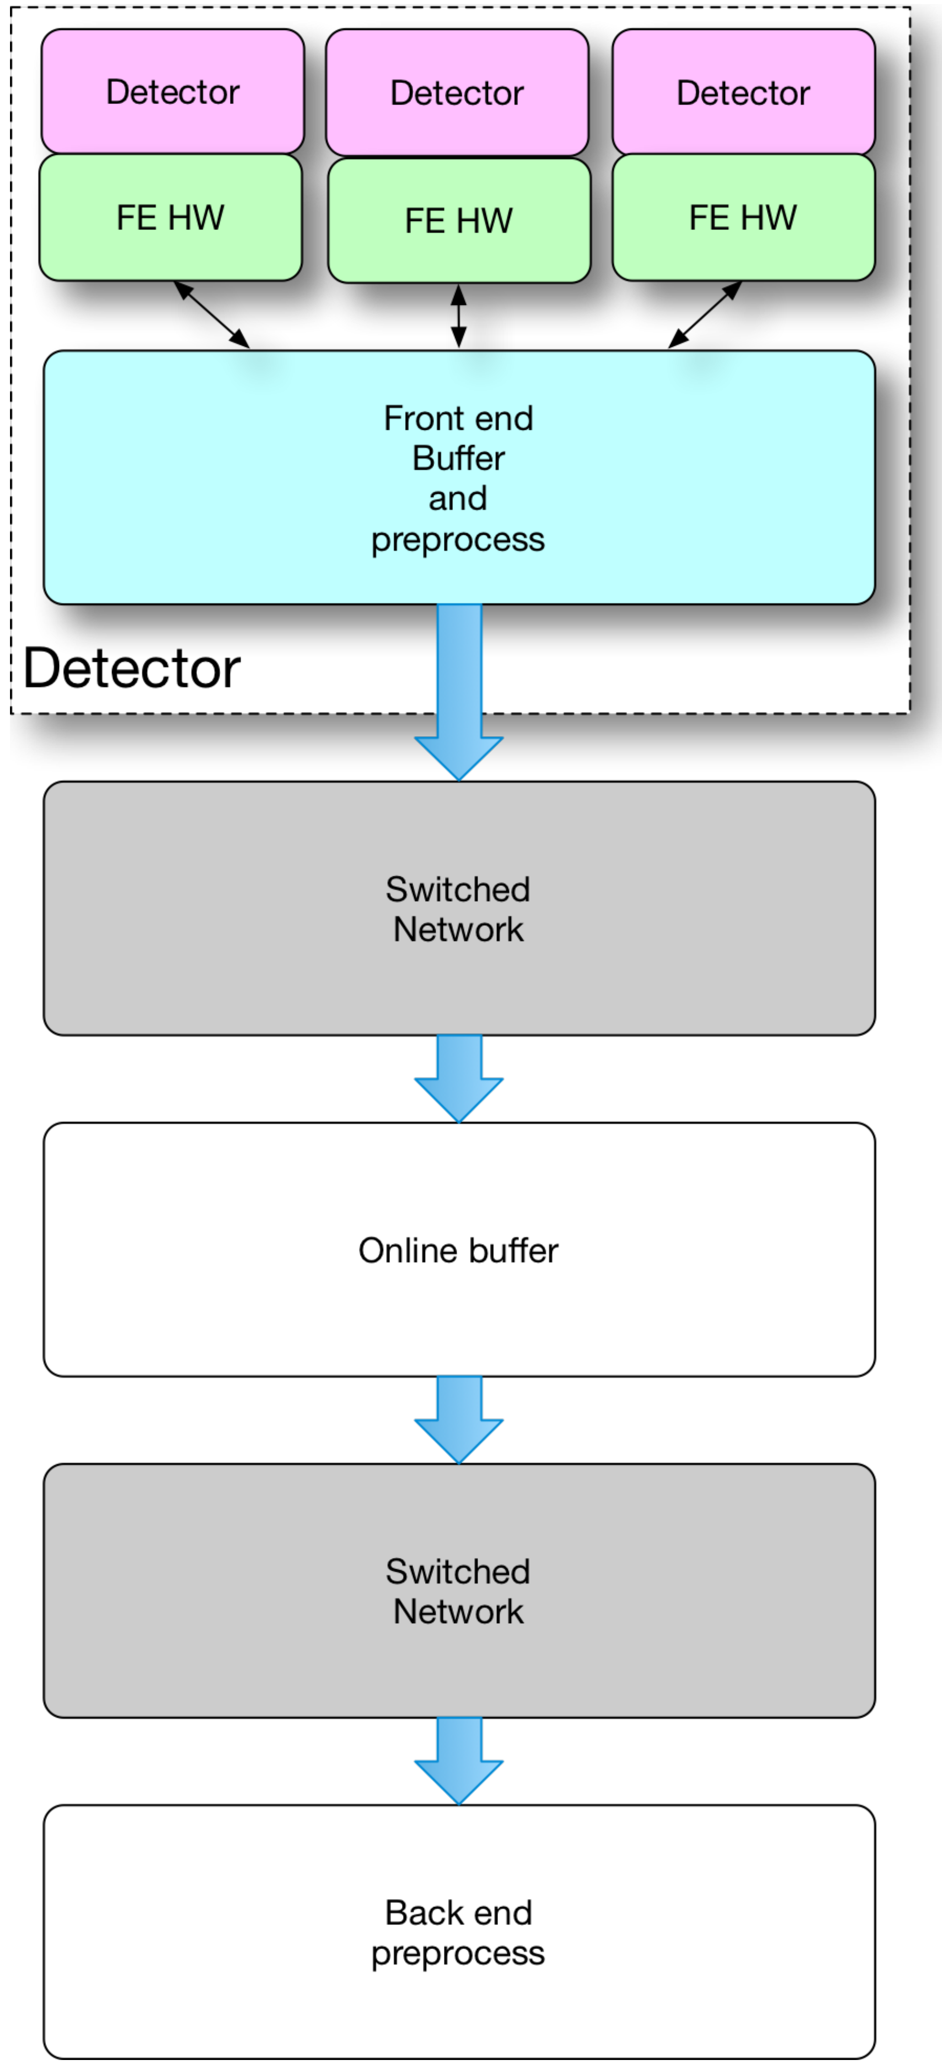

Streaming system

- Electronics close to detector digitises analog signals, preprocesses data and creates stream over fibre.
- Fast online storage holds data while:
  - Virtual trigger defines regions of interest.
  - Event builder defines parts of individual events.
- Slower near-line storage holds data for longer period for slower tasks to access it.
- Data is archived to long term storage for later analysis.
- For this to all work we need:
  - A stream-oriented timing system.
  - Streaming data transport protocol over fiber.
  - Front end electronics that outputs standard format.
  - A stream oriented random-access data storage tier.
  - A framework for data tagging.
  - A framework for near-line tasks.

