REBAM Summary

January 11, 2019 – Watchman Collaboration Meeting

University of Pennsylvania

Outline

- Scope of REBAM
- Interfaces
- Point estimate of infrastructure and needs
- Open questions

Disclaimer

There are larger issues that impact the estimates and assumptions made here:

- AG control room
- Remote shifting
- Amount of storage UG

these will be approached at the end of the talk.

REBAM Scope

- Readout, Event Building, and Monitoring Working Group
- Interfaces:

  Upstream: DAQ Hardware, Tank , Site, Water, HV

  Downstream: Data flow, Offline data processing/analysis
- Ultimately REBAM is responsible for ensuring quality data taking for the physics analyses by connecting all subsystems upstream

Overview schema

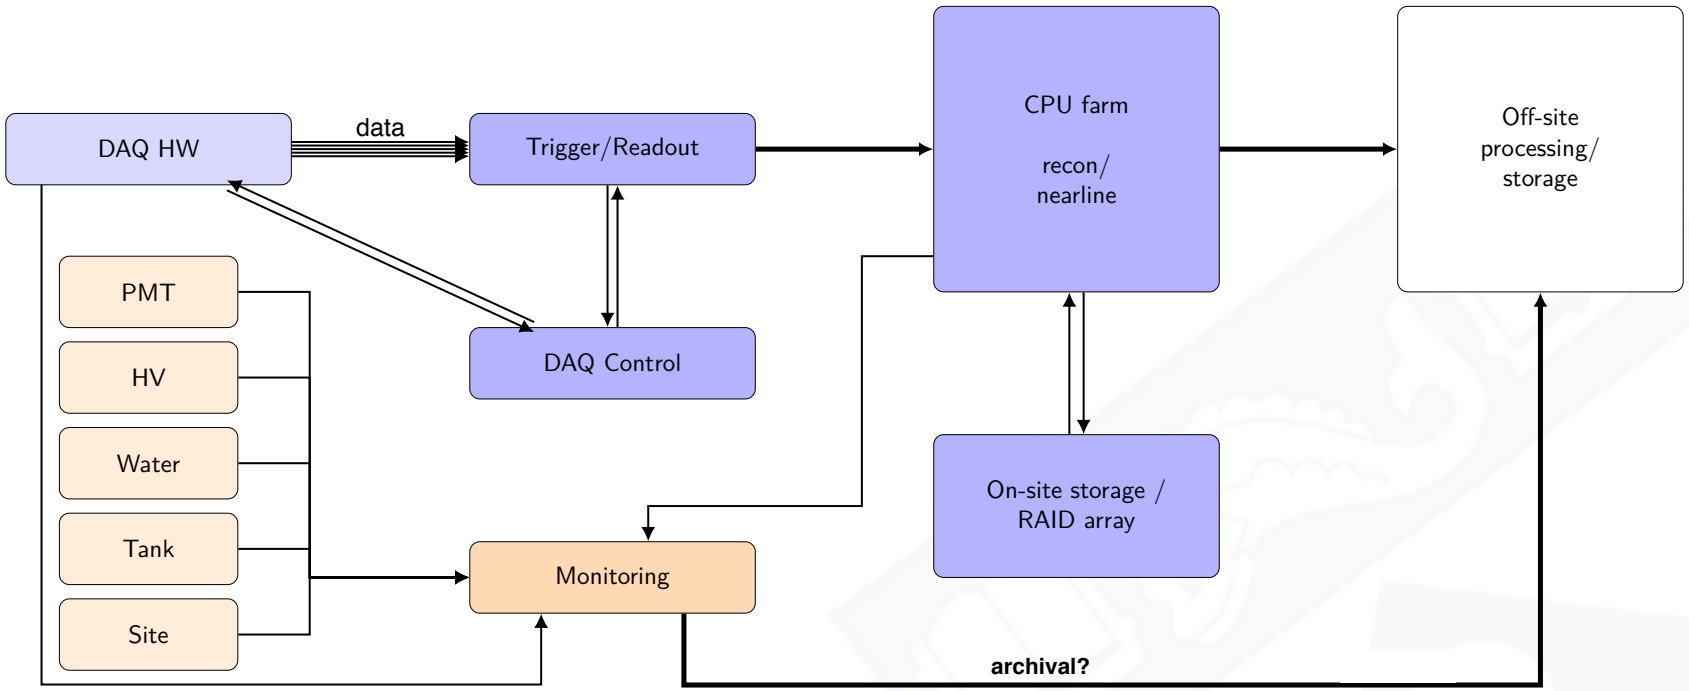

Readout

Readout

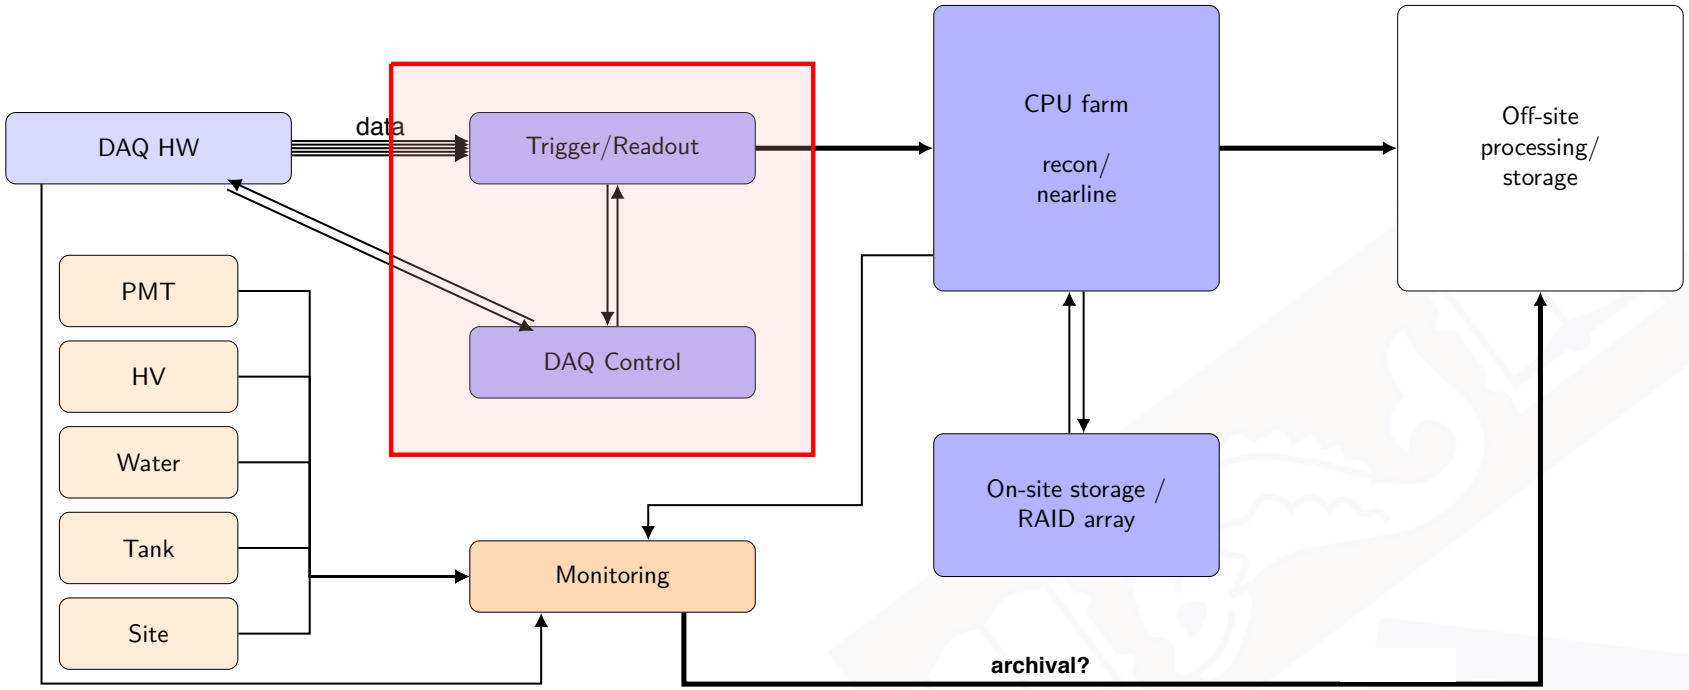
DAQ HW
data
Trigger/Readout
DAQ Control
CPU farm
recon/
nearline
Off-site
processing/
storage
PMT
HV
Water
Tank
Site
Monitoring
On-site storage /
RAID array
archival?

