Responses to reviewers questions

DAQ design team

CERN, 4th November 2016

1. Produce an orgchart with leadership & person power in each box -is there a steering group?

[

See next slides

DUNE Organization

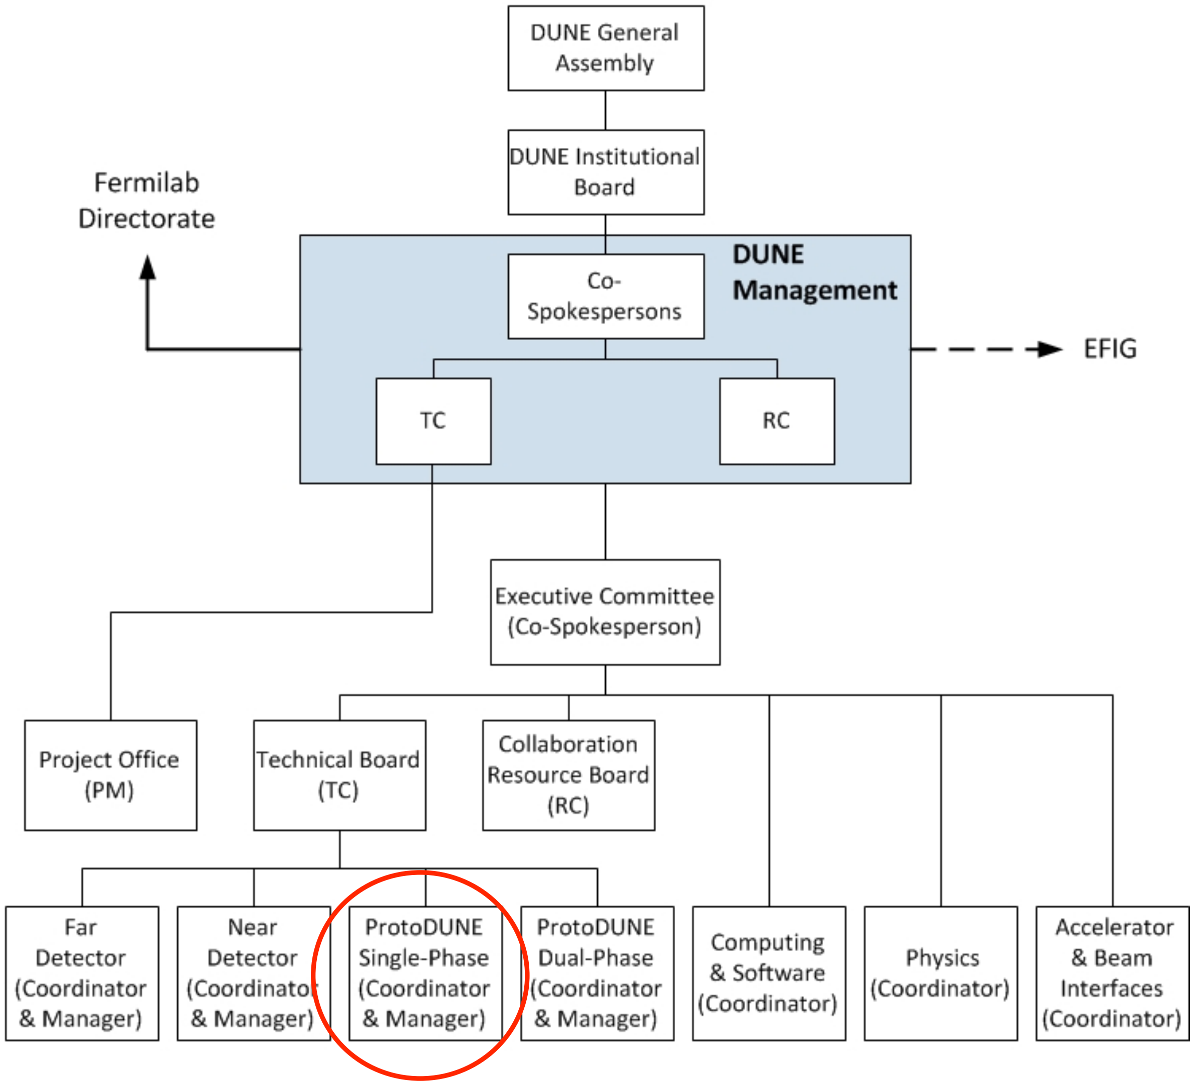

ProtoDUNE-SP Organization Chart

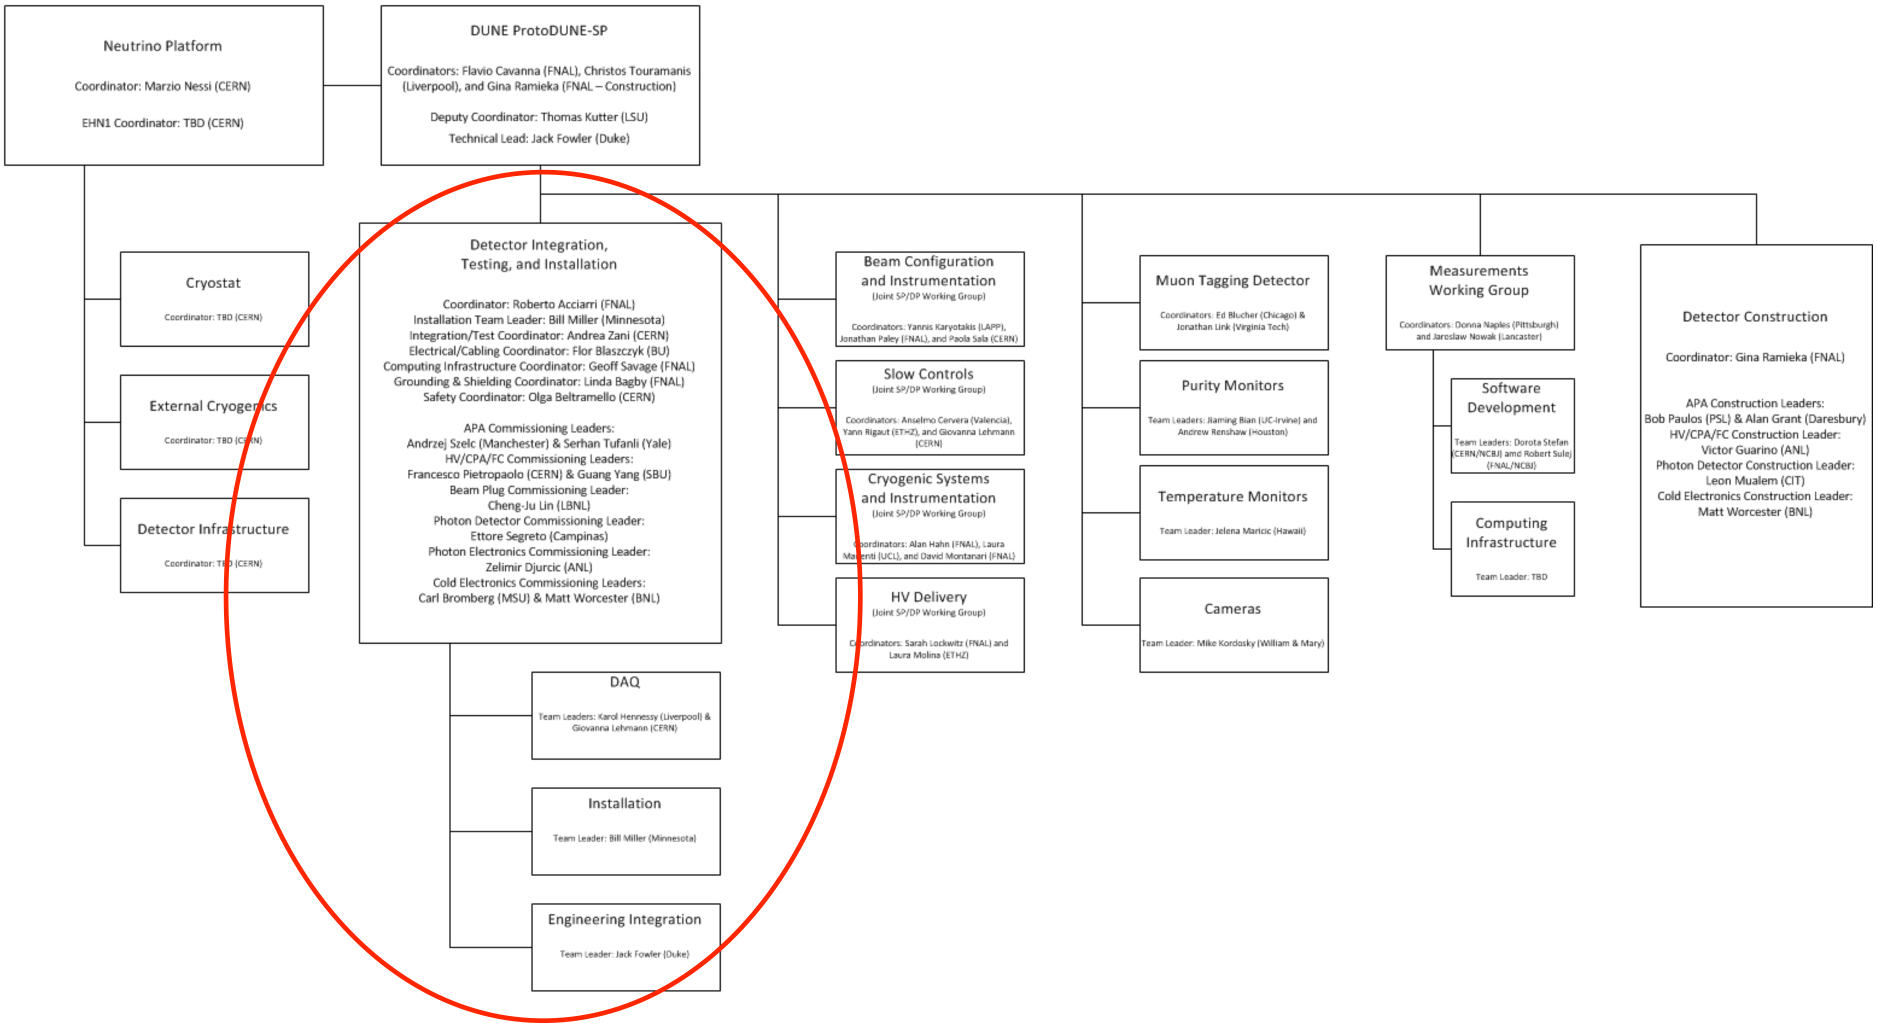

Detector Integration, Testing, & Commissioning

- CERN will provide overall coordination for the activities in the EHN1 area
- This is the on-ground team that will look after ProtoDUNE-SP installation, commissioning, and operation
- With one exception, all of these individuals have agreed to re-locate to CERN for extended periods over the next two years

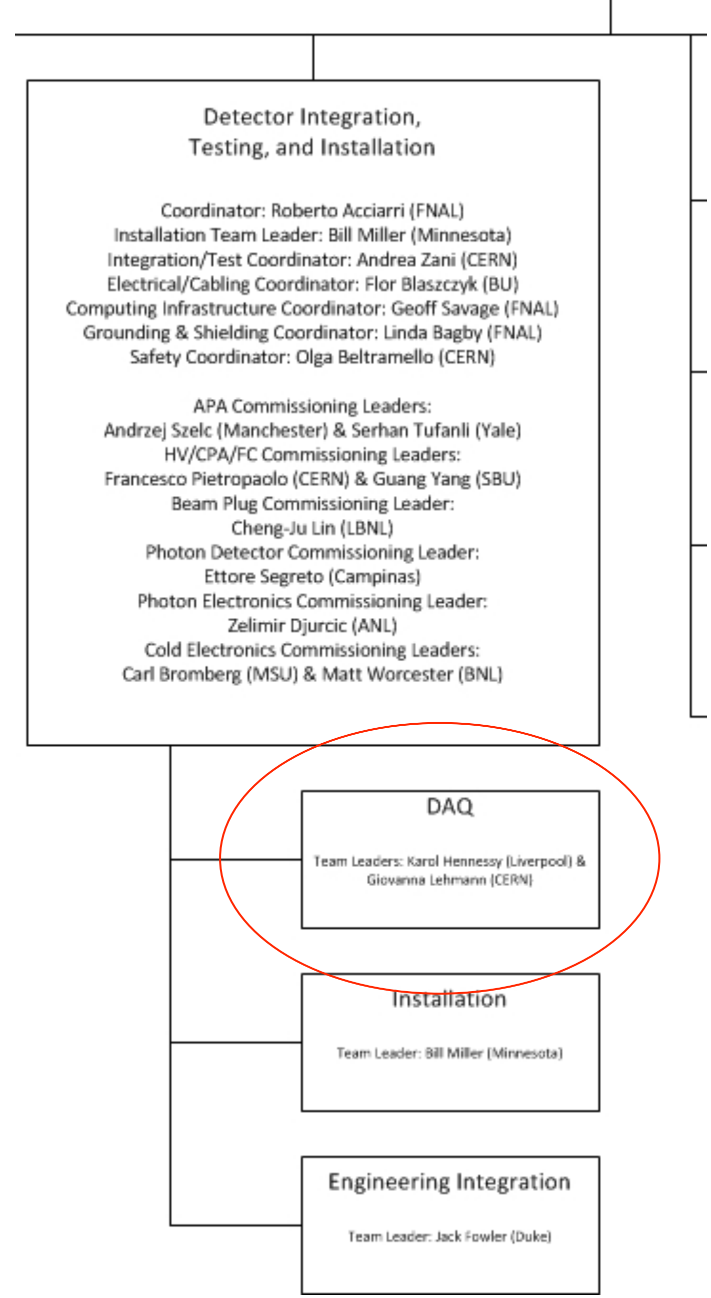

ProtoDUNE-SP DAQ Organization

DAQ
CERN, Liverpool

RCE Readout
SLAC, UC-Davis

FELIX Readout
CERN, NIKHEF, PNNL

ARTDAQ
FNAL, Oxford, RAL

SSP Readout
ANL, Warwick

Trigger/Timing
Bristol, Penn

Run Control
CERN, FNAL

Cosmic Tagger Readout
Virginia Tech

Monitoring
Sheffield, Sussex

Beam Instr. Readout
CERN, FNAL

ProtoDUNE-SP DAQ On-ground Team

K. Hennessey, G. Lehmann

DAQ
CERN, Liverpool

J. Wang, SLAC R.S.

RCE Readout
SLAC, UC-Davis

FELIX Readout
CERN, NIKHEF, PNNL

B. Abi, G. Barr, F. Azfar

ARTDAQ
FNAL, Oxford, RAL

Z. Djurcic, M. Haigh

SSP Readout
ANL, Warwick

N. Fiuza De Barros, D. Newbold

Trigger/Timing
Bristol, Penn

W. Ketchum

Run Control
CERN, FNAL

C. Mariani

Cosmic Tagger Readout
Virginia Tech

Monitoring
Sheffield, Sussex

J. Paley, CERN B.D.

Beam Instr. Readout
CERN, FNAL

2. Produce a global diagram (with tables if needed) with all links, boxes, bandwidths etc.

(Next page)