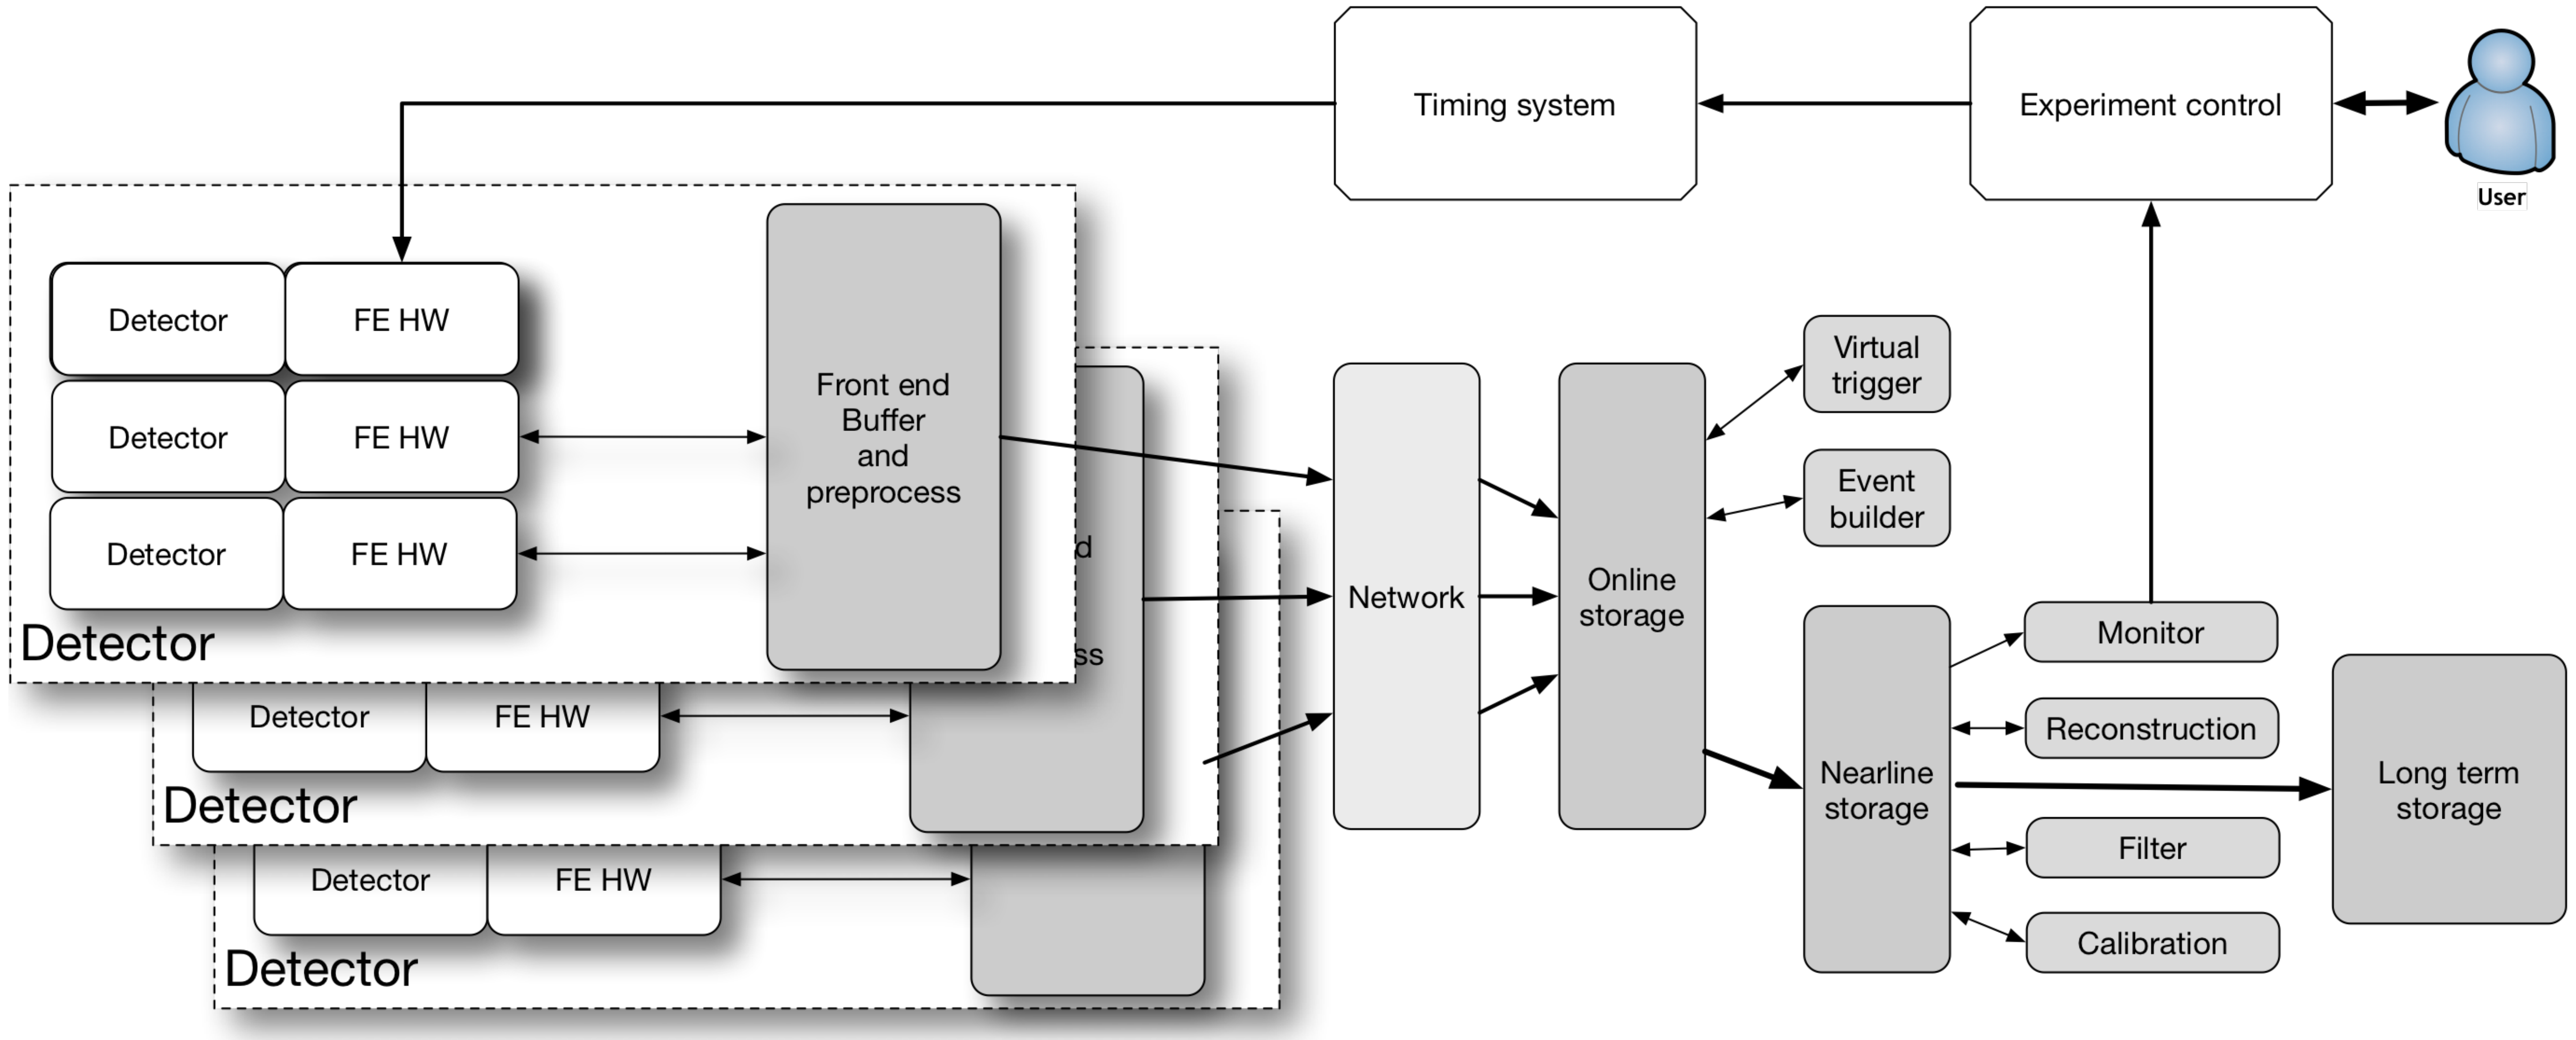

Where are we now- hardware?

- In streaming concept key hardware elements are:
  - A data source outputting on fiber.
  - A front end buffer with FPGA.
  - A high speed low latency network.
  - An online compute resource to buffer and process data.
  - A timing system.
- We have set up an R&D area at JLab and have a test stand using existing hardware.