Readout Interface with DAQ

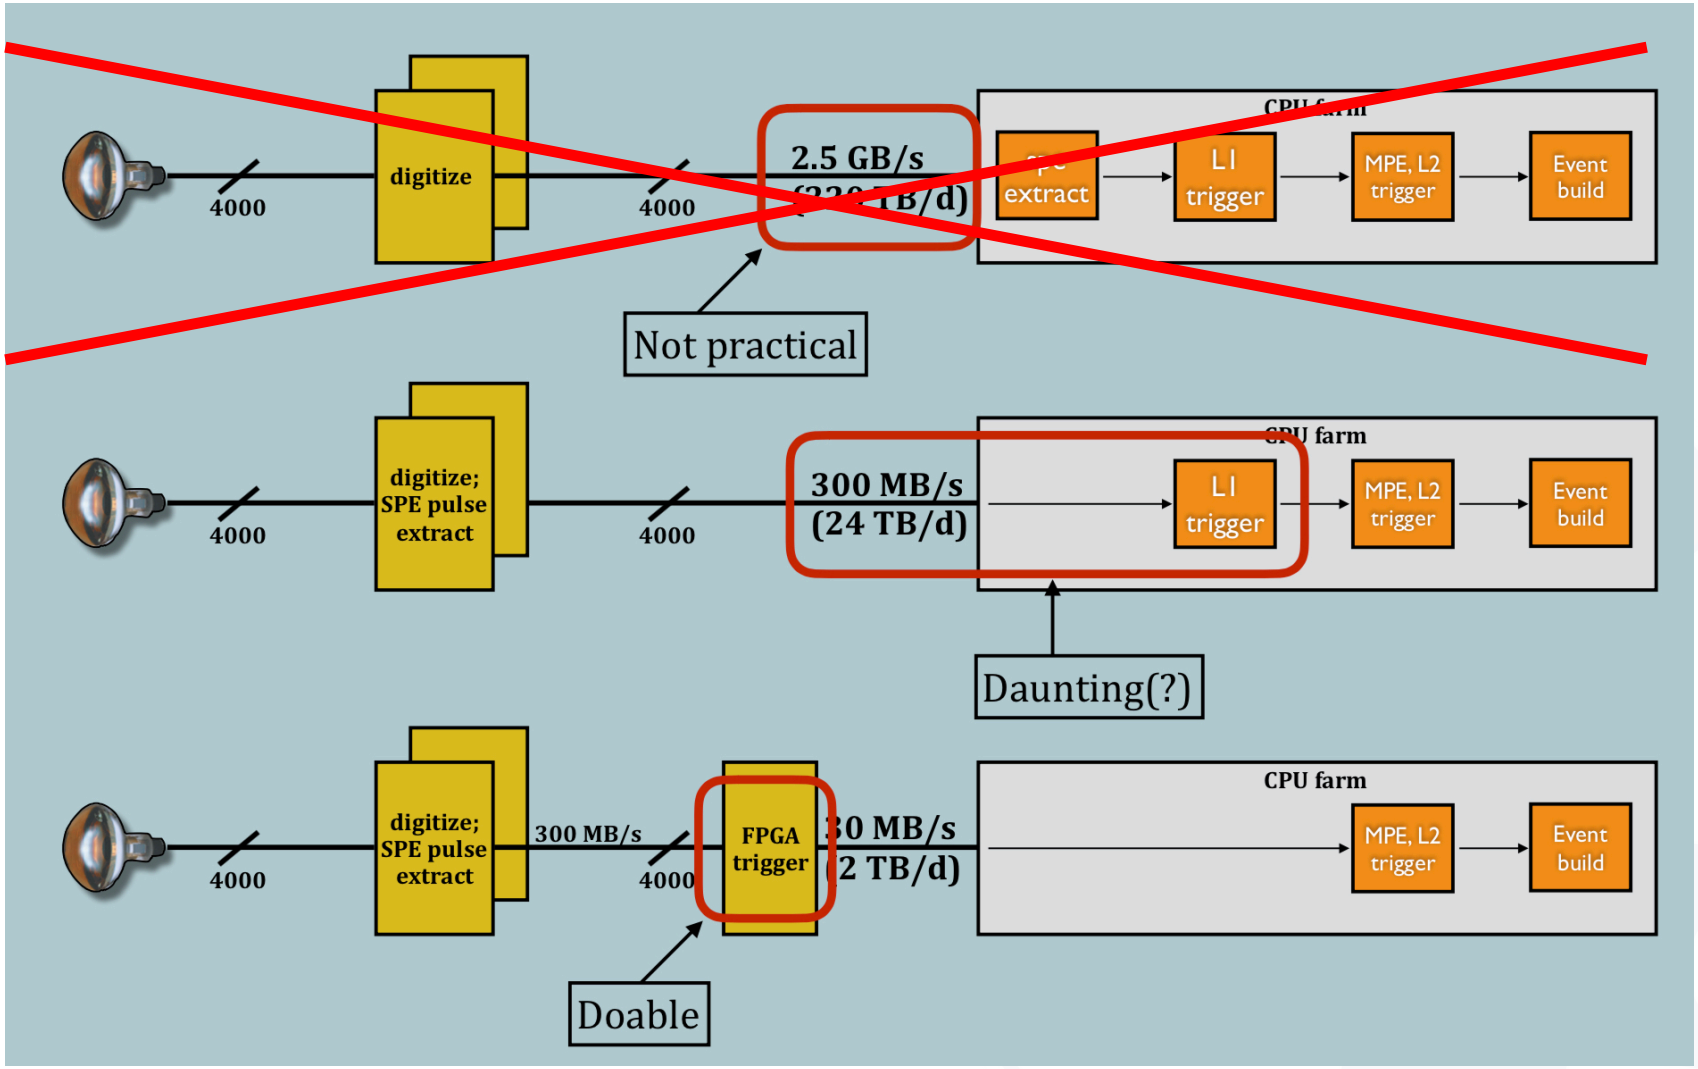

From Doug Cowen's pre-meeting talk

Comparison

- Similar requirements in terms of readout
- Option A : No/Limited hardware event trigger
  - Receiving (Q,T) for SPE, waveforms for events possibly non-SPE
  - Data rate estimated $\sim$ 300 MB/s (24 TB/d)
  - Trigger in software (integrated or not with event reconstruction)
  - Heavier processing requirements
  - Do we still want to store the waveforms?
    - Reduce final storage space by compression
- Option B: Hardware trigger
  - Receiving (Q,T) for SPE + rise time for events possibly non-SPE
  - Data rate estimated $\sim$ 30 MB/s (2 TB/d)
  - Lower storage and processing requirements, offset by cost of trigger

DAQ Control

- Hardware:
  - 2 operator monitoring computers/event display
  - 1 DAQ control computer
  - Monitors to go along
- Software:
  - Various options available (relatively similar in terms of requirements).
  - Several people familiar with ToolDAQ
  - One of hardware options has/will soon have interface with ToolDAQ
- A few questions:
  - How many pipelines into the DAQ readout?
    - 1 per crate, single line (if hardware trigger)
    - Can restrict ethernet bandwidth/fiber vs copper

Trigger

- Hardware:
  - Reduces pressure on software/computing
  - Unsure if completely feasible on hardware
  - DAQ hardware responsibility?
  - Could be:
    - Talk-back (take only primitives and pass decision back to hardware of what to ship)
    - Push-forward (receive everything and decide if should be stored or dropped)
- Software:
  - Receive everything and process along nearline reconstruction
  - Need to ensure enough RAM/SSD buffer to process pipeline (further discussed later on event building)
  - More flexible on implementation, but can only be push-forward

(Rough) Hardware cost estimate

Item	Description	#	Cost (unit)
DAQ operator + monitoring	Xeon E5-1620v2, 8GB RAM, 1TB HDD, decent graphics card	3	$2,000
Monitors	Seiki 39 inch	3	$600
		Total	$7,800

