Mu2e Trigger & DAQ Design and Challenges

Tomo Miyashita
Caltech
On Behalf of the Mu2e Collaboration

DPF 2017
July 31st, 2017

Outline

- The Mu2e Experiment
- Mu2e TDAQ Introduction
- General DAQ Strategy
- Hardware Design
- Software Design
- Software Progress
- Next Steps

The Mu2e Experiment I

- Search for charged lepton flavor violation (CLFV) based at Fermilab
  - Looking for the coherent conversion of a μ^- to an e^- in the field of a nucleus:

$$\mu^- + A(Z,N) \to e^- + A(Z,N)$$

  - Allowed in the Standard Model through neutrino oscillations, but extremely suppressed ($\mathcal{B}(\mu^- N \to e^- N) < 10^{-52}$)
  - Therefore, any signal at our sensitivity would be a sign of new physics

Target Sensitivity:

$$R_{\mu e} = \frac{\Gamma\left[\mu^- + A(Z,N) \to e^- + A(Z,N)\right]}{\Gamma\left[\mu^- + A(Z,N) \to \nu_\mu + A(Z-1,N+1)\right]} < 6.7 \times 10^{-17} (90\% \mathrm{CL})$$

4 orders of magnitude better than current limits: **SINDRUM II [W. Bertl et al., Eur. Phys. J. C 47, 337-346 (2006)]**

- Need $\sim 10^{18}$ stopped muons over 3 year run
  - => $\sim 10^{10}$ stopped muons per second
- Signal is mono-energetic electron at 104.9 MeV
  - Main background is Decay In Orbit (DIO) events

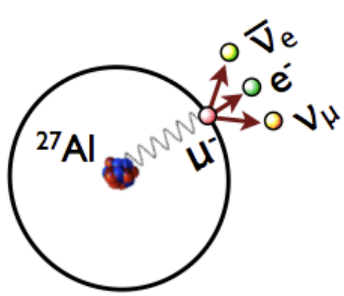

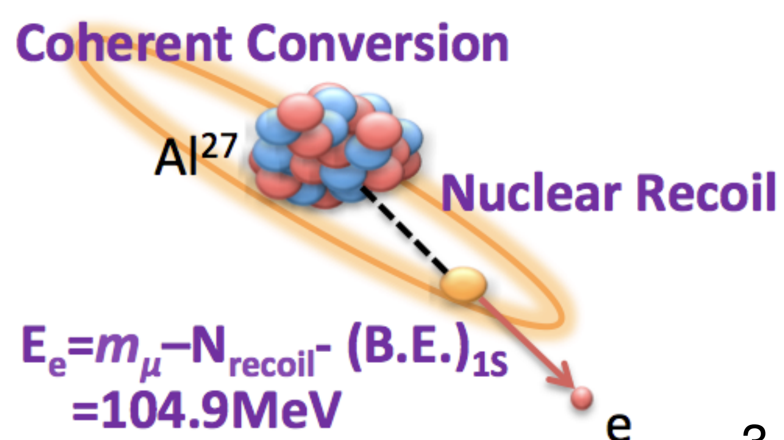

The Mu2e Experiment II

- Experiment Layout:

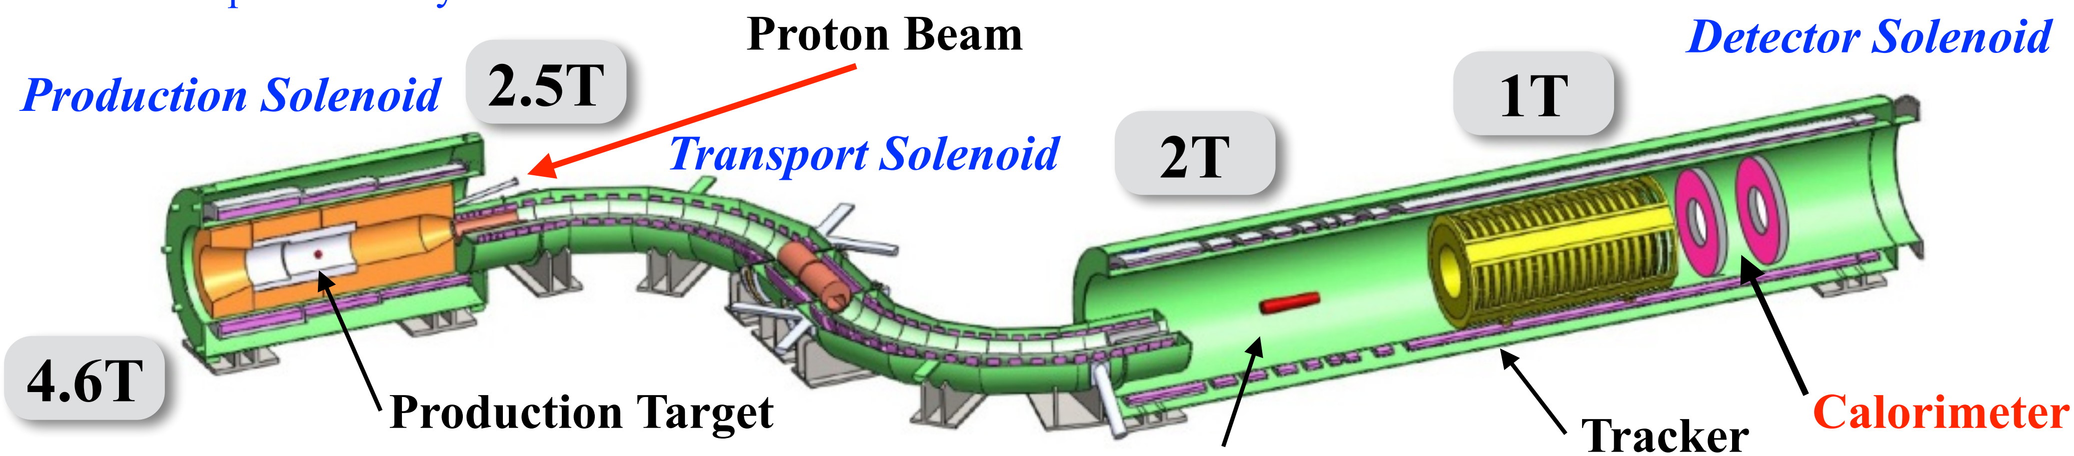

- Production Target + Production Solenoid
  - High intensity, pulsed, 8 GeV proton beam strikes tungsten production target producing pions
  - Pions are captured by the graded magnetic field and decay to muons
- Transport Solenoid
  - Selects low momentum, negative muons
  - Absorbers and Collimators eliminate high energy negative particles, positive particles, and line-of-sight neutrals
- Stopping Target, Detector, and Detector Solenoid
  - Muons are stopped on an aluminum foil target
  - Tracker measures momentum and trajectories of electrons from muonic atoms
  - Calorimeter measures energy/time
- Cosmic Ray Veto detector surrounds detector solenoid

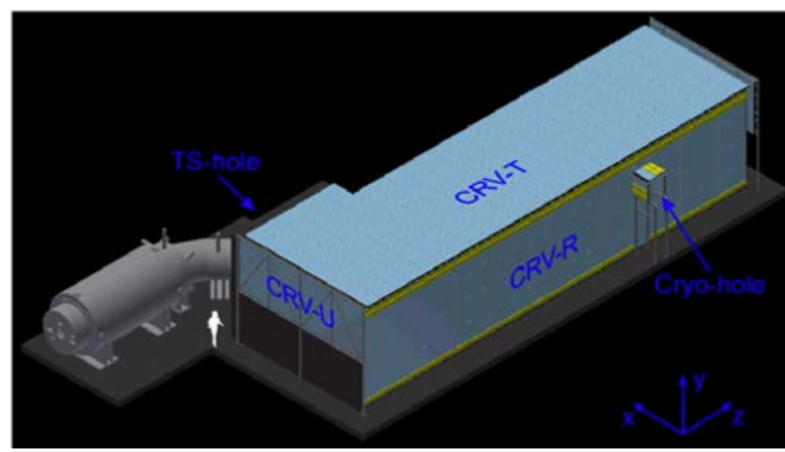

Event Window

- One Cycle of Muon Beamline:

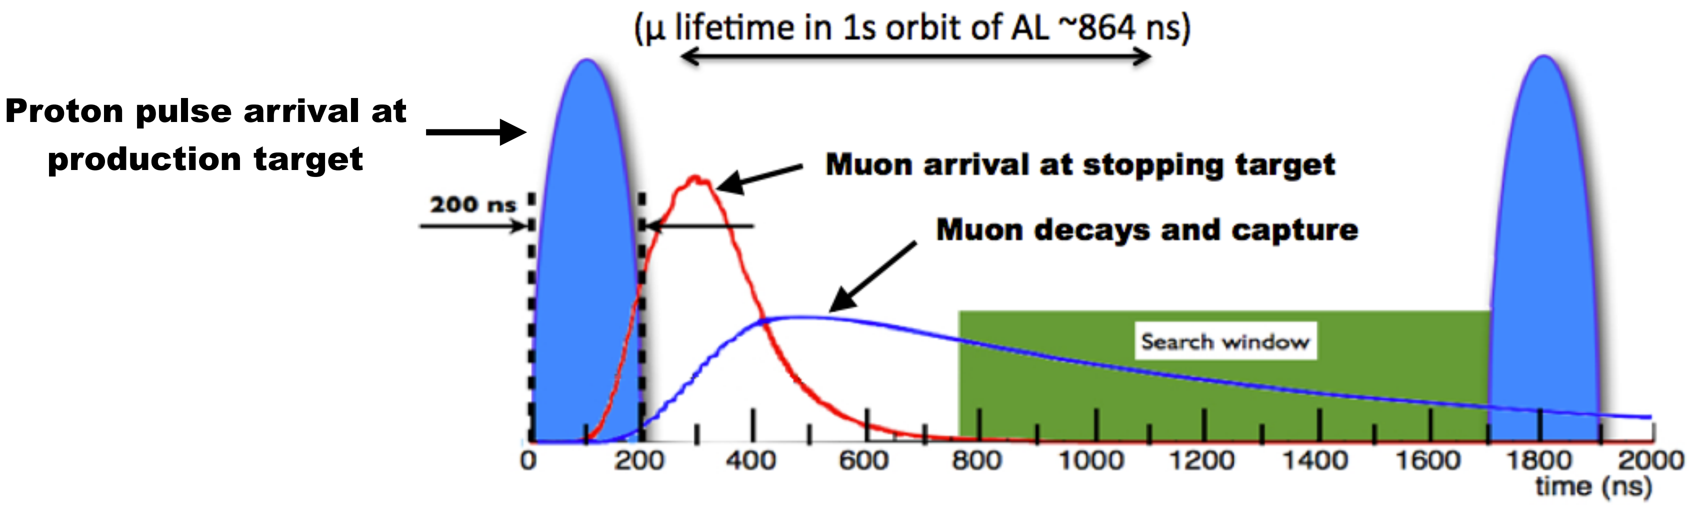

- Muons are accompanied by e^-, e^+, π, anti-protons, etc
  - These create prompt backgrounds
  - Wait for them to decay
- Extinction = (# protons between bunches) / (protons per bunch)
  - We require: Extinction < 10^{-10}
  - Use Extinction Monitor to verify that extinction is at acceptable level
- Stopping Target Monitor is responsible for determining the number of stopped muons
  - Must have an accuracy of at least 10% over the course of the experiment