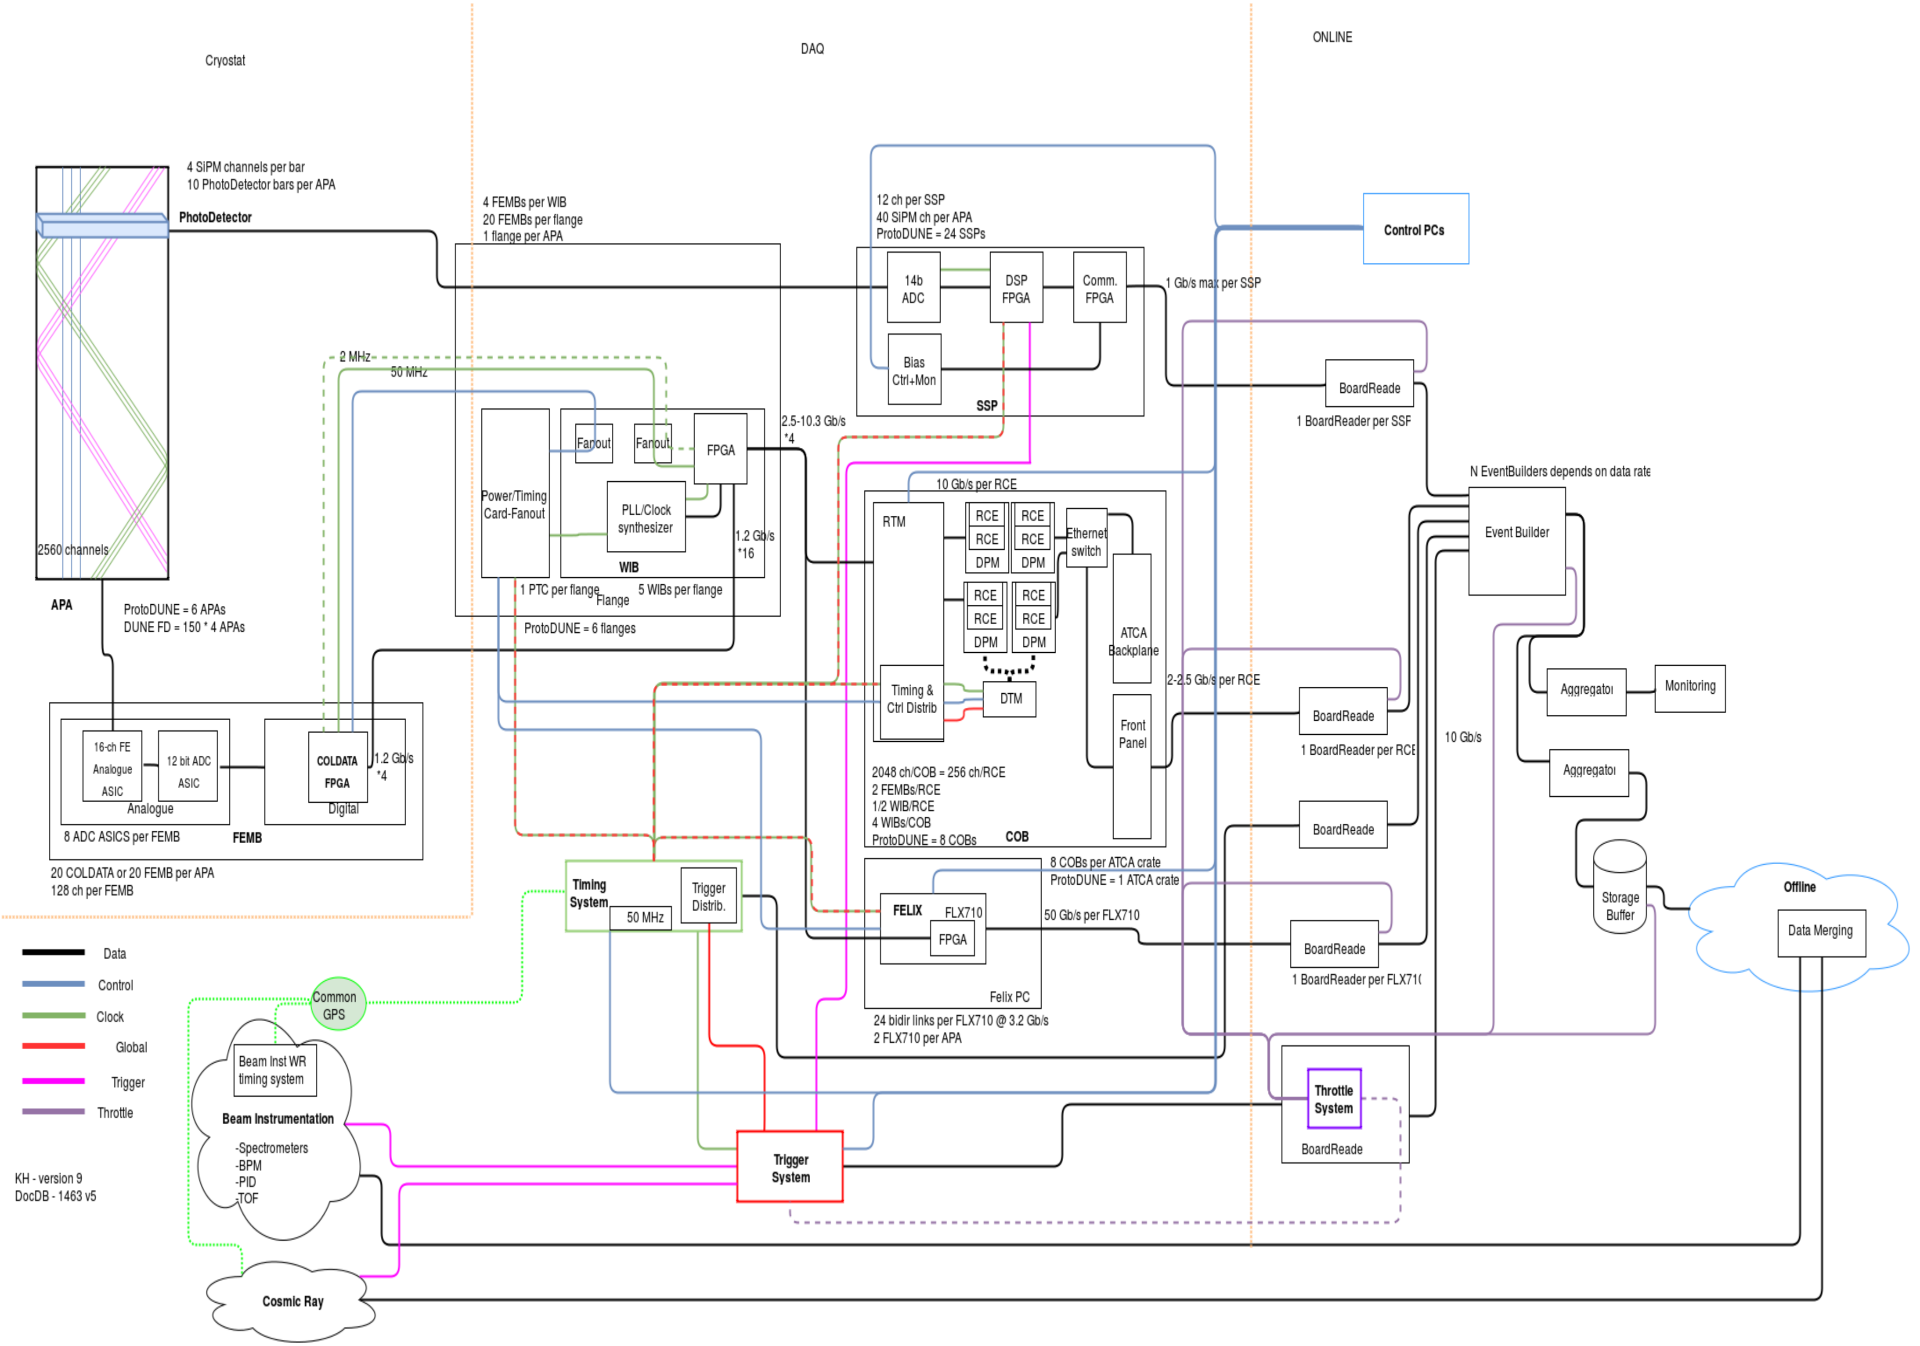

Cryostat
DAQ
ONLINE
4 SiPM channels per bar
10 PhotoDetector bars per APA
PhotoDetector
2560 channels
APA
ProtoDUNE = 6 APAs
DUNE FD = 150 * 4 APAs
4 FEMBs per WIB
20 FEMBs per flange
1 flange per APA
2 MHz
50 MHz
Power/Timing Card-Fanout
Fanout
Fanout
FPGA
PLL/Clock synthesizer
WIB
1 PTC per flange
Flange
5 WIBs per flange
ProtoDUNE = 6 flanges
1.2 Gb/s *16
2.5-10.3 Gb/s *4
12 ch per SSP
40 SiPM ch per APA
ProtoDUNE = 24 SSPs
14b ADC
DSP FPGA
Comm. FPGA
Bias Ctrl+Mon
SSP
1 Gb/s max per SSP
Control PCs
BoardReade
1 BoardReader per SSF
N EventBuilders depends on data rate
Event Builder
10 Gb/s per RCE
RTM
RCE
RCE
DPM
RCE
RCE
DPM
RCE
RCE
DPM
RCE
RCE
DPM
Ethernet switch
ATCA Backplane
Front Panel
Timing & Ctrl Distrib
DTM
2-2.5 Gb/s per RCE
2048 ch/COB = 256 ch/RCE
2 FEMBs/RCE
1/2 WIB/RCE
4 WIBs/COB
ProtoDUNE = 8 COBs
COB
8 COBs per ATCA crate
ProtoDUNE = 1 ATCA crate
BoardReade
1 BoardReader per RCE
BoardReade
10 Gb/s
Aggregator
Monitoring
Aggregator
Storage Buffer
Offline
Data Merging
16-ch FE Analogue ASIC
12 bit ADC ASIC