- Should account for 1 FTE (expert engineer or technician) over ∼ 2 years
- Plus participation from GD/PD

Event Building

Readout

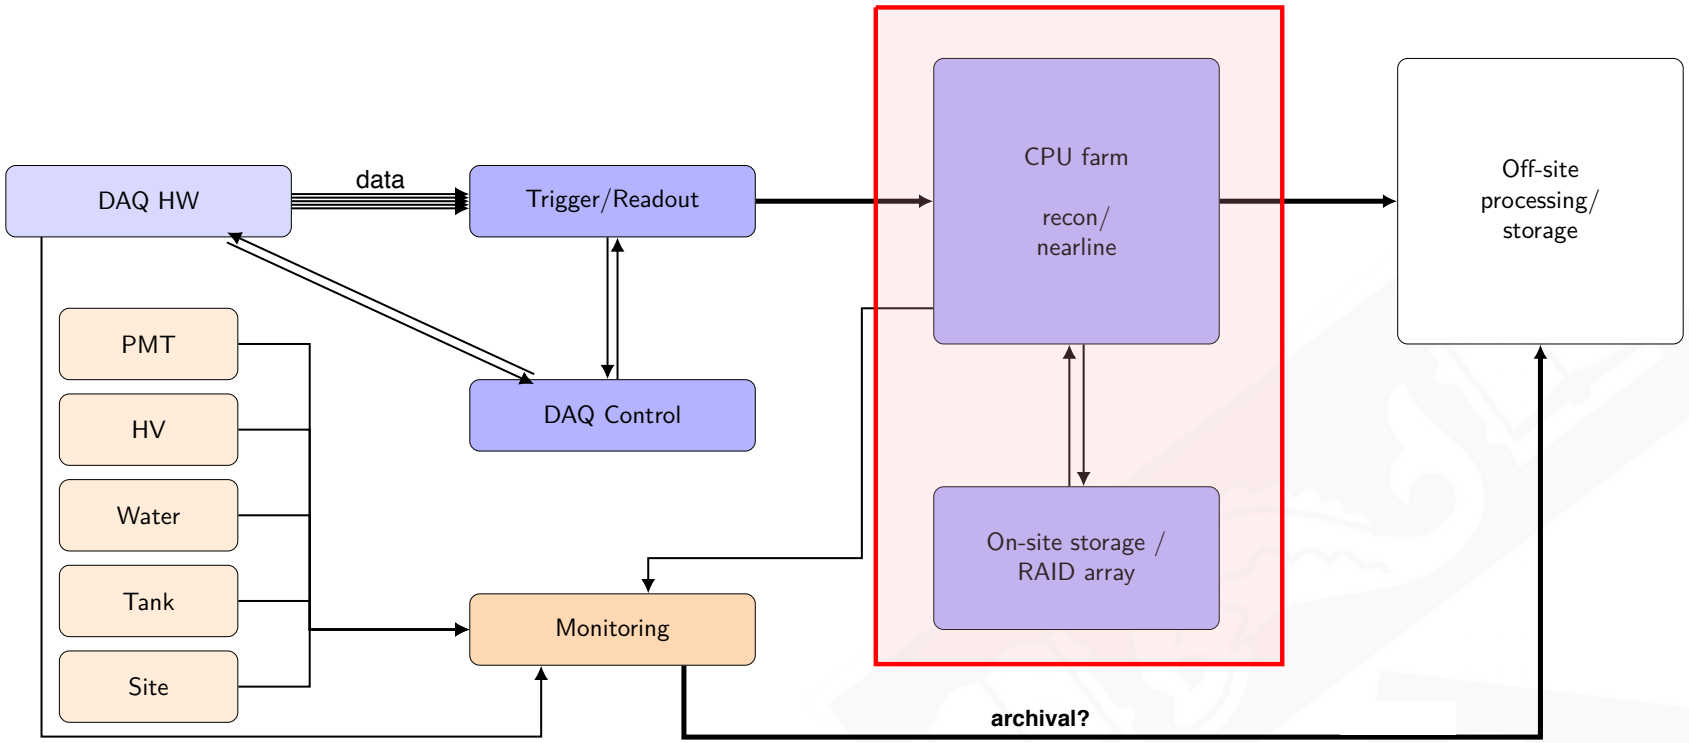
DAQ HW
data
Trigger/Readout
CPU farm
recon/
nearline
Off-site
processing/
storage
PMT
HV
Water
Tank
Site
DAQ Control
On-site storage /
RAID array
Monitoring
archival?