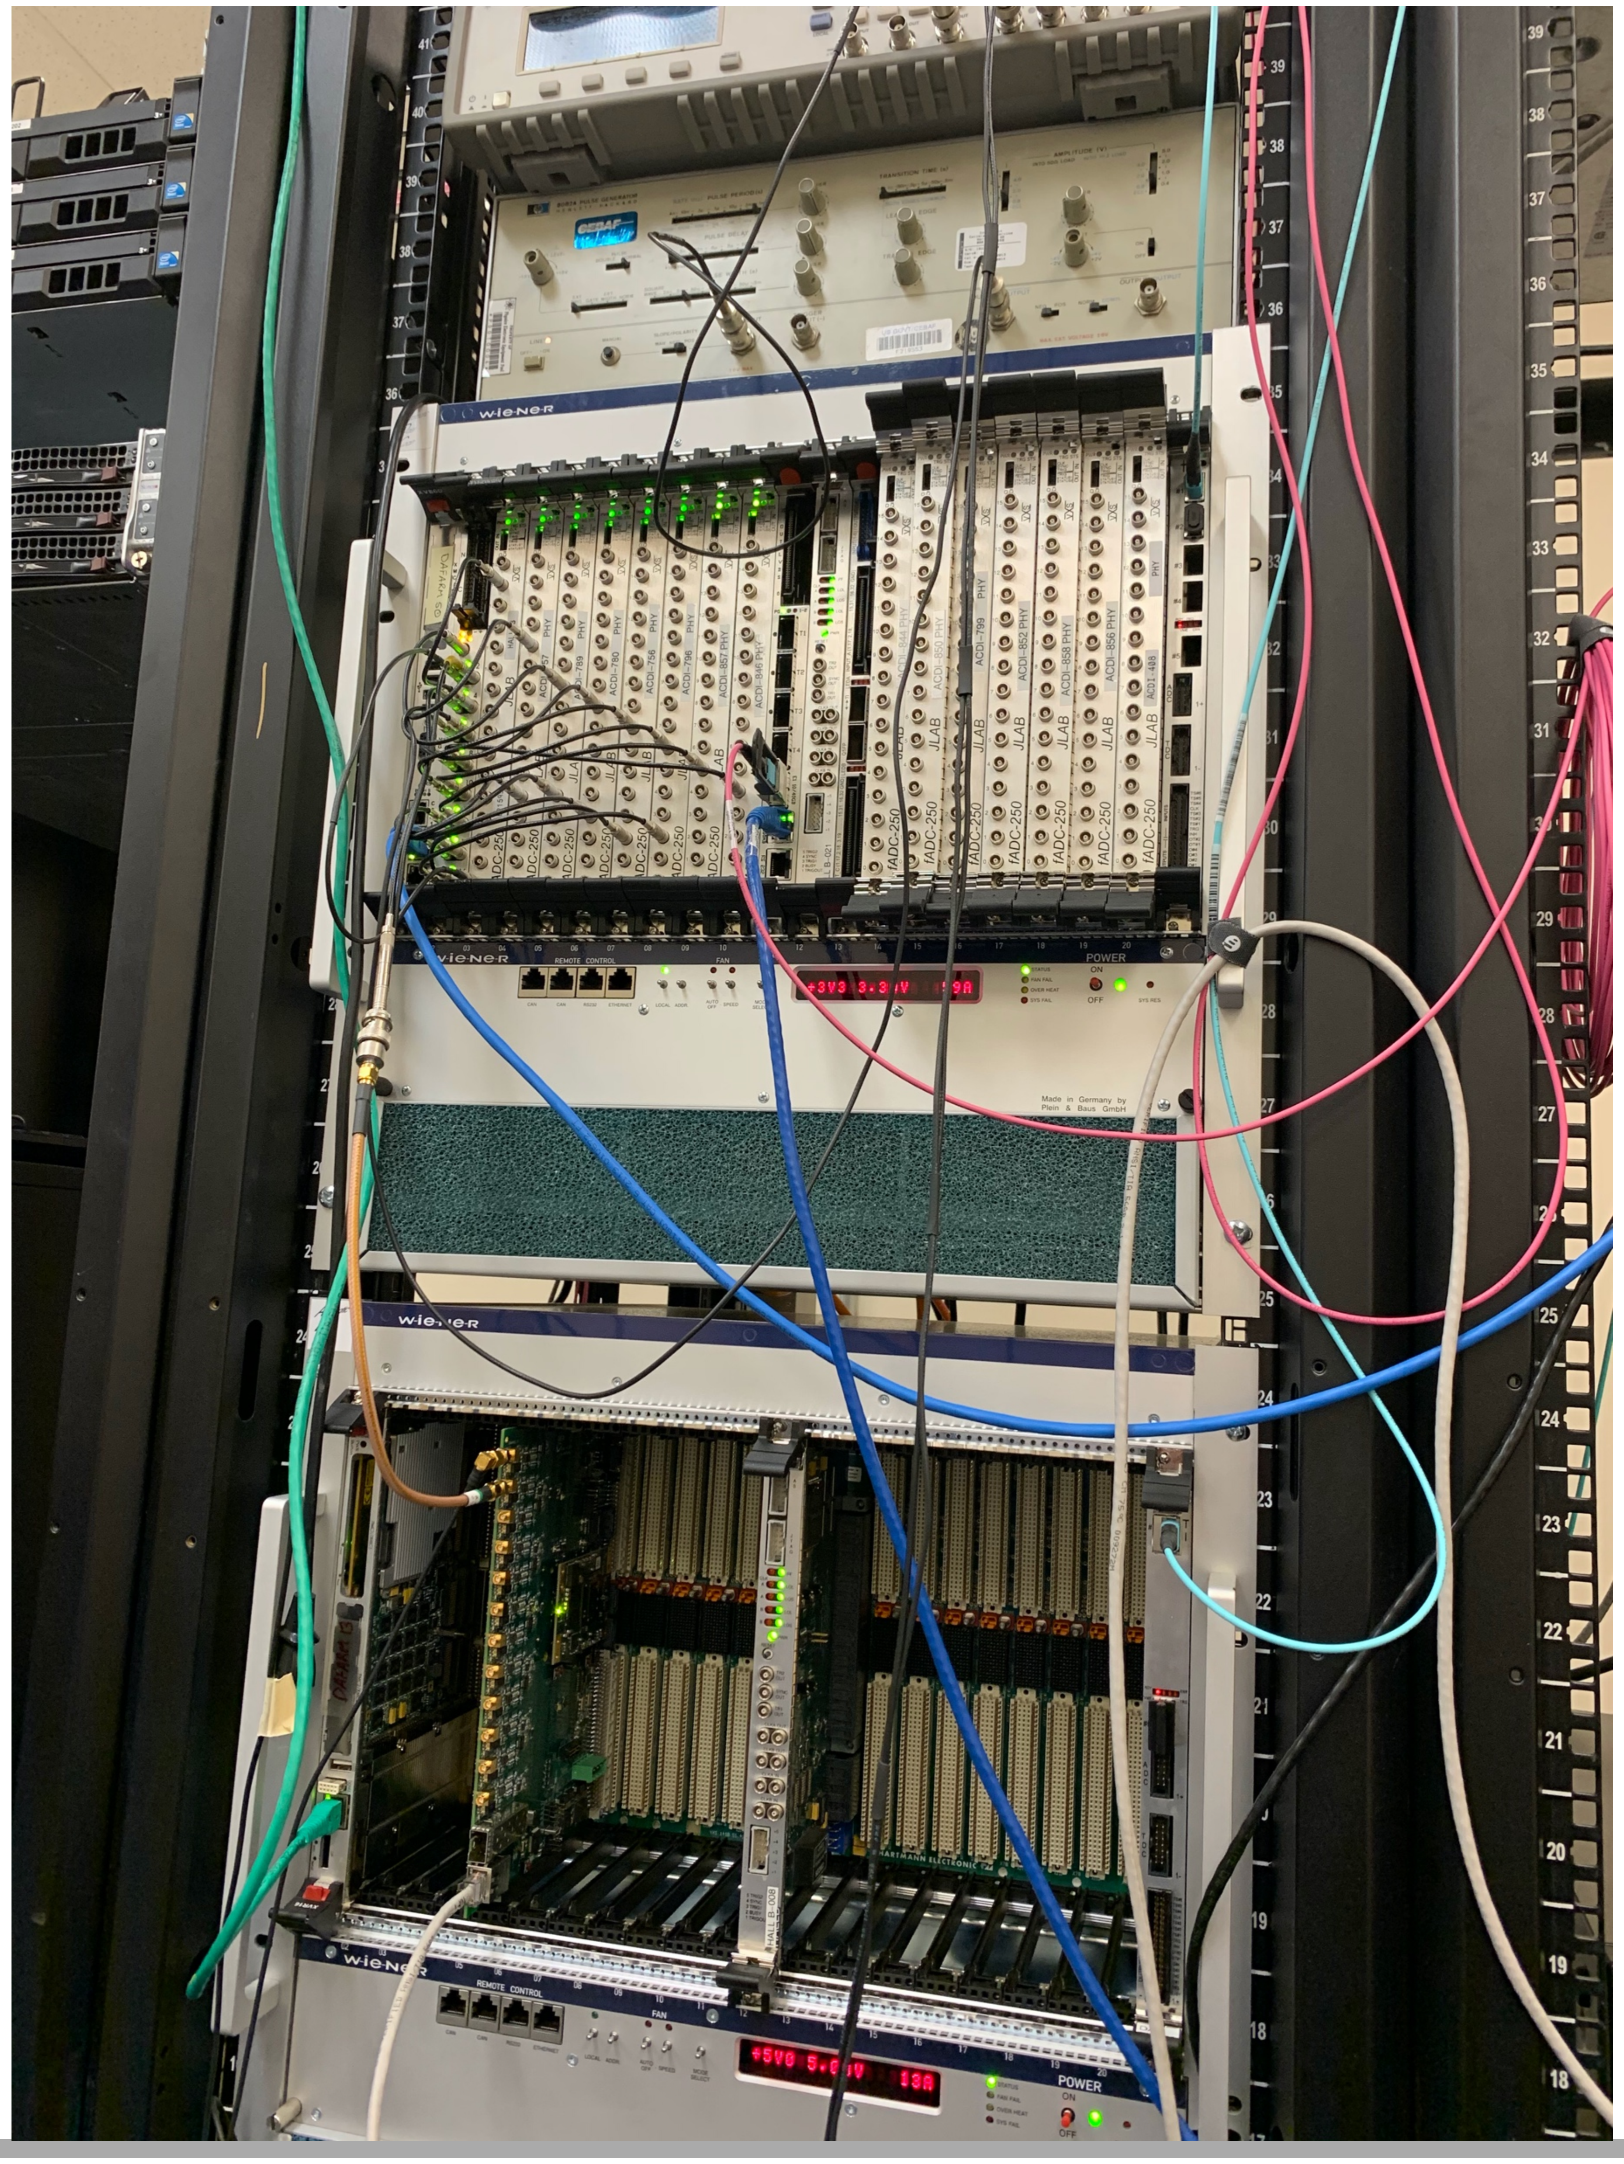

Where are we now- software?

- Two software tracks:
  - JLab developed streaming code
    - Receive data over TCP or UDP, convert to publish/subscribe protocol.
    - Process data using JANA2 C++ Framework.
    - Working in lab test-stand.
  - INFN (Italy)/JLab rapid prototype.
    - INFN has a streaming software system TRIDAS for neutrino experiment.
    - Receive data over TCP or UDP, convert to TRIDAS protocol.
    - Level 2 virtual trigger using TRIDAS.
    - Level 3 virtual trigger+analysis+calibration+monitor using JANA2.
    - We plan a test of detector readout this Spring.
  - JANA2 is a general-purpose framework that can support ML algorithms.