Analogue
COLDATA FPGA
Digital
1.2 Gb/s *4
8 ADC ASICS per FEMB
FEMB
20 COLDATA or 20 FEMB per APA
128 ch per FEMB
Timing System
50 MHz
Trigger Distrib.
FELIX
FLX710
FPGA
Felix PC
50 Gb/s per FLX710
24 bidir links per FLX710 @ 3.2 Gb/s
2 FLX710 per APA
BoardReade
1 BoardReader per FLX71(
Throttle System
BoardReade
Data
Control
Clock
Global
Trigger
Throttle
Common GPS
Beam Inst WR timing system
Beam Instrumentation
-Spectrometers
-BPM
-PID
-TOF
Trigger System
KH - version 9
DocDB - 1463 v5
Cosmic Ray

Bandwidth displayed from view of main switch

To switch: 0.3Gb each
From switch: 0

To switch: 10Gb each
From switch: 0

To switch: 1Gb each
From switch: 0

8 TPC
To switch: 0.3Gb each
From switch: 0.3Gb each

2 SSP
To switch: 1Gb each
From switch: 1Gb each

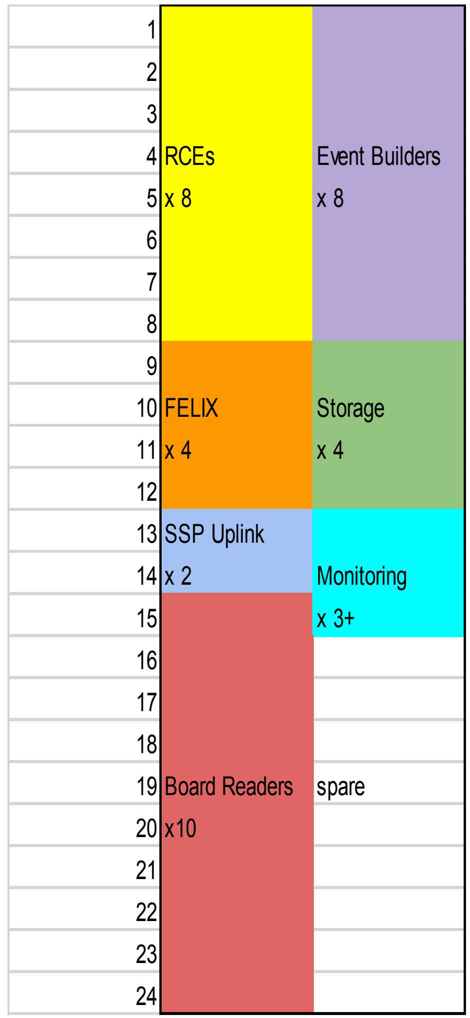

To switch: 0.3 Gb each
From switch: 0.3 Gb each

To switch: 0.6 Gb each
From switch: 0

To switch: 0
From switch: <1 Gb each

3. For exploitation talk tomorrow (or after), would like to see how things phase together (who depends on who) towards vertical slice