Overview

- We want to run some level of reconstruction near the detector
- Discussions from Michael Smy about a possibility using a SK style nearline reconstruction/trigger
- Processing requirements similar for both scenarios (with or without hardware trigger). Difference in buffer storage requirements.
  - With pulse shapes, 300MB/s (24 TB/day) require extra underground storage, even with compression.
  - With hardware trigger, can loosen storage but offset with hardware construction.
- Question to site WG: What connectivity is available to outside the facility?

Straw man hardware design

- Modest processing farm, able to reconstruct muons (and followers)
- Using SK[1] as an indication:
  - $\sim$ 64 cores
  - $\sim$ 8 GB RAM per core
    - Investigate possibility to trade RAM by SSD
  - 10GB ethernet link to RAID array
- Overall storage on RAID array with 50 TB
  - **Assuming ability to keep data for $\sim$ 3 days in case of connectivity loss to surface**
  - Additional storage allows for longer UG isolated operation

[1] Based on discussion with Michael Smy

(Rough) Cost estimate

Item	Description	Quant.	Cost (each)
RAID array	64 TB (56 usable with Raid5 + spare)	1	$11,375
CPU Farm	PowerEdge 40 core, 128 GB SSD, 64 GB RAM	3	$4,821
Network Switch	12 port 10GBps SFP+, 8 port 1GBps	1	$2,000
		Total	$27,838

- Need up to date quotes (oldest is 3 year old)
- Compiling alternative scenarios based on decisions/requirements
- Assuming 2 FTE $\times$ 2 year for this part (plus additional GD/PD)