Introduction I

- The Trigger and Data Acquisition (TDAQ) subsystem is responsible for collecting data from the:
  - Tracker (TRK), Calorimeter (CAL), Cosmic Ray Veto (CRV),
  - Stopping Target Monitor (STM),
  - Extinction Monitor
- Must perform online processing and filtering before delivering data to offline
- Also responsible for generating and distributing timing information to synchronize detector subsystems
- Needs to provide a means of controlling operating modes (data readout, calibration injection, etc.) for each event window
- Responsible for monitoring and operator interfaces

Introduction II

- Online Data Processing must:
  - Provide timing distribution system with recoverable clock jitter <500 ps
  - Provide >35 GB/s input bandwidth
  - Provide processing for 200K events/sec and achieve a reduction ratio of >100:1 with 90% efficiency
  - Provide infrastructure for online data quality monitoring

DAQ Strategy I

- We use a scalable software-based online data processing system and take advantage of commercial hardware
- Traditionally, experiments have developed their own custom event building and filtering software from scratch with minimal sharing of effort between experiments
- While large experiments such as CMS and ATLAS had the resources necessary to do this, it is impractical for smaller experiments such as Mu2e
- To address this problem, the Realtime Software Infrastructure Group at Fermilab has been developing the artdaq toolkit for event-building and filtering at the DAQ level
  - artdaq incorporates functionality from the art toolkit
  - artdaq is used by several Fermilab experiments

DAQ Strategy II

- Multiple benefits from implementing the data acquisition framework using artdaq
  - Allows use of commercial computers for online processing
  - Efficient use of multicore computers
  - Simplifies parallelism
  - Easily scalable
  - By using the art framework, artdaq makes interfacing with offline software easier and facilitates sharing of code between online and offline development

