




# Data Acquisition with GPUs for Muon $g-2$ at Fermilab

Wes Gohn

University of Kentucky

2 August 2017






# Muon g-2 Experiment Overview



- Goal is to measure the anomalous magnetic moment of the muon to 140 ppb, which is a factor of 4 better than has been previously measured.
- Muon fills are injected into the ring at a rate of 12 Hz.
- The precession frequency of the muons is measured by detecting decay positrons in 24 segmented calorimeters inside the ring.
- We finished a five-week commissioning run on July 7.





# g-2 Detector Systems




- 24 Calorimeters
  - 1 uTCA crate for each calorimeter
  - 54 channels \* 24 calos = 1296 channels of digitized data.
  - Data provided by 12 Cornell waveform digitizers.
- 4 Fiber Harps
  - 7 channels \* 4 harps = 28 channels
  - Data provided by Cornell waveform digitizers
- Quads and Kickers
  - Write 4 quad channels and 15 kicker channels
  - Data provided by Cornell waveform digitizers
- 3 Trackers
  - Data from Multihit TDCs sent from FC7s in a uTCA crate
- IBMS and quads
  - Running on CAEN digitizers
- Slow control SCS3000 devices



# Rate requirements



- Accommodate 12 Hz average rate of muon fills that consist of sequences of eight successive 700  $\mu$ s fills with 10 ms fill-separations.