Monitoring

Readout

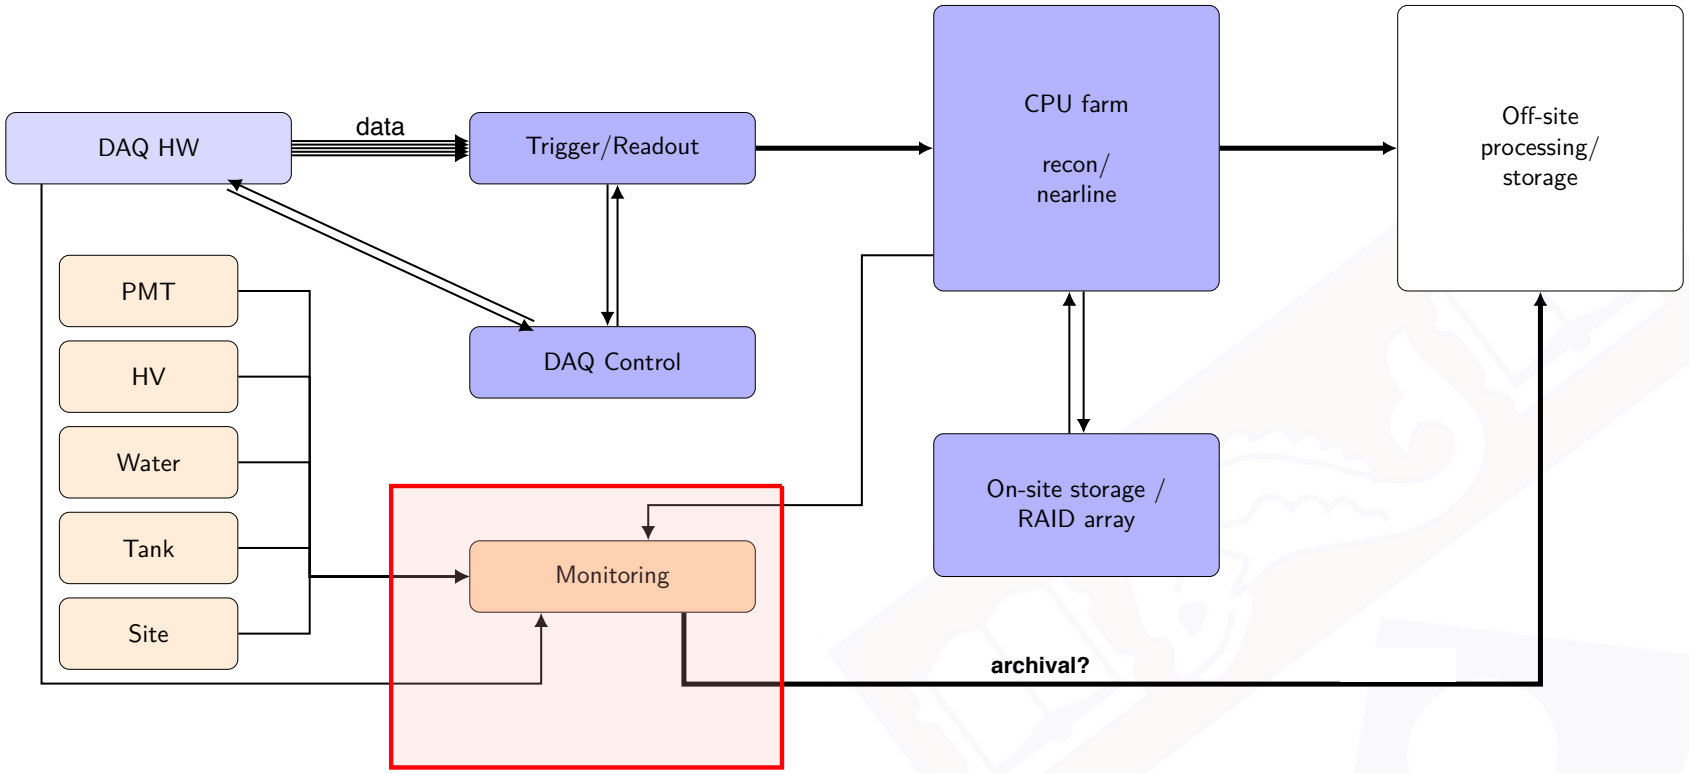
DAQ HW
data
Trigger/Readout
CPU farm
recon/
nearline
Off-site
processing/
storage
PMT
HV
Water
Tank
Site
DAQ Control
On-site storage /
RAID array
Monitoring
archival?

Monitoring

- Needs further discussion to define responsibilities and interfaces with relevant subsystems
- Two types of monitoring:

  Physics: Nearline data quality, Channel occupancy, Data flow monitoring

  Detector: Slow controls, computing infrastructure

  - We want to know if any hardware fails

Physics Monitoring

- Need to define metrics for quality of data being observed
- Inputs from different sources:
  - Directly from data stream:
    - PMT occupancy, trigger rate
  - From reconstruction
    - Spatial distribution of events

Detector monitoring/Slow controls

- DAQ hardware/PMT/HV
  - PMT voltage, current, baselines, Hardware state, deck temperature
  - Needs coordination with relevant subgroups to define formats and responsibilities
- Tank/site
  - Water temperature, resistivity, flow rate, deck temperature, any other relevant monitored quantities on site
  - Needs coordination with water, tank and site groups

(Rough) Cost estimate

Item	Description	Quant.	Cost (each)
Physics monitoring	2x Intel Xeon E5-2420v2, 32GB RAM, redundant PSU	1	$4,793.00
Slow Controls	Same as above	1	$4,793.00
		Total	$11,586

- Slow controls shouldn't require much in terms of processing power, but may have specific requirements for sensor readout
- Assuming 1 FTE $\times$ 2 year for setup and integration

Putting it all Together

Point Cost Estimate Summary

Item	Description	#	Cost (unit)
Readout/DAQ Control			
DAQ operator + monitoring	Xeon E5-1620v2, 8GB RAM, 1TB HDD, decent graphics card	3	$2,000
Monitors	Seiki 39 inch	3	$600
Event Builder/Reconstruction			
RAID array	64 TB (56 usable)	1	$11,375
Processing farm	40 core, 128 GB SSD, 64 GB RAM	3	$4,821
Network Switch	12 port 10GBps SFP+	1	$2,000
Monitoring			
Monitoring machines	2x Intel Xeon E5-2420v2, 32GB RAM, redundant PSU	2	$4,793.00