Hardware Architecture I

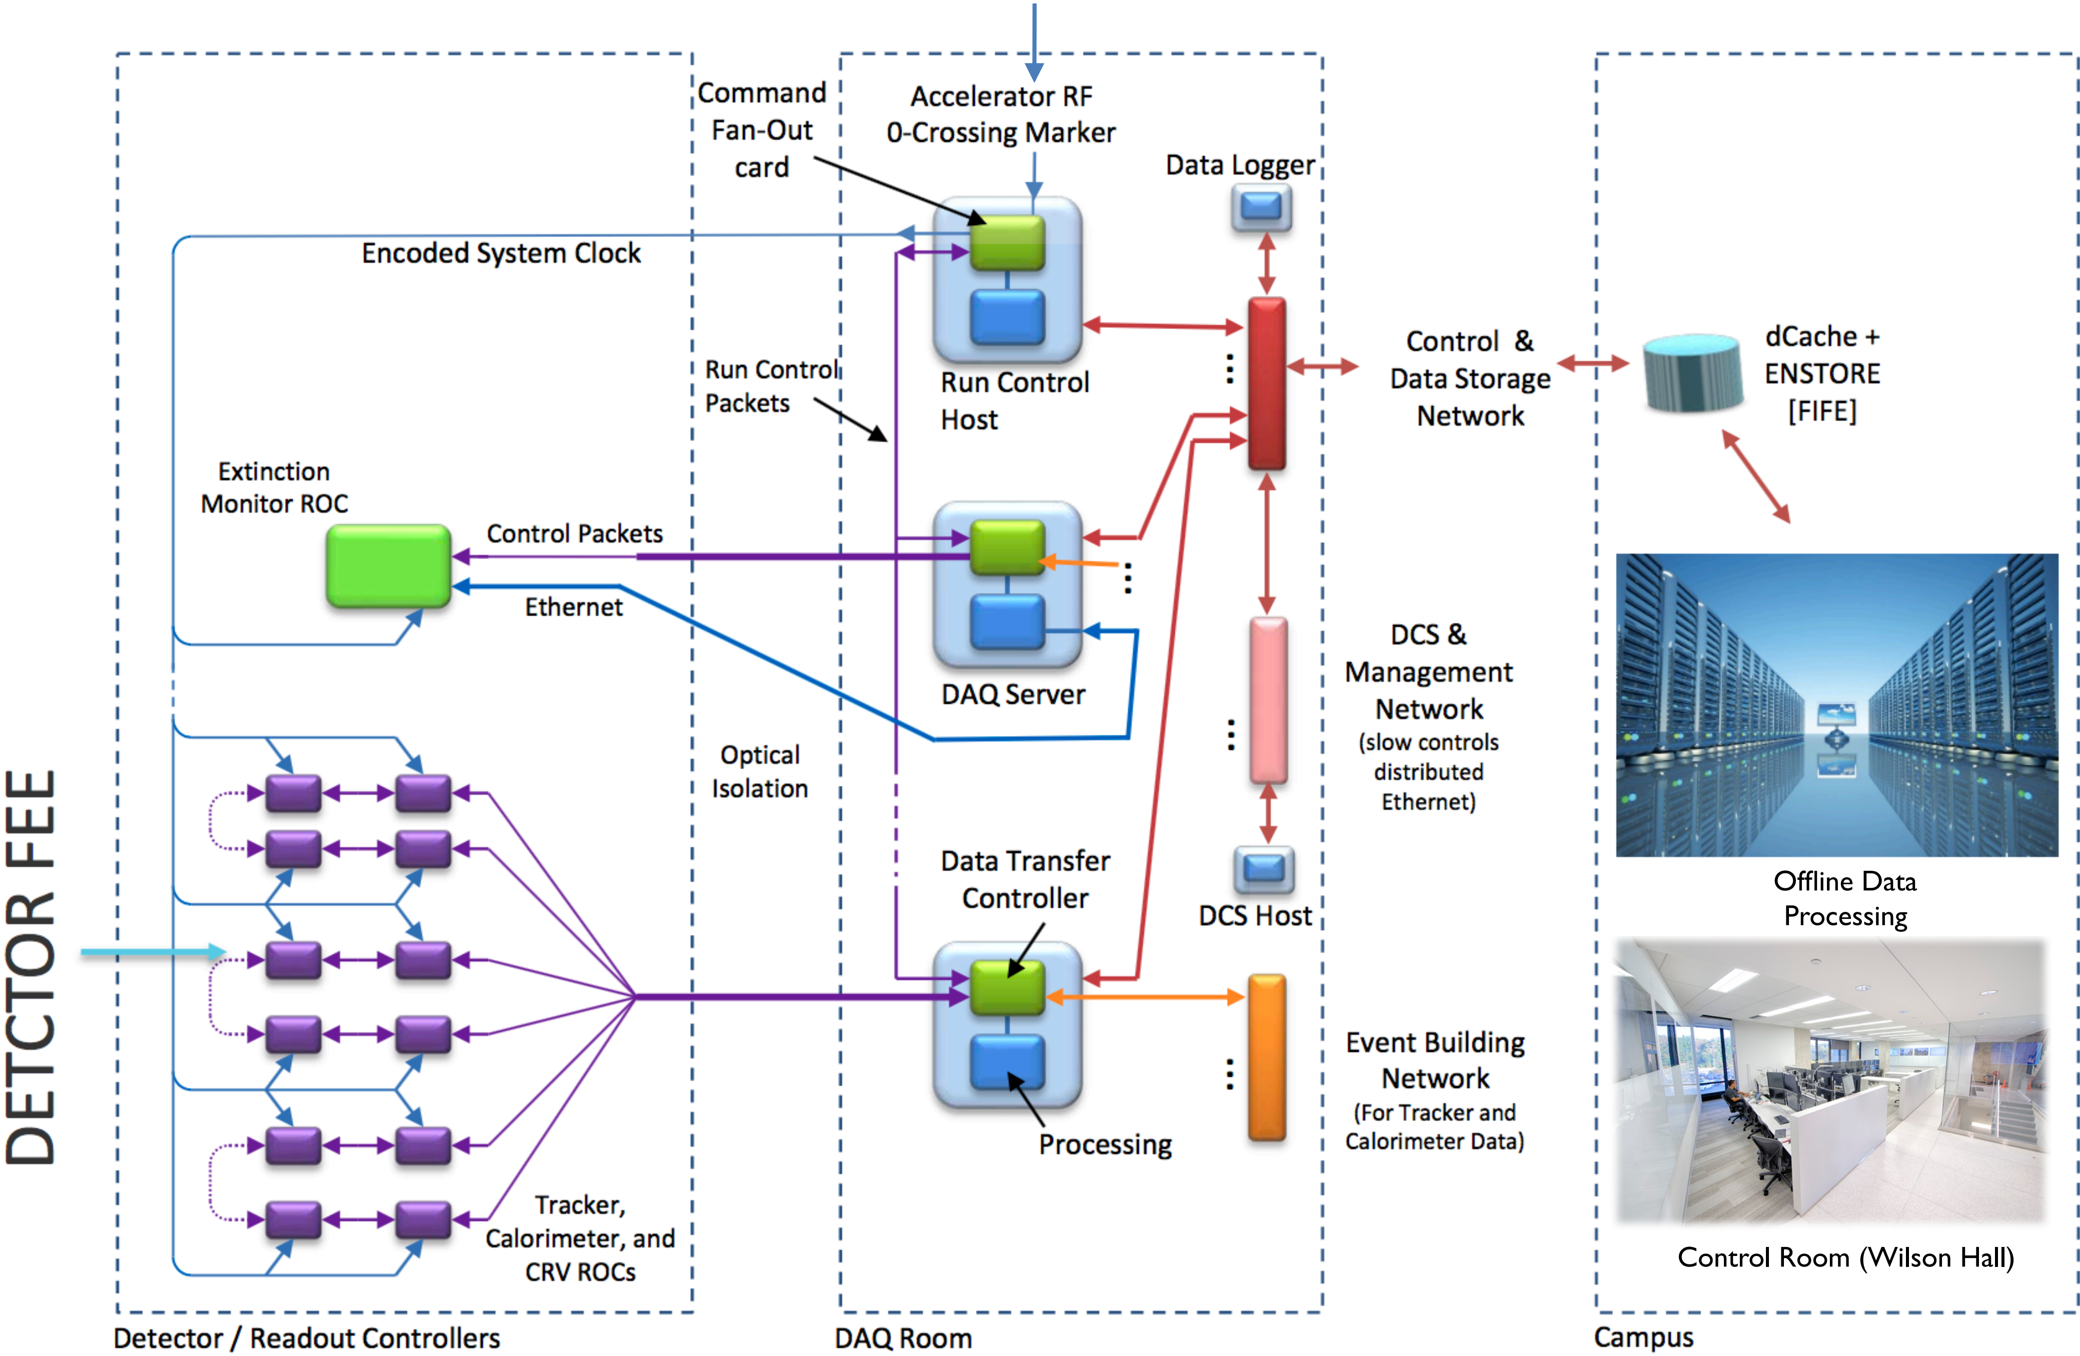

Hardware Architecture II

A ReadOut Controllers (ROCs) read out data from the various subsystems

B These ROCs are connected via optical links and readout by Data Transfer Controller (DTC) PCIe cards

C The PCIe cards will be connected to DAQ servers which read packets from the DTCs, convert them and pass them through the processing/filtering pipeline

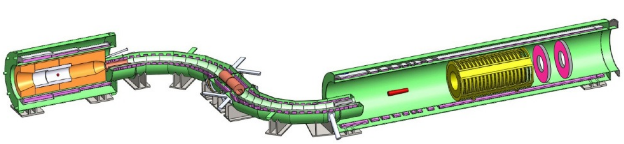

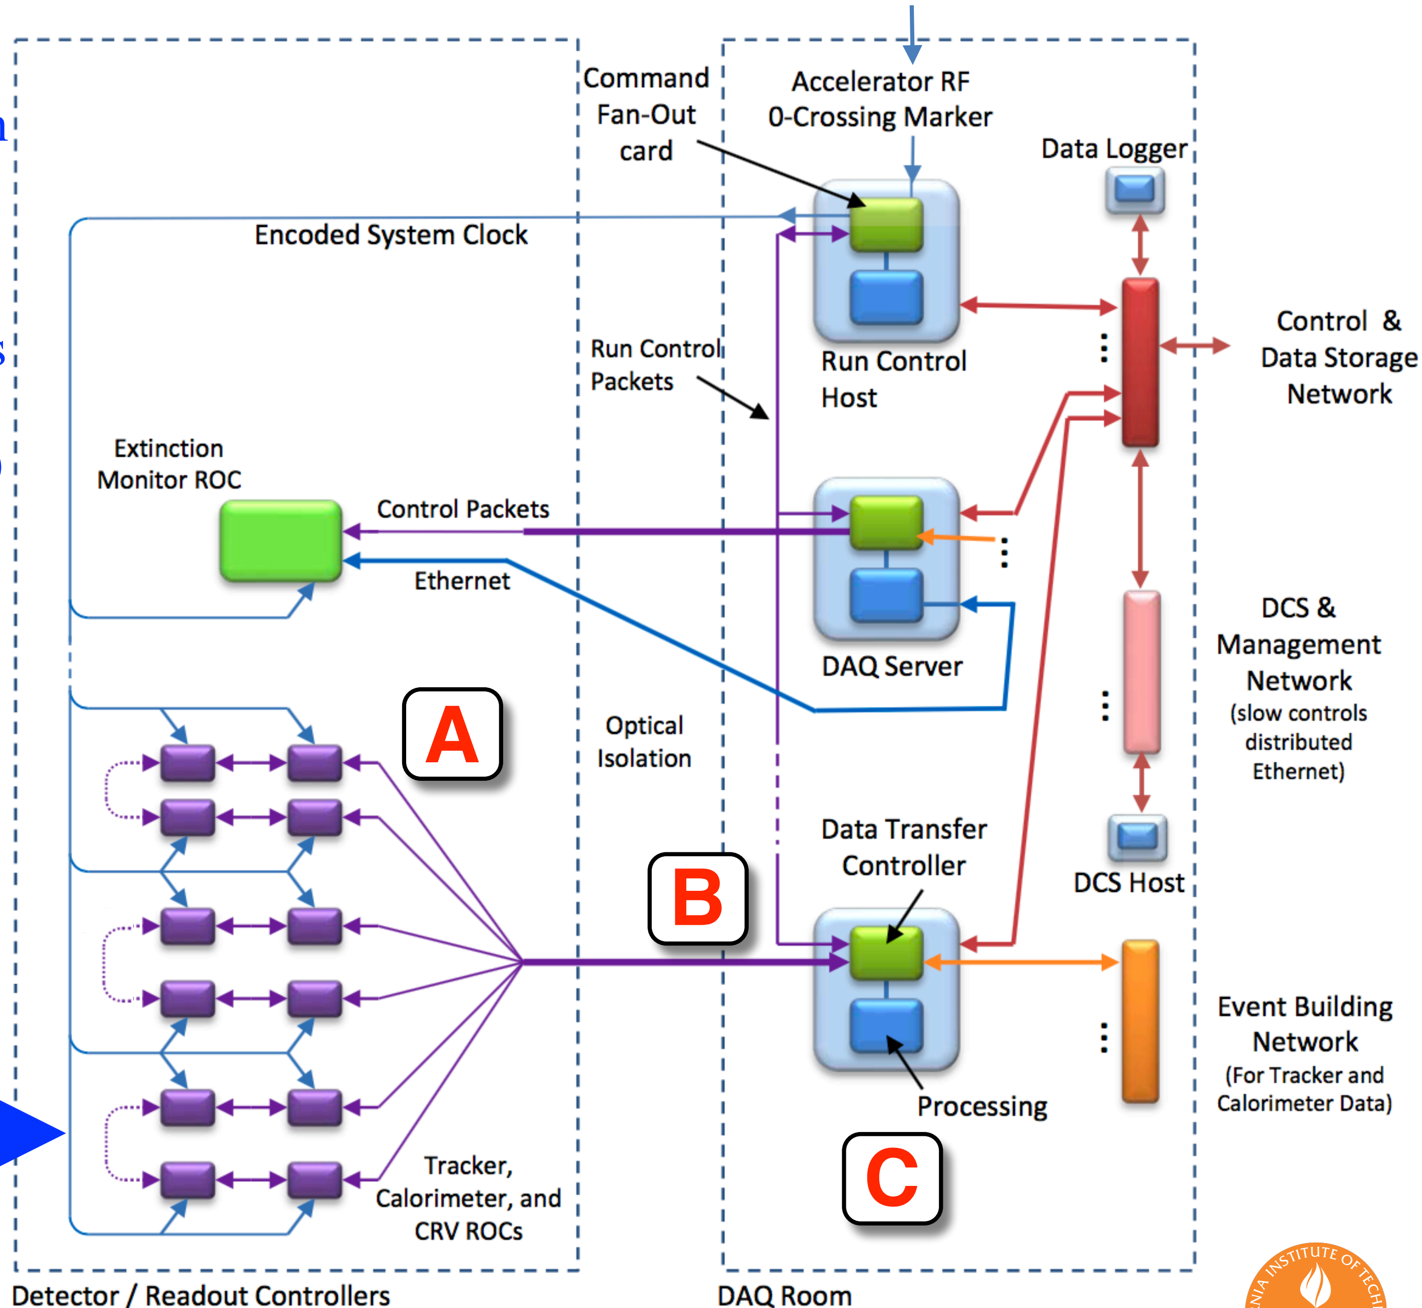

Hardware Architecture III

D An Event Building network is used to route raw data by time slices (event windows) and deliver all the data for a single time slice to the same DAQ server

E The Command Fanout Module (CFO) synchronizes ROCs by providing system clock and event window marker using one signal

F The Detector Control System allows experimenters to monitor the status and health of the detector, and provides slow controls for voltages, etc.