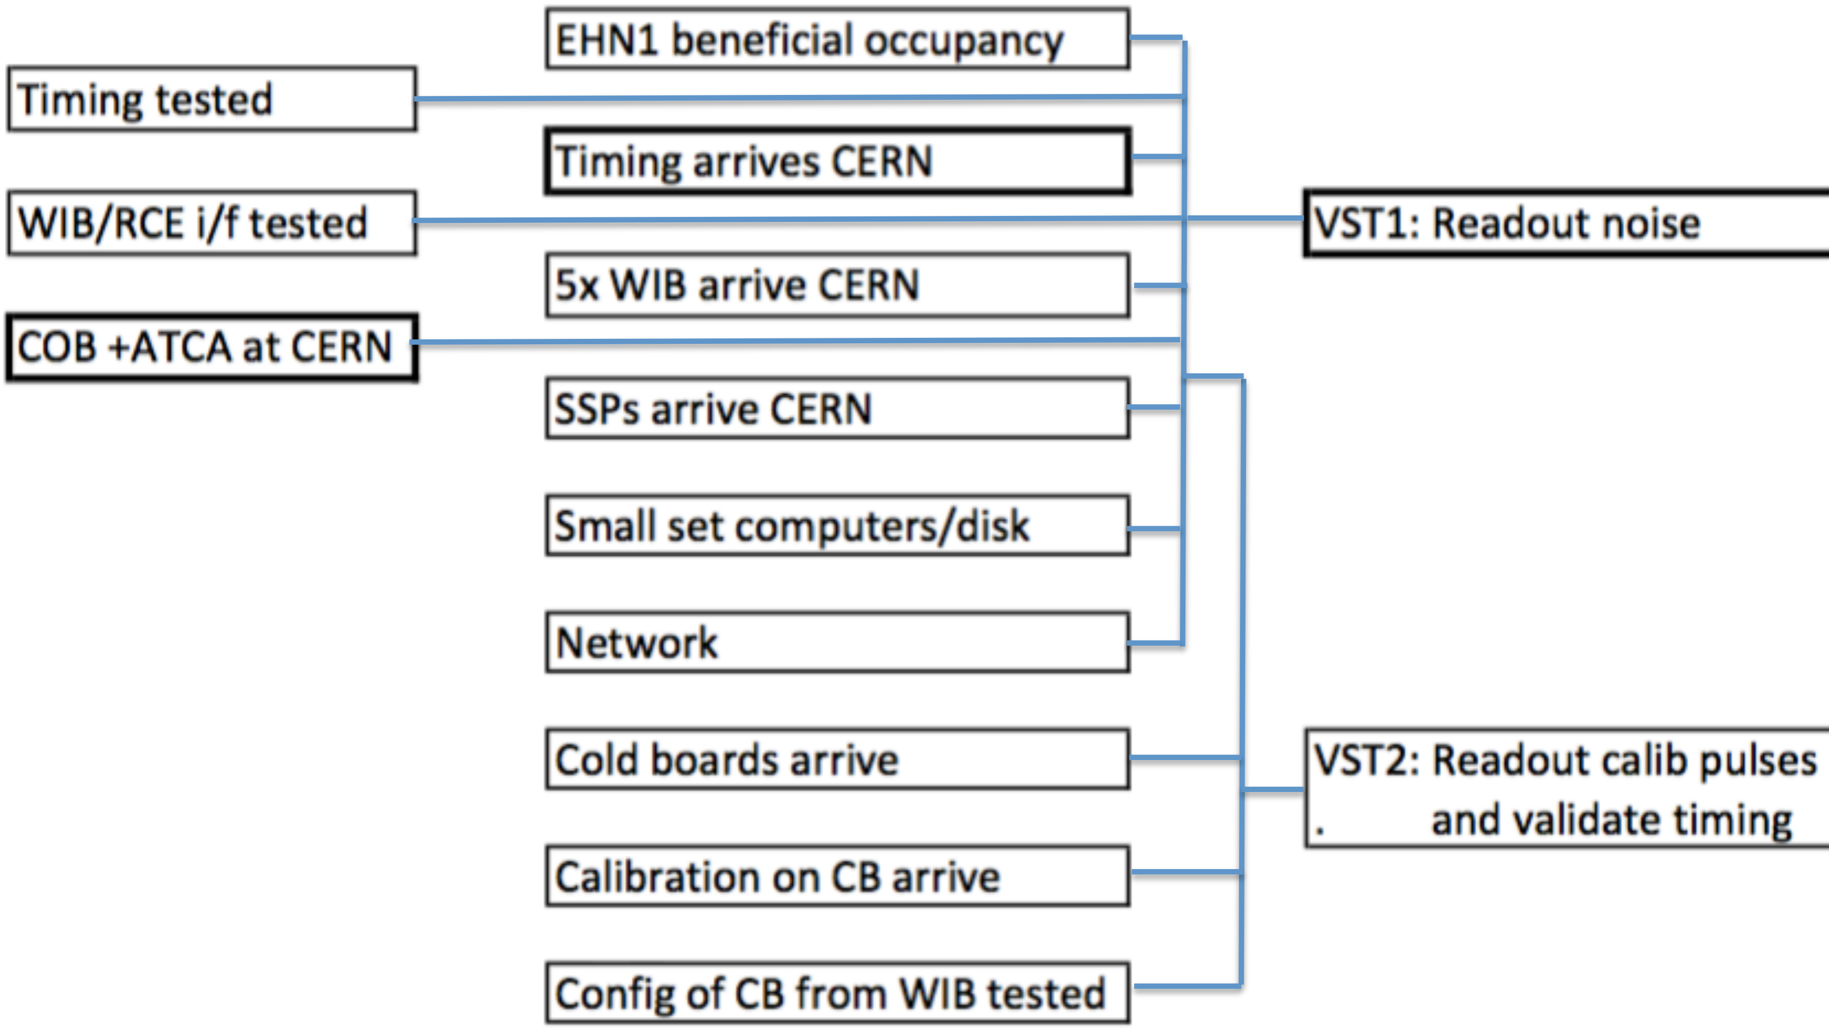

4. For exploitation talk (or after), would like to understand how expert functionality will be delivered on necessary timescale for experts

We have already implemented changes compared to the 35t to ensure experts are available much more and there are more of them

- Increase number of people who are 'almost full time' at CERN working on this. e.g. K. Hennessey, G Lehmann-Miotto, G. Savage, W. Ketchum. There are a lot of people (several representatives from each main component listed for Q1) signed up to come for periods of O(3months) in addition.
- Run control is strategic choice to be developed at CERN. This gives the onsite experts immediate access (and training opportunities) to interface to most of the other parts of the system; thus accumulating and sharing expertise effectively.
- Step though weekly meeting list giving expertise.

5. Timing diagram showing tolerances, also showing where timestamps get put in, how things get aligned

[See next slides]

Timing Alignment

- Step 1: Timing system provides 'absolute reference' across system
  - A phase-adjusted clock and timestamp, identical in each system
  - Triggers / calibration pulses marked with a single reference timestamp at source
  - (Blocks of) data samples are marked with a timestamp
  - Data blocks to board reader carry the trigger timestamp and event number in the header
- Step 2: For each detector, define a sampling window around trigger
  - Some data before trigger, and window sized to capture earliest / latest possible signals
    - e.g. for TPCs, slightly more than one drift period
  - Trigger latency is ~1μs; data path latency varies – detector-dependent trigger / data offset
    - e.g. data arrives before trigger in SSP, after trigger (due to compression stage) in RCEs
  - In SSP, the offset between trigger and data is fixed, compensated for on a per-board basis
  - In RCE and FELIX, the (compressed) blocks are timestamped, correlated with trigger
- Step 3: 'Fine alignment' done offline after timing calibration
  - May be time-dependent, calibration-dependent or have sub-sample precision
- Timing tolerance
  - Phase alignment should be a small fraction of the fastest sampling period
    - 150MHz sampling in the SSPs -> alignment precision of ~1ns; matches the effective precision of BI timestamps
  - Timing jitter should be a small fraction of alignment precision