Total	$53,024
Total with 50% miscellaneous	$\sim$ $80000

Some notes on estimate

- `Miscellaneous` refers to parts not listed that we will need (cables, extra switches, extra machines, more storage)
- No contingency has been included
  - Need backup machines ready to take over in case of failure
- Each operator/monitoring machine requires 2U. Raid requires 4U.
  - Including redundancy/spare machines should in principle need 2 racks
- Haven't considered yet power/cooling requirements
- Cabling considerations between detector and control room not included

Manpower estimate summary

Number of FTE corresponds to experienced egineer/technician

- 1 FTE $\times$ 2yr for interface with DAQ control
- 2 FTE $\times$ 2yr for nearline reconstruction and trigger
- 1 FTE $\times$ 2yr for monitoring implementation/integration with other systems
- In total, estimate the equivalent of 8 FTE $\times$ yr over the span of the implementation of the whole system (3-4 years)
- Estimate effort being complemented by a team of 4 GS/PD $\times$ year

Other Considerations And Questions

Site and Data Flow

- Where does the purview of REBAM end?
  - At the fence of the site?
  - Where is data going to be stored long term? How will it be shipped there?
  - Where's offline processing done?
  - What requirements and implications into the design at site?
- Do we need an AG control room?
  - Run all shifts underground until remote shift CRs are available and validated?
  - For now assuming no AG control room
- Do we plan to set up remote shifting?
  - Should this be considered funded by operations or construction?
  - Should we plan this immediately along construction or focus on installation on site and set up remote shifting later?
  - Need to define requirements in terms of equipment, latency, security to site access

Software

- DAQ Control
  - A single proposal has been put forward: ToolDAQ
  - I am not very familiar with it personally, but it was highly recommended
  - Will be used in ANNIE, so there's value in reusing existing code (eg.: if K0TO DAQ HW is downselected)
  - Decision may have constraints from DAQ hardware solution pursued
- Nearline reconstruction
  - SK nearline reconstruction (Michael Smy)
  - ToolAnalysis (Matt Wettstein)
- Data format
  - Offline format should be the same (or a subset) of the MC format
  - Might be worth some effort in estimating if what we (RAT-PAC) have now is adequate
  - Online format will likely depend on the details of DAQ hardware

Storage and data shipping

- How long do we want to be able to store data UG?
  - Estimate based 50 TB.
  - Depending on details of data format/compression options should allow for $\sim$ 3 days
  - Do we need more? How long?
- What kind of connection is available UG/AG?
  - Is it shared with anyone else?
  - What is the effective bandwidth available?
  - Need to consider how much data we accumulate per day vs how much can we ship off-site

Yet More Questions/Requirements

Database Need one locally for reconstruction. What is stored? Which format?

- RAT-PAC uses CouchDB (in SNO+ we moved to PostgreSQL)
- Additional server and manpower needed

UPS Separate UPS for control room? Use detector UPS?

- Constraints on requirements from site, LLNL, collaboration
- Not included on the cost estimate yet.

Remote power control Gracefully shut down the detector if problem arises and no immediate access possible

- Automatic shutdown if UPS reaching end of charge

Installation Outsource? Draft from within the collaboration?

- Enough people familiar with this kind of infrastructure available?
- In SNO+ 5 grad students/postdocs did the whole thing in $\sim$ 2 weeks

Prototyping

- DAQ HW interface
  - Vertical slice test with DAQ hardware is the obvious moment to carry this out
  - Urgent item. Should be coordinated with DAQ hardware
- Monitoring
  - Test integration as systems and interfaces become available
- Estimate allocation of 1 FTE for this effort

Conclusion

- Determination of requirements being assembled
  - Have straw-man estimate based pre-meeting discussions
  - Many points still under discussion
- Manpower needed on several fronts
  - Slow controls
  - Reconstruction
  - DAQ software implementation/testing/prototyping
- Further discussion on interfaces and requirements needed
- Will be dedicating some time for REBAM on DAQ meetings
- Will start drafting text for CDR
- Feel free to contact me if have any questions, suggestions or offers for contributions