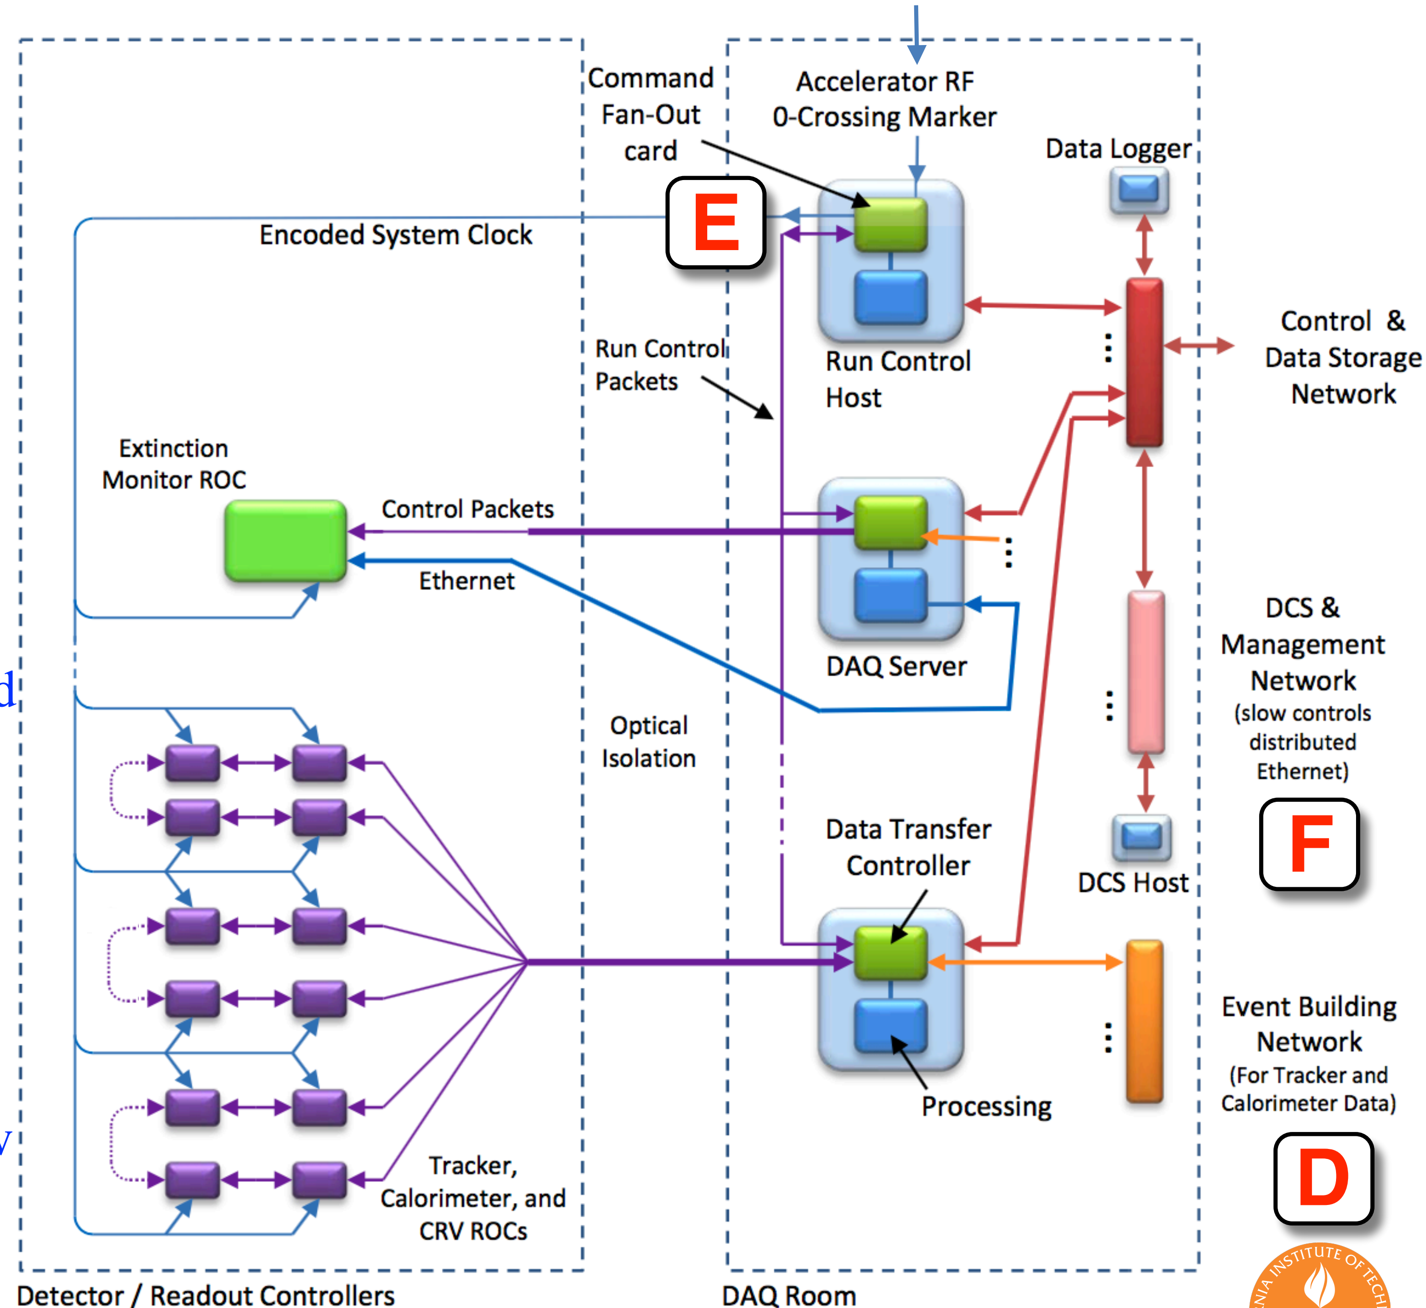

Hardware Architecture IV

- Server connections to ROCs:
  - Two optical links to each ROC
  - After any single link failure, we can still readout all ROCs
- Sub-detectors are responsible for front-end Readout Controllers (ROCs) and signal fanout inside the vacuum
- TDAQ uses a data pull system
  - Data flow is initiated by the Run Control Host via the Command Fanout module (CFO)
- Architecture supports both streaming (tracker and calorimeter) and triggered (CRV) readout

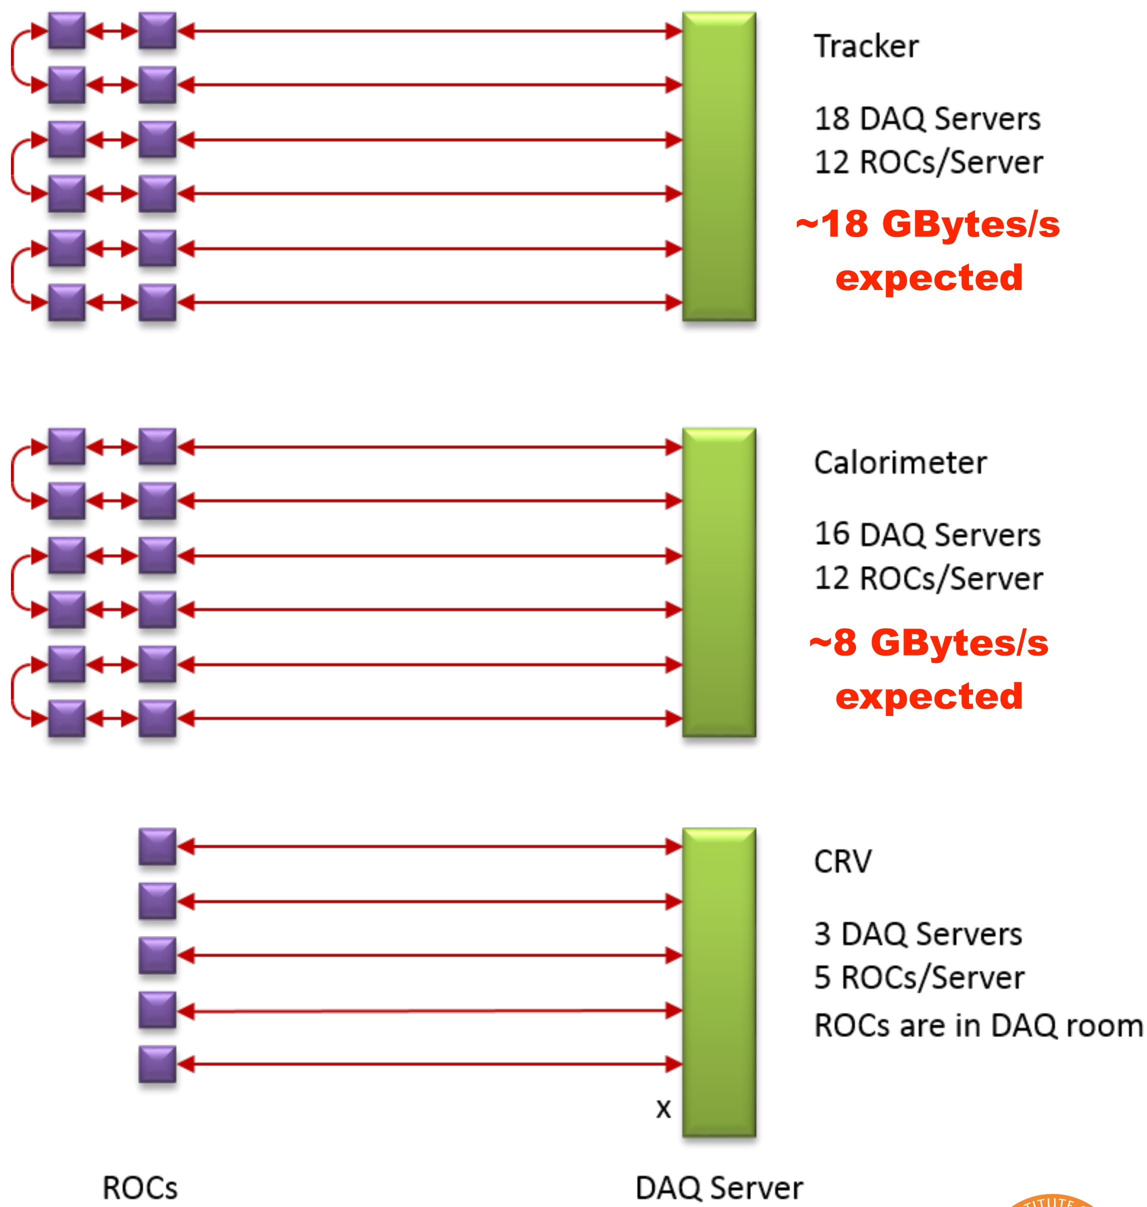

Hardware

- DAQ Servers (40)
  - 3U rack-mount computers
  - Integrated DAQ and online processing (~200k Events/s)
  - Easily scalable so, if necessary, we could increase processing performance by adding servers

- Data Transfer Controller Modules (1 per server)
  - Implemented using commercial PCIe card with FPGA, memory, and an 8-port optical interface (3.125 Gbps links)
  - 1 GByte/sec readout bandwidth
  - Same hardware used for CFO

- Hardware Event Building:
  - DTCs are connected to each other through a 48 port 10G Ethernet switch
  - For each event, each DTC passes data from its fraction of the detector to the event building network which compiles full-detector events and distributes them to DTCs in a round-robin order

Pilot System

- We have built a 6 server TDAQ pilot system
- Useful in understanding bandwidth capacity in the system
- The Pilot system is capable of delivering 60 kHz of simulated Mu2e events to processor nodes in playback mode
    - Equivalent to 360 kHz in full-size system
- Using fully simulated data packets, this pilot system will serve as a valuable testbed for the subsystems' reconstruction and filtering algorithms
- Currently working on adding CFO server to pilot system
    - Prototype timing measurements show peak-to-peak jitter of recovered clock at ~150ps with RMS jitter<50ps, satisfying requirements

Software

- As mentioned earlier, we use the artdaq framework for online processing
- Framework provides built in support for features common to many experiments
- The collaboration adds support for features specific to Mu2e such as:
  - Raw data formats from the ROCs
  - Libraries to interface with the DTCs
  - art “DataProducts” needed by filtering algorithms
  - art modules implementing filtering algorithms