Timestamp Application

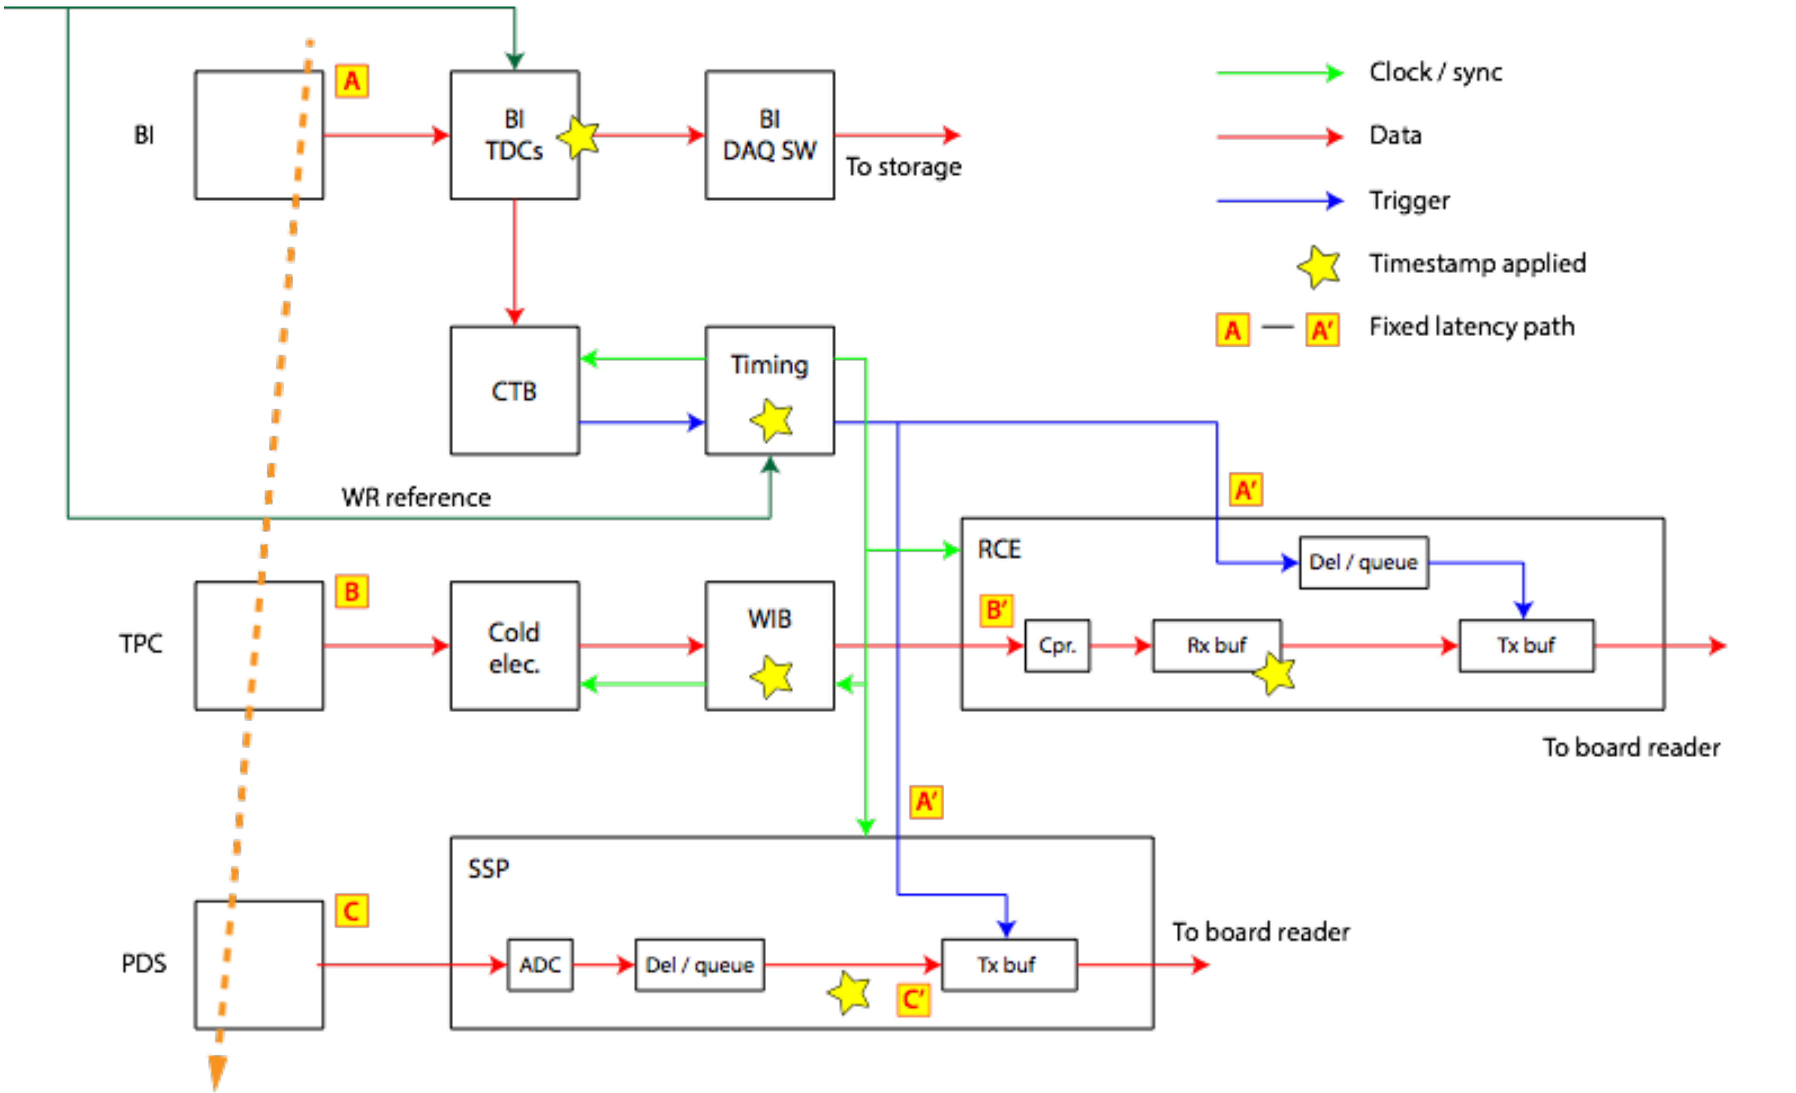

6. How much data can the system take from SSP for full waveform/continuous readout around a beam trigger

- 1.1Gbit/s is used to read out all the headers based on the rates we have been given.
- If we allocate a further 0.1Gbit/s to the photon detector readout, the size of the waveform that could be read on every channel is 6us.
- We consider 2Gbit/s from the SSPs a manageable amount of data that we can afford to collect in normal data taking mode.
- For special runs, much more bandwidth can be allocated to the photon detectors, so we should be able able to take data for much longer.5.