Software Progress

- We have developed a communication protocol between the ROCs and the DTCs based on 128 bit packets
    - For each subsystem, a header packet precedes a set of data payload packets.
- Using a full digi-level simulation of the detector response to physics events, we have generated binary files containing simulated hits for the TRK and CAL
    - We can either load these packets into a physical DTC and pass them through the data processing chain, or run in fully simulated DTC mode and pass the packets to artdaq in software
- Once the packets are in the artdaq processing chain, we can now read them in within an art module, parse out the relevant physics info, and generate the same art DataProducts used in the Offline environment
    - Art events containing these DataProducts can then be passed through filter algorithms developed in the Offline environment

Mu2e Trigger Algorithms

- We have now developed, and are continuing to improve, several prototype filter modules within the offline framework:
  - Tracker Trigger
    - Consists of a combination of several filters using only tracker information
  - Calorimeter + Tracker Trigger
    - Uses only TRK hits within a ±50ns window around good CAL cluster with at least 50 MeV deposited
  - Standalone Calorimeter Trigger
    - Conversion Electron (CE) trigger
    - Track Candidate Trigger (mainly Decay In Orbit events with p>90 MeV/c) for studying TRK trigger and track reconstruction efficiency

Trigger Optimization

- Given our processing requirements, we need to achieve a maximum of ~4 ms/event:

$$200\text{K}\ \frac{Events}{s},\ 40\ \text{nodes},\ 20\ \frac{art\,Threads}{Nodes} \rightarrow \mathbf{4}\ \frac{ms}{\frac{Events}{art\,Threads}}$$

- We are now testing the performance of our filter algorithms in the online data processing environment where we can play back fully simulated data packets from physical DTC cards
    - We iterate by improving performance in the offline development environment, then testing our improved filters in the online environment
    - There are a number of software optimizations that can be made
    - If necessary, we can also achieve an ~20% improvement by buying more servers
    - We could also use preprocessing in FPGAs to further improve performance

Next Steps

- Our focus now is on algorithm development
  - Coordinating with sub-detector groups to optimize their filter algorithms for online processing
  - Exploring the use of FPGA-based preprocessing
- Also working towards end-to-end hardware+software tests using prototype TRK/CAL ROCs
- Timeline:
  - 2019: Purchase computing hardware
  - 2020: Cosmic ray tests
  - 2022: Begin running with beam

Backup Slides

Command Fanout (CFO) Module

- Implemented using same PCIe card as DTC and custom daughterboard with 8 optical transceiver ports
- Installed in Run Control Host
- Receives μBunch marker from accelerator, then generates encoded System Clock
- CFO follows an experiment-provided Run Plan to generate a "Heartbeat" packet in advance of each Event Window
    - The Heartbeat Packet provides timestamp and specifies operating mode for next Event Window

DAQ Pull Architecture

- DAQ is based on pull system:
  - DAQ requests and reads data from subsystems
  - No current plans for push system where subsystems request readout
- Pull system is very flexible
  - List of subsystems and readout frequency (integration time) can be easily tailored on the fly to fit the situation (e.g., on-spill, off-spill, calibration, etc.)
- If necessary, a subsystem can run a low-level hardware trigger and send its state along with the data, or choose to send an empty event if nothing interesting has been found
  - For example, the CRV can check whether a muon has crossed the tracker and encode this information in the data stream, or clear the CRV data if nothing went through
- The DAQ can initiate the calibration systems and record the corresponding data with appropriate readout frequency

Timing Signals

Encoded System Clock (40MHz), along with the Heartbeat packet defines contiguous Event Windows. Marker is alternating 25/75 and 75/25 duty cycle clock periods and is time-zero reference for Readout Controller Internal Timestamp counters.

Control/Data Ring – 3.125 Gbps (DAQ ↔ Readout Controllers). A Heartbeat packet is sent to give fine granularity TDC offset of RF 0-Marker relative to the System Clock marker and to specify Readout Controller behavior for the next Event Window. This control packet includes the Event Window tag, and is delivered at least 850ns before the Event Window marker to which it refers.

Readout Controller acquisition clock (e.g. 100 or 200 MHz - varies by detector) is phase-locked to System Clock - drives Readout Controller Internal Timestamp counter and ADCs/TDCs.

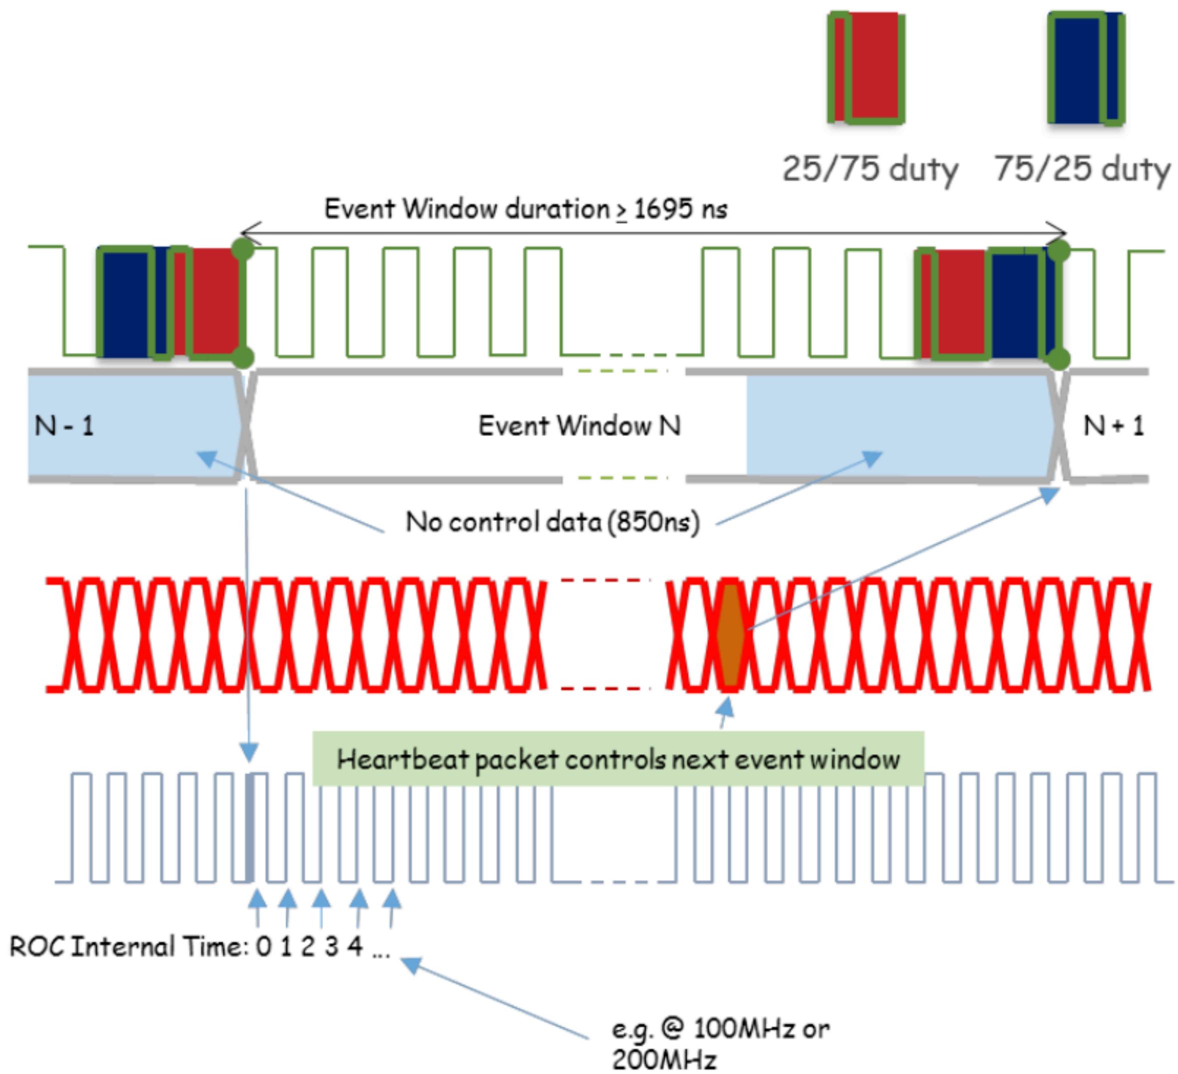

Rack Plan

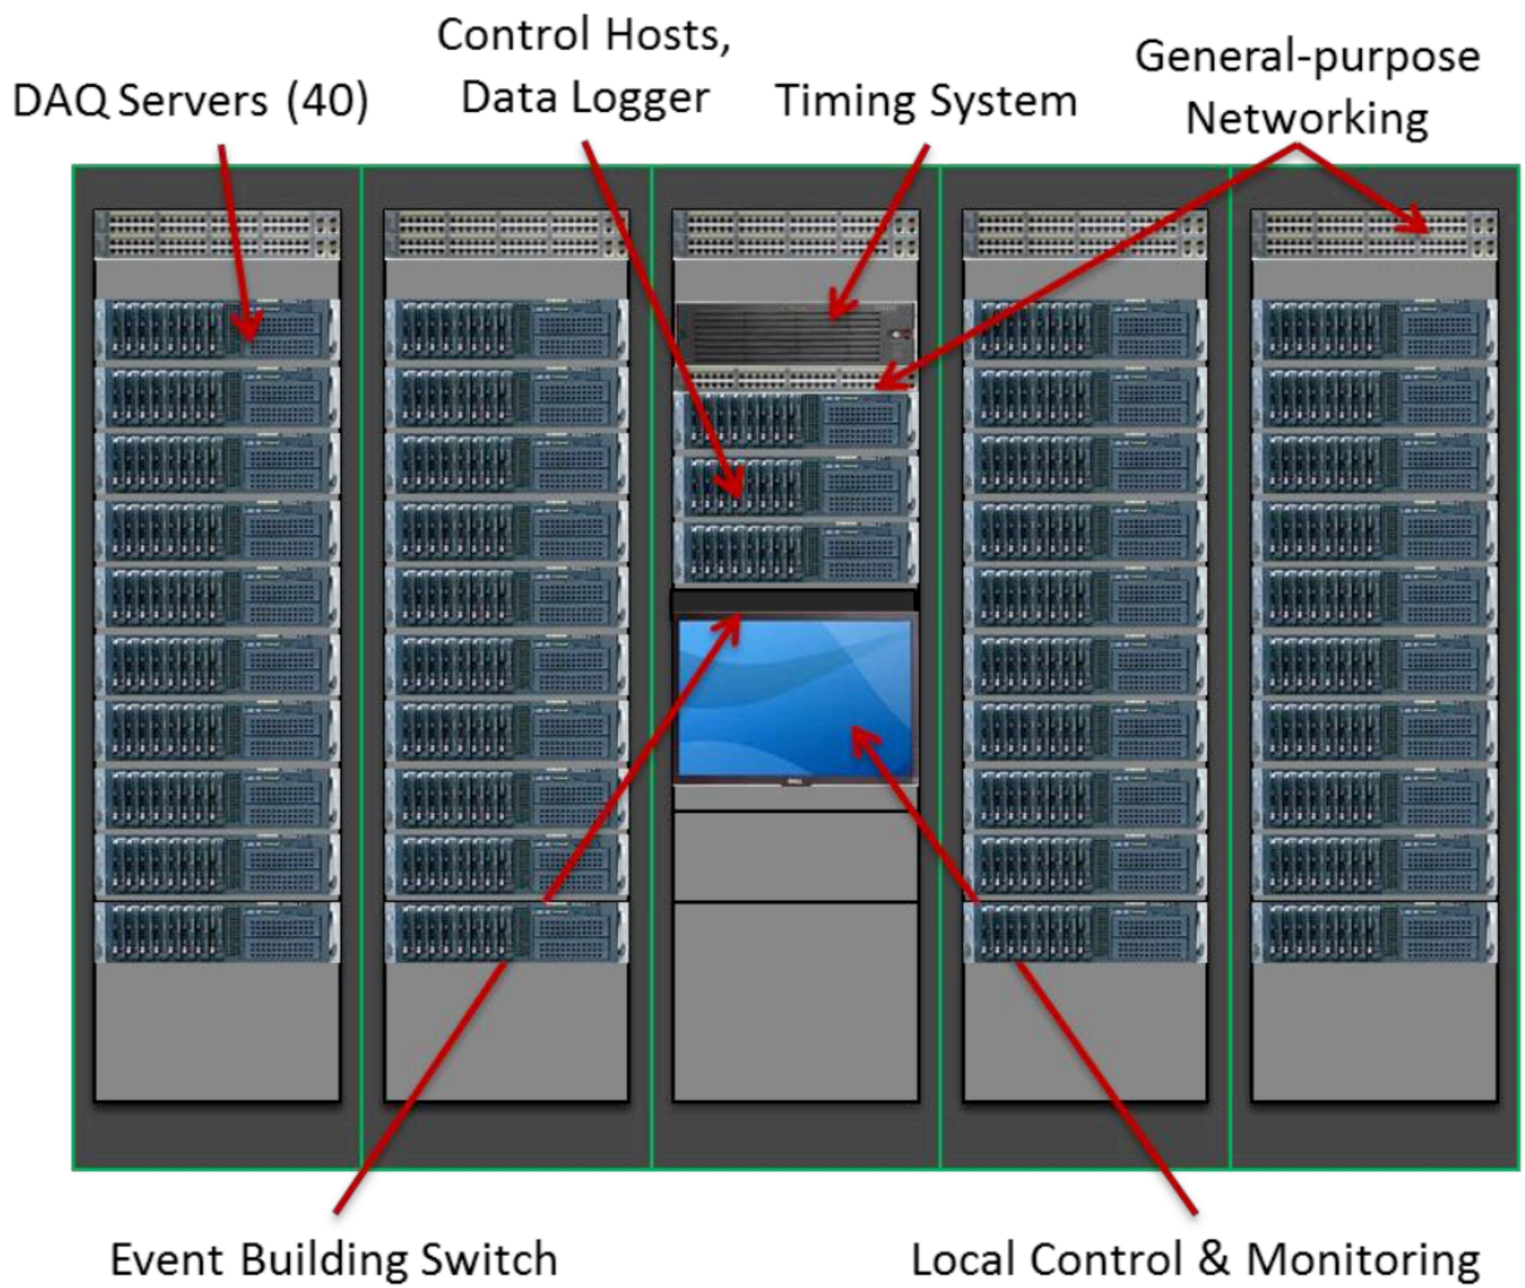

DAQ Servers (40)
Control Hosts,
Data Logger
Timing System
General-purpose
Networking
Event Building Switch
Local Control & Monitoring

Estimated Average Data Rates

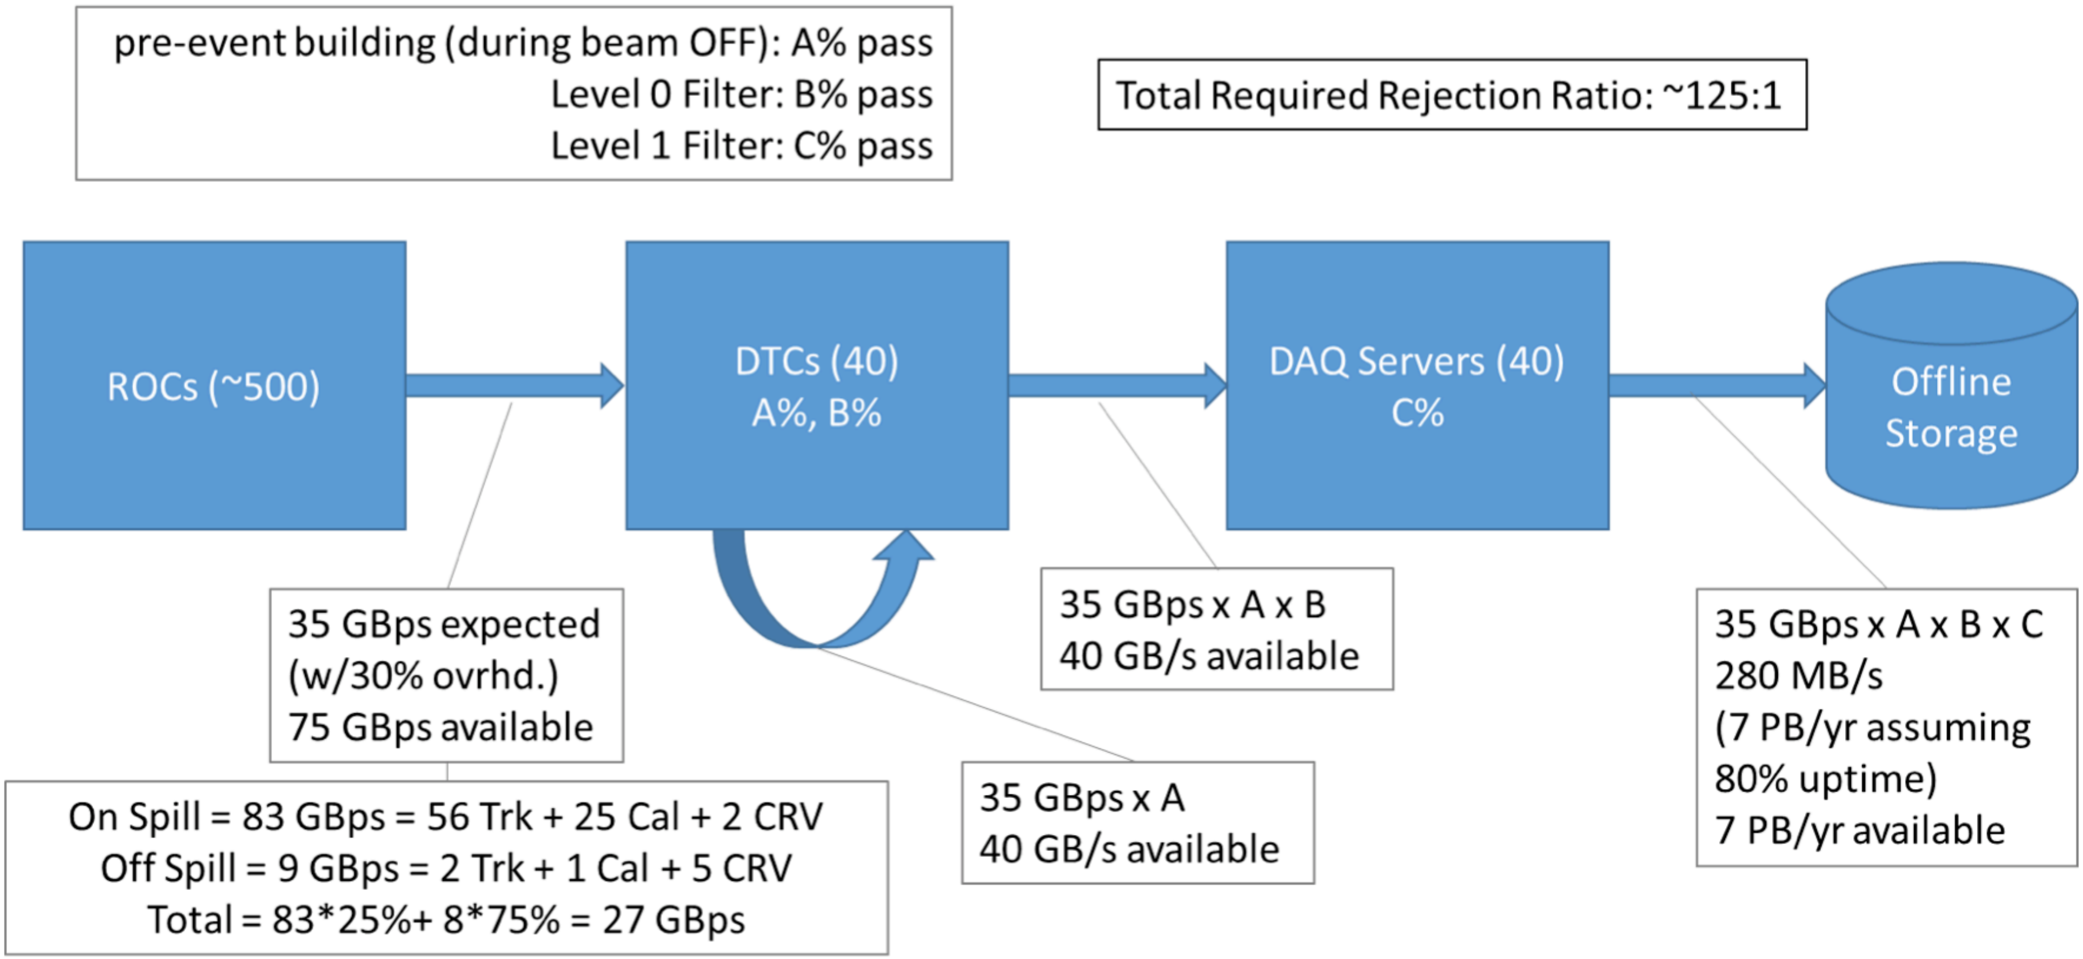

Control Room

- Mu2e Remote Control Room - Wilson Hall, 1st floor West, Fermilab
  - Shared use - LBNF, MicroBooNE, MINERvA, MiniBooNE, Minos, Muon g-2, Mu2e, NOvA