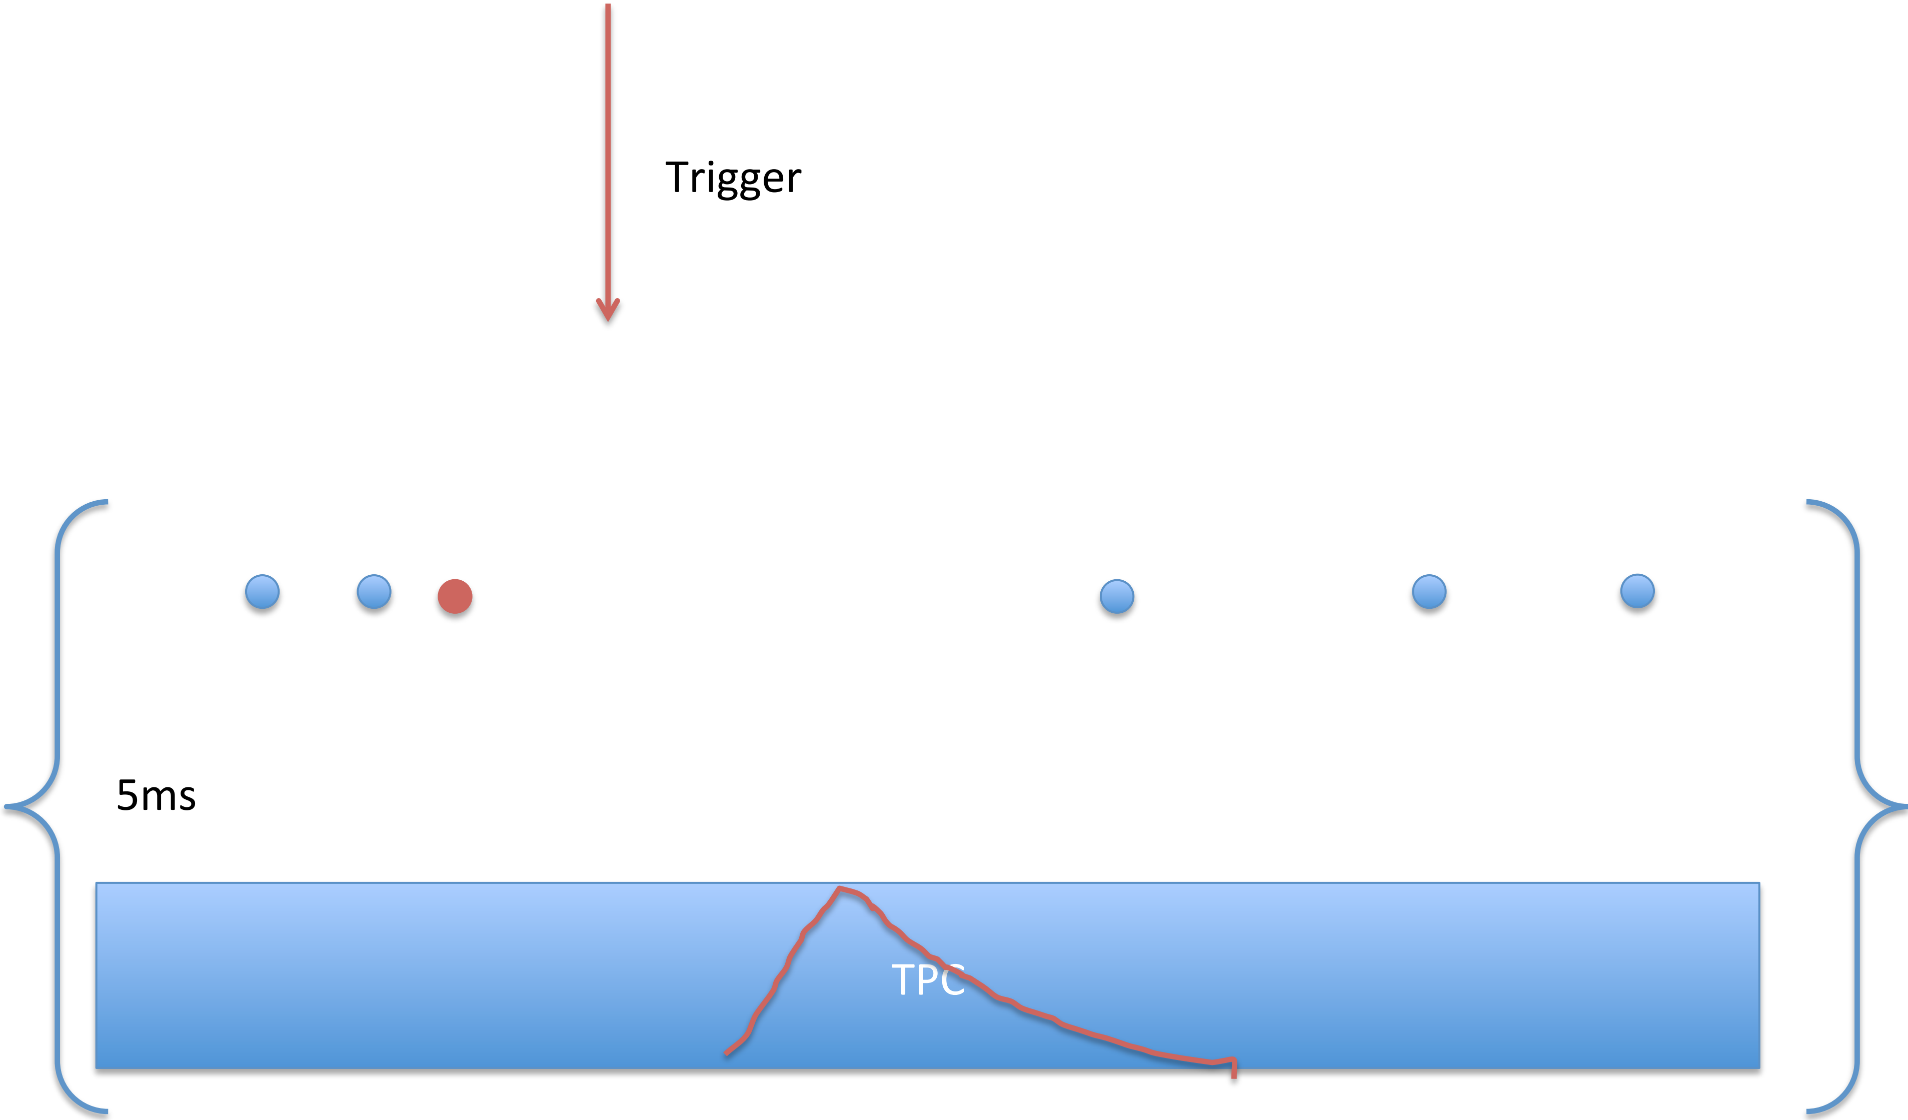
Trigger
5ms
TPC

7. Where and when exactly is the beam information merged into the data stream? Locally or at Tier-0? CRT data?

- Our initial answer is that this merging will wait until tier-0, because that is the simplest (which is a guiding principle for our DAQ design, see M.Thomson's talk)
- If it emerges that merging beam data is necessary to adequately do online data quality monitoring, we have several ideas for providing this locally (e.g.
  - Write beam data to a database and associate it keyed by event time at the point of reconstruction
  - BI information could be extracted like other conditions data during the processing stage, not implying necessarily a complete rewrite of data.
  - Write a parallel file for each spill

This is an area where new people can come in with good ideas over the next year, so it could be fixable more elegantly than indicated here.

Backup

Backup for question 1

Monday, October 10, 2016

10:00 - 10:45	ProtoDUNE-SP status	
	10:00	**RCE** *5'*
	10:05	**FELIX** *5'*
	10:10	**artDAQ** *5'*
	10:15	**Timing** *5'*
	10:20	**Trigger & Backpressure** *5'*
	10:25	**Run Control** *5'*
	10:30	**Online Monitoring** *5'*
	10:35	**Integration/Installation** *5'*
	10:40	**Interfaces** *5'* *- Front-End + WIB* *- Photon Detection System* *- Slow Control* *- Configuration Management + Databases* *- Online Computing* *- Beam Instrumentation*
10:45 - 11:05	Testing Centres	
	10:45	**UK** *5'*
	10:50	**FNAL** *5'*
	10:55	**CERN** *5'*
	11:00	**PNNL** *5'*

Communication and knowledge of leadership of each part is exchanged through weekly meetings status rundown.