Detector Control System (DCS)

- Detector Control System (DCS) will allow experimenters and detector experts to monitor the status and health of the Mu2e detector
  - Will provide graphical user interface for monitoring currents, temperatures, magnetic field strength, pressure, and so on
- The DCS will also allow users to adjust voltages, turn subsystems on and off, etc
- EPICS has been selected as the underlying control and monitoring framework (modeled after NOvA and uBooNE)
- Baseline plan is to provide web-based monitoring and alarming

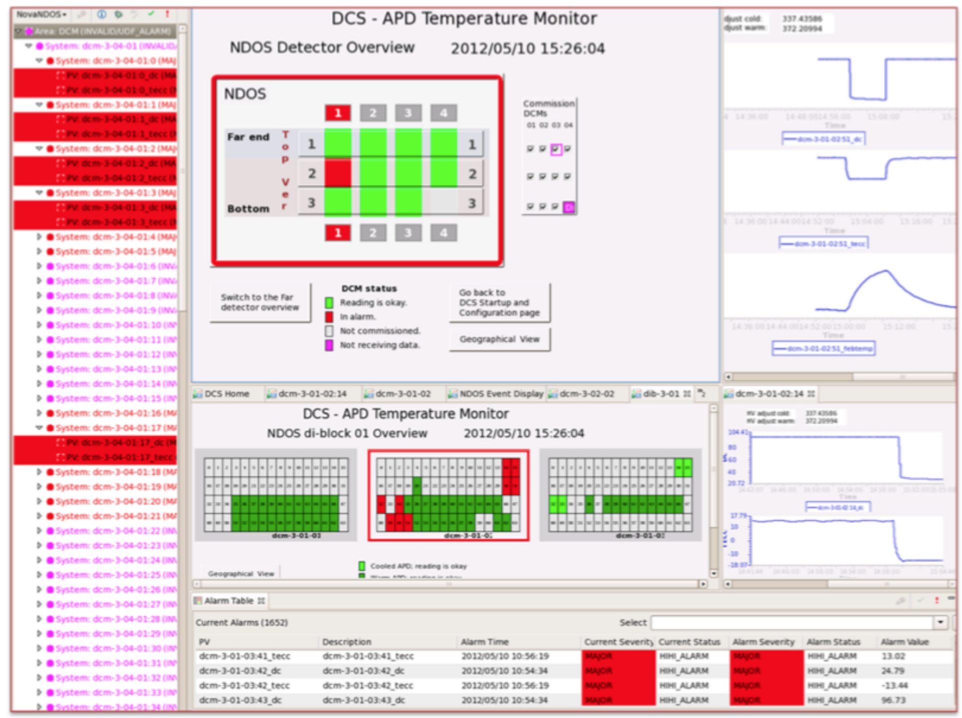

Control System Studio control and monitoring GUI from NOvA.

artdaq Software Architecture

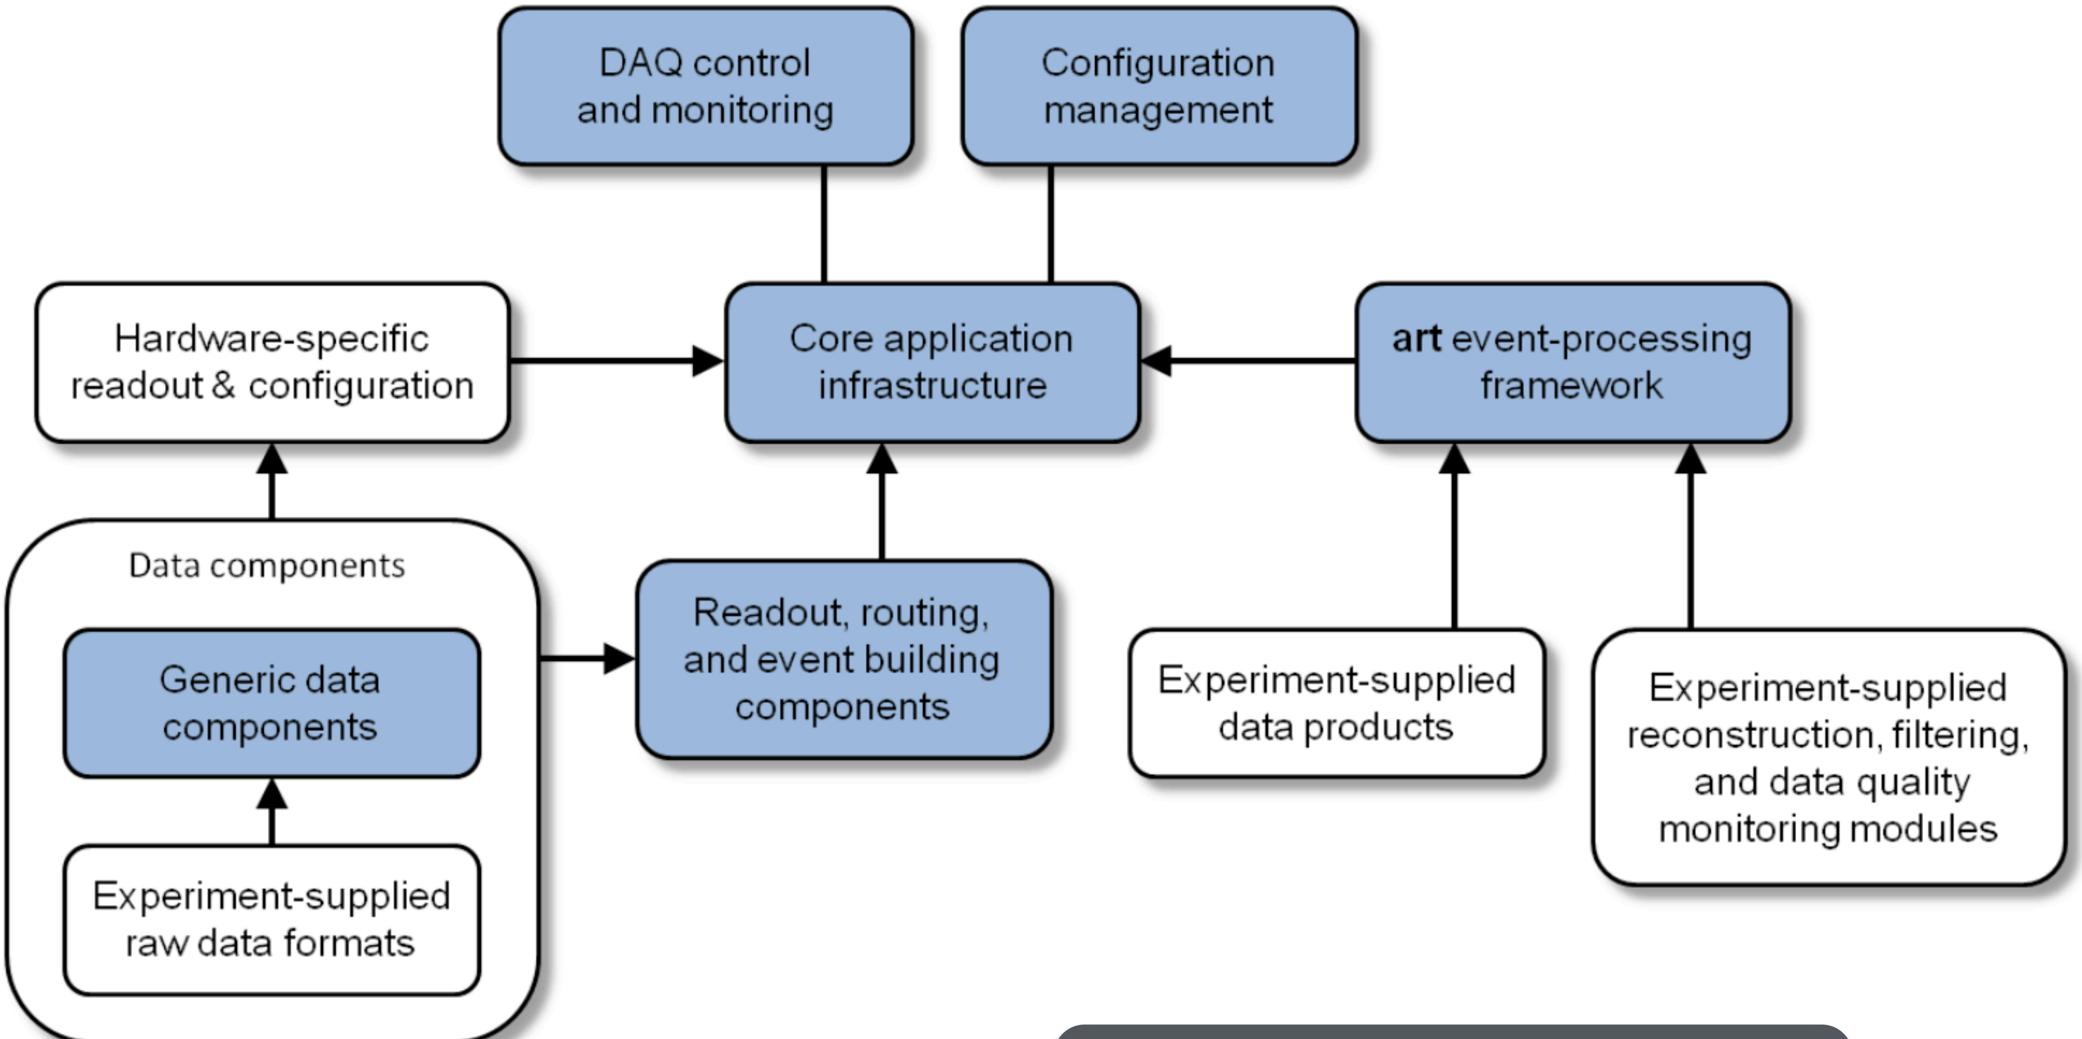

BLUE: Provided by artdaq
WHITE: Provided by experiment

artdaq Software Architecture II

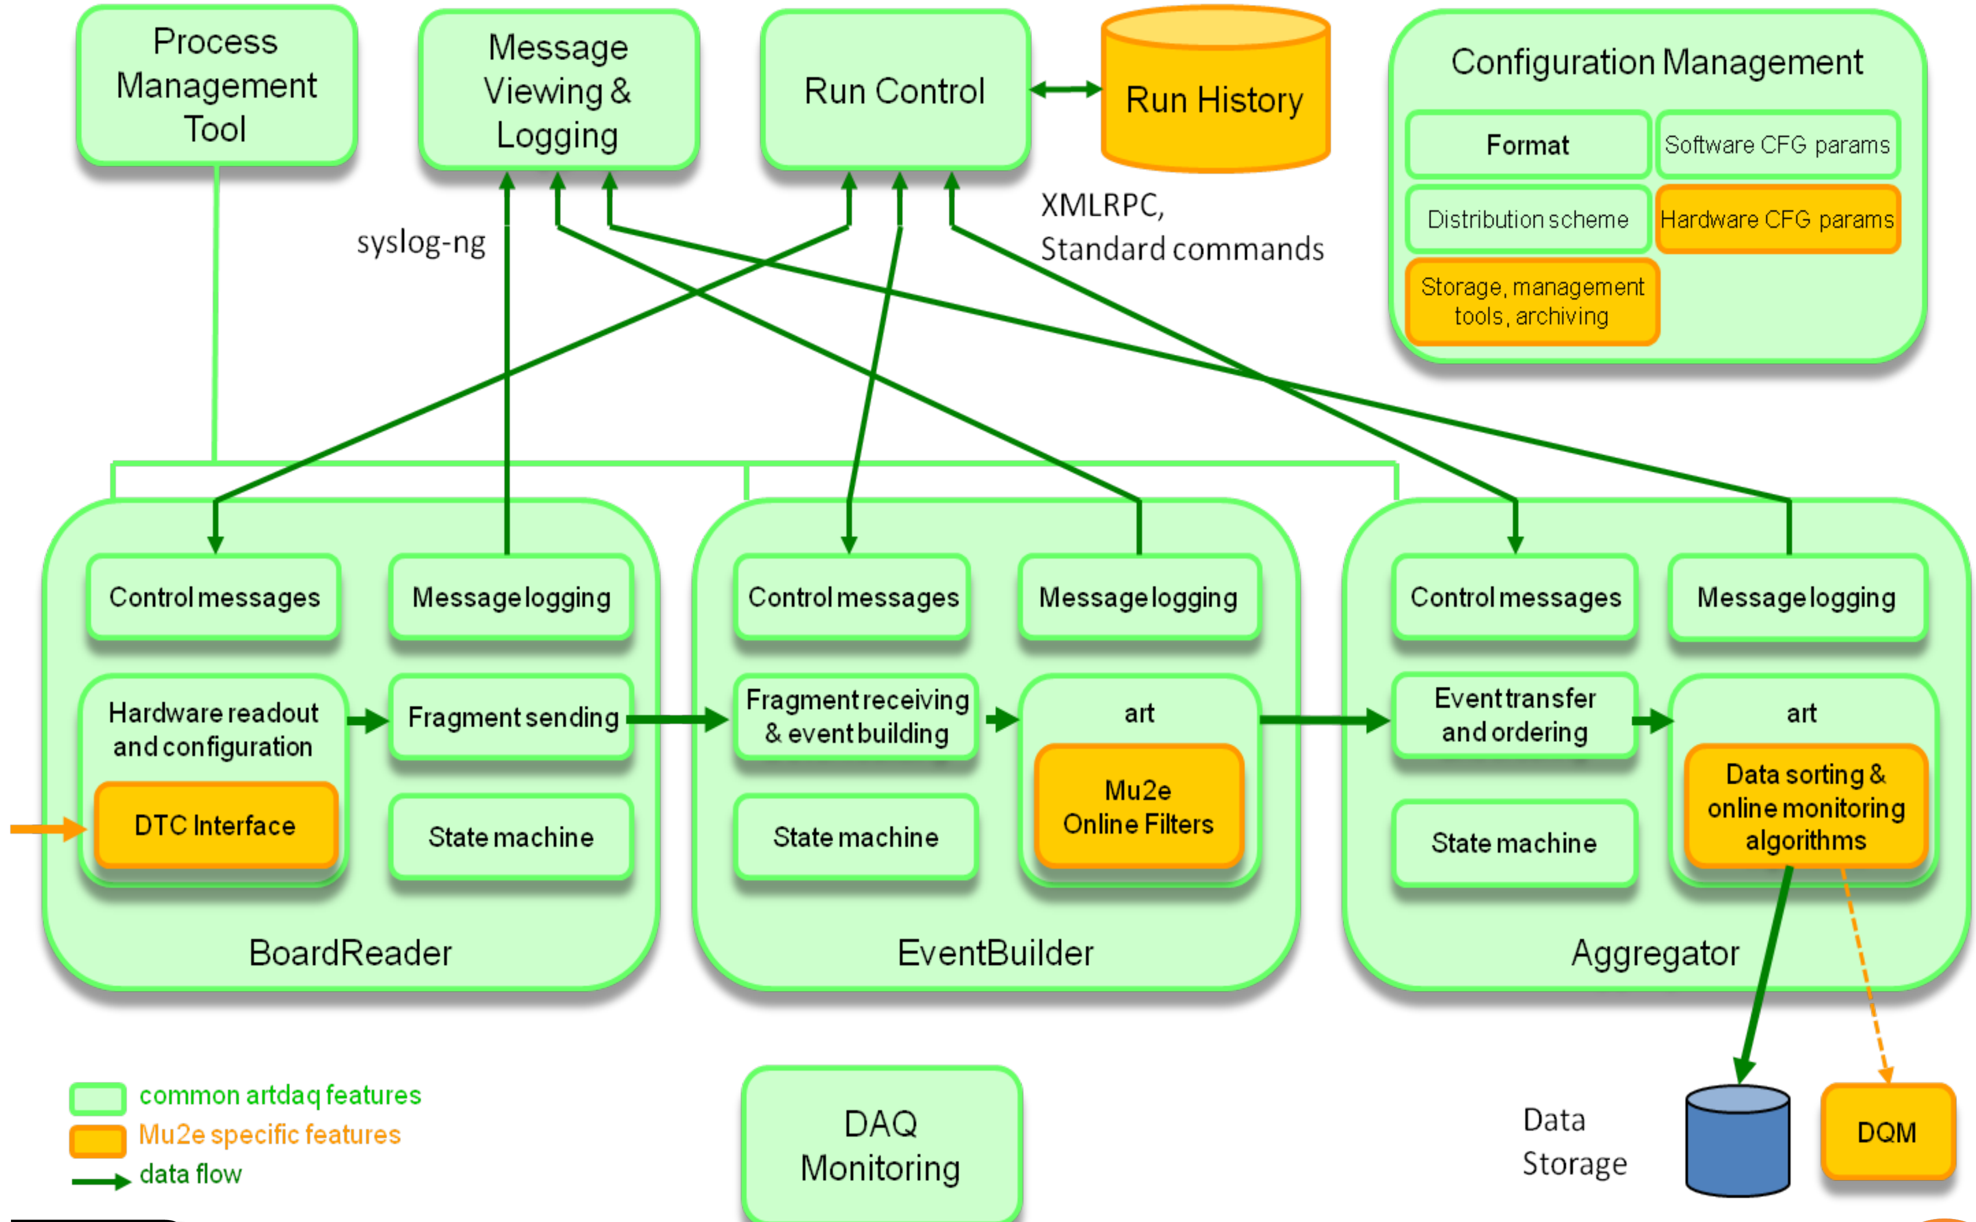

Mu2e Trigger Design I

- Trigger will run in streaming mode
  - Data continuously read from detector and time slices (events) are selected and written to tape for further processing
- Trigger will consist of a set of software filters fully configurable at run time
  - Also considering possible preprocessing in FPGAs
- Data taking periods will be divided into two modes: on-spill and off-spill
  - On-spill covers periods when 8 Gev proton bunches are colliding with the production target
  - Off-spill covers all other periods: between bunches, calibration periods, commissioning, etc

Mu2e Trigger Design II

- Current plan for on-spill operation:
  - Time slice of TRK + CAL data is requested and read by DAQ, assembled by the event builder, and sent to a processing unit (CPU core)
  - If any of the TRK or CAL software filters are satisfied, data from the CRV is fetched and the full event is persisted (CRV has a ~1s buffer, so the decision must be made during that time window)
  - CRV data is not available during on-spill processing by the CPU
- Similar plan for off-spill operation
  - Currently studying CRV off-spill data rate
- Extinction Monitor and Stopping Target Monitor have independent triggers and provide data streams to TDAQ (the time and event header is provided by the TDAQ)

On-Spill Filter Summary

- Upstream / downstream track [TRK]:
  - High-momentum track selection from tracker information only
- Proton tracks (if needed) [TRK]:
  - Proton track selection for tracker dE/dx calibration and alignment
- Lower momentum electrons [TRK]:
  - Low-momentum electron selection for tracker alignment studies
- High-energy cluster [CAL]:
  - Conversion electron selection from CAL information only
- Dual calorimeter - tracker [CAL + TRK]:
  - Conversion electron selection using both TRK and CAL information
- Random Trigger:
  - Random trigger to study background
- Physics trigger [TBD]:
  - Additional trigger for other physics studies

Off-Spill Filter Summary

- CRV cosmics [CRV]:
  - Cosmic muon selection for CRV and background studies. Would also be used to flag muons going through the TRK and CAL for the corresponding calibrations
- Tracker cosmics / calorimeter cosmics [TRK / CAL]:
  - Tracker alignment and calorimeter energy calibrations from cosmic muons going through the tracker / calorimeter. Might be redundant with the CRV cosmic filter if the latter is not prescaled
- Tracker electronics [TRK]:
  - Calibration of the tracker electronics
- Calorimeter laser and electronics [CAL]:
  - Laser pulse sent into the crystals or electronic pulse sent into the electronics for calibration
- CRV calibration [CRV]:
  - Trigger former in the CRV electronics for calibration, efficiency, and stability studies
- Calorimeter source [CAL]:
  - Calorimeter source calibration