We think that no steering group is needed, the weekly run-through of status among all collaborators is sufficient.

Upper level ProtoDUNE management has just gone through an evolution (past few weeks) and we are optimizing how we are plugged in to that new structure

Backup for question 5

- In ProtoDUNE, since the headline numbers (5ms drift, 2MHz digitization) are much slower than a normal detector, this is exceptionally easy.

- The latencies of the incoming trigger decisions and data arrival times are fixed and are very small compared to this.

- So alignment is done by
  1. Establishing a clean trigger from the beam (coincidence of two beam counters
  2. Collect data with TPC and photon system. Measure drift distance of start of track (we need to do this to about 1ms accuracy which is easy), use this to verify that track ends appear in correct location in detector and correct if necessary
  3. This also requires drift velocity measurement, obtained by cosmic tracks passing from cathode to anode plane
  4. By averaging over several events, look for accumulation of photon detector hits in the expected time bin in SSPs or nearby.

From CD1R: (About 1 year out of date)

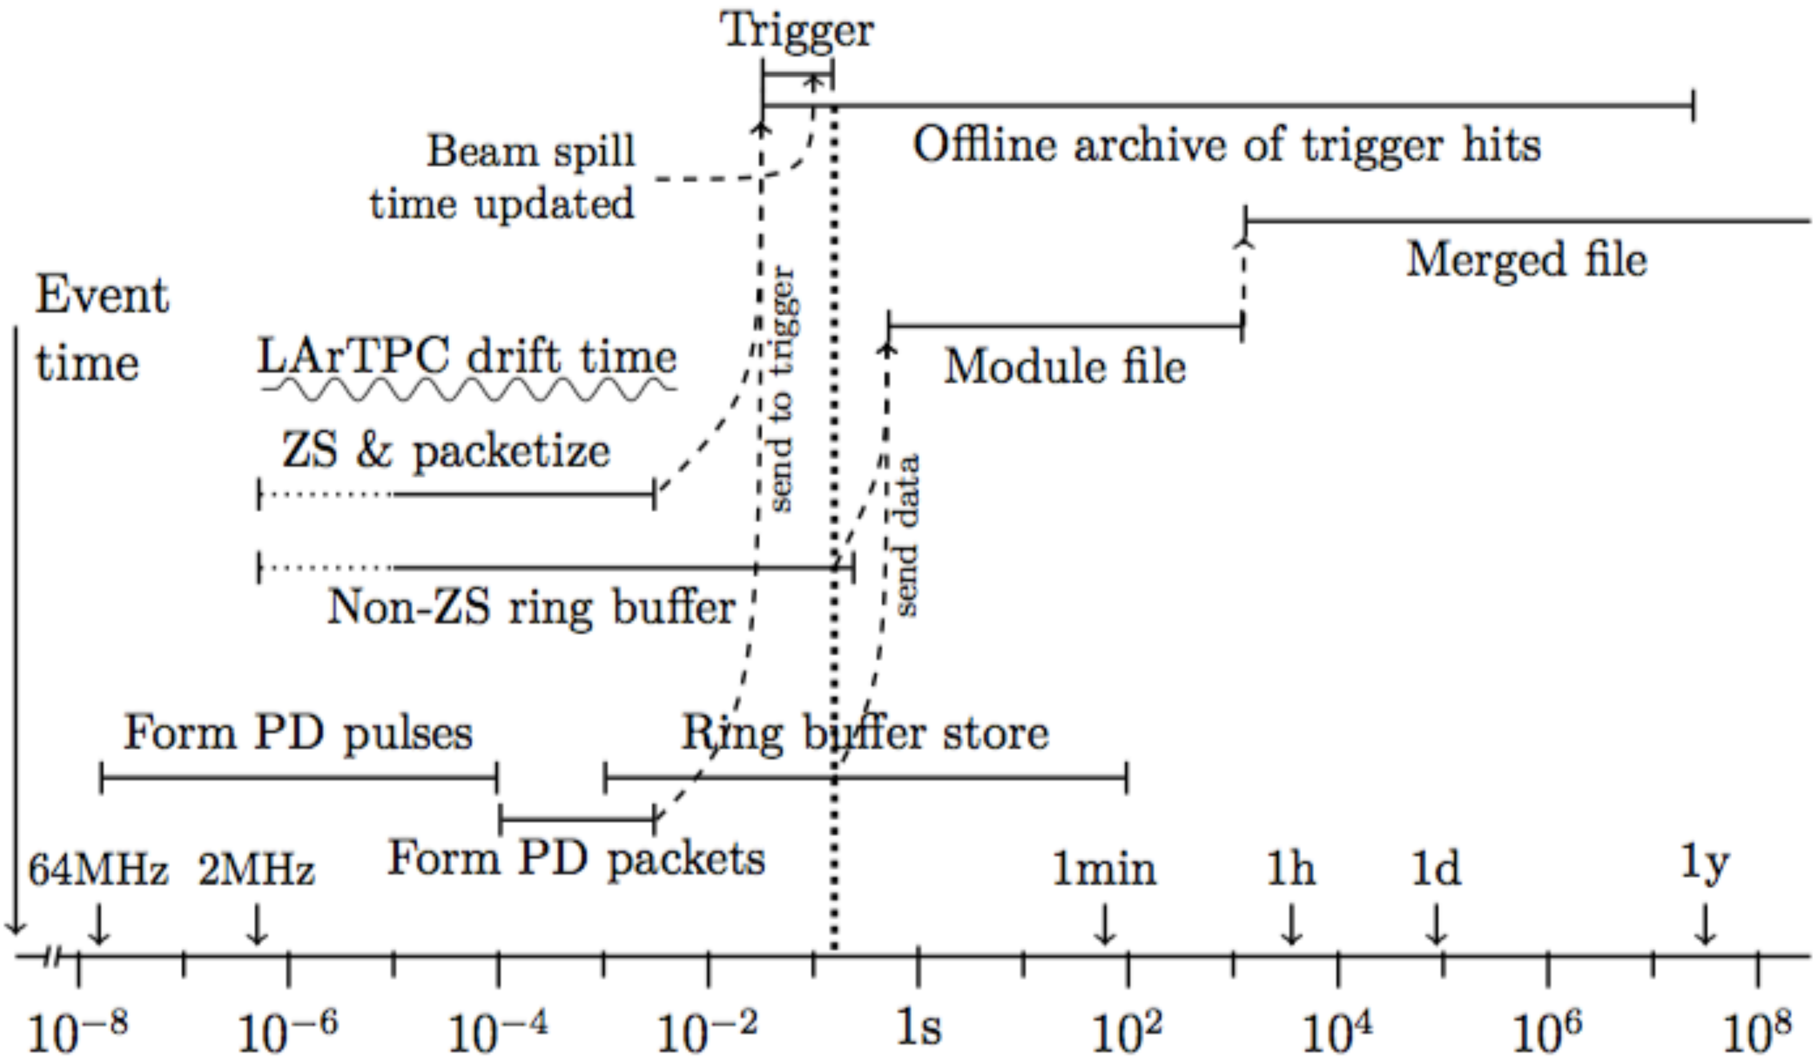

Figure 4.12: Main DAQ steps displayed relative to the event time.