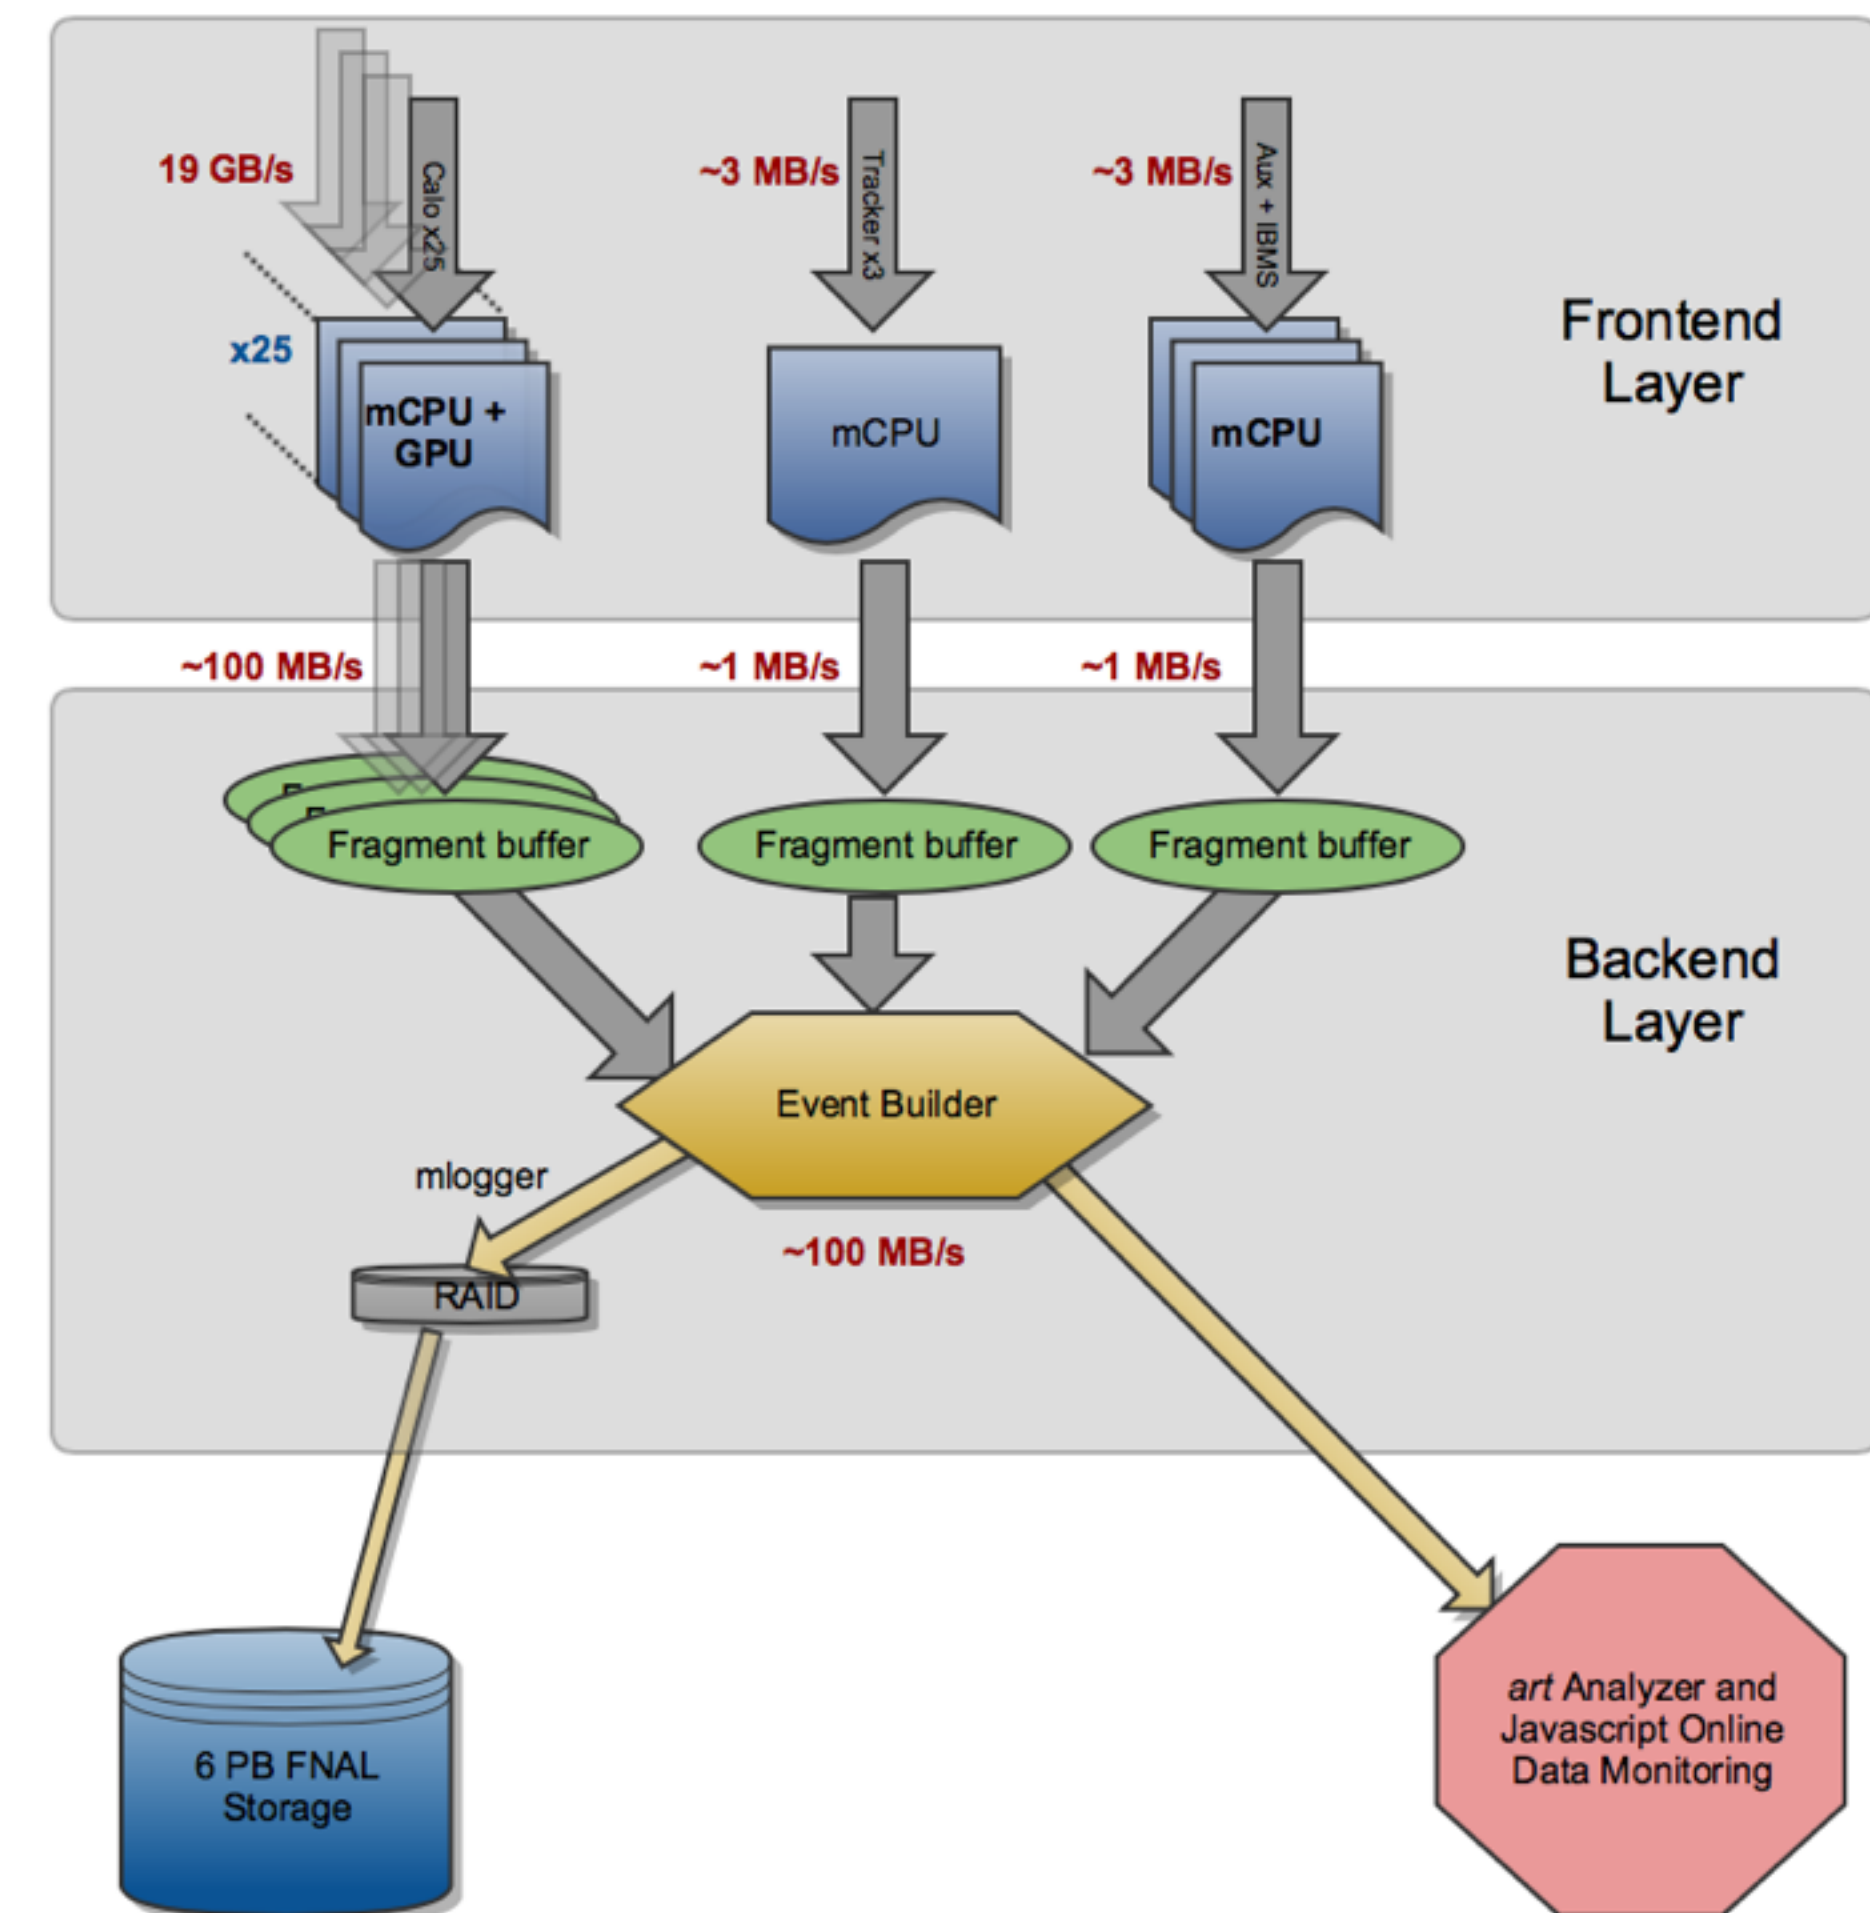
- Time-averaged rate of raw ADC samples is 20 GB/s, which must be reduced by a factor of 100.
- Data is processed in GPUs to accomplish this task.
- Total data on tape after 2 years of running will be 10 PB.

Source	MB Per Fill	MB Per Second
Raw data	1,600	19,400
T-Method	9.4	112.5
Q-Method	4.0	48.5
Prescaled Raw	1.6	19.4
Tracker	0.75	9
Laser Monitor	0.08	1
Auxiliary	0.33	4
<b>Event Builder:</b>	<b>16.2</b>	<b>194.4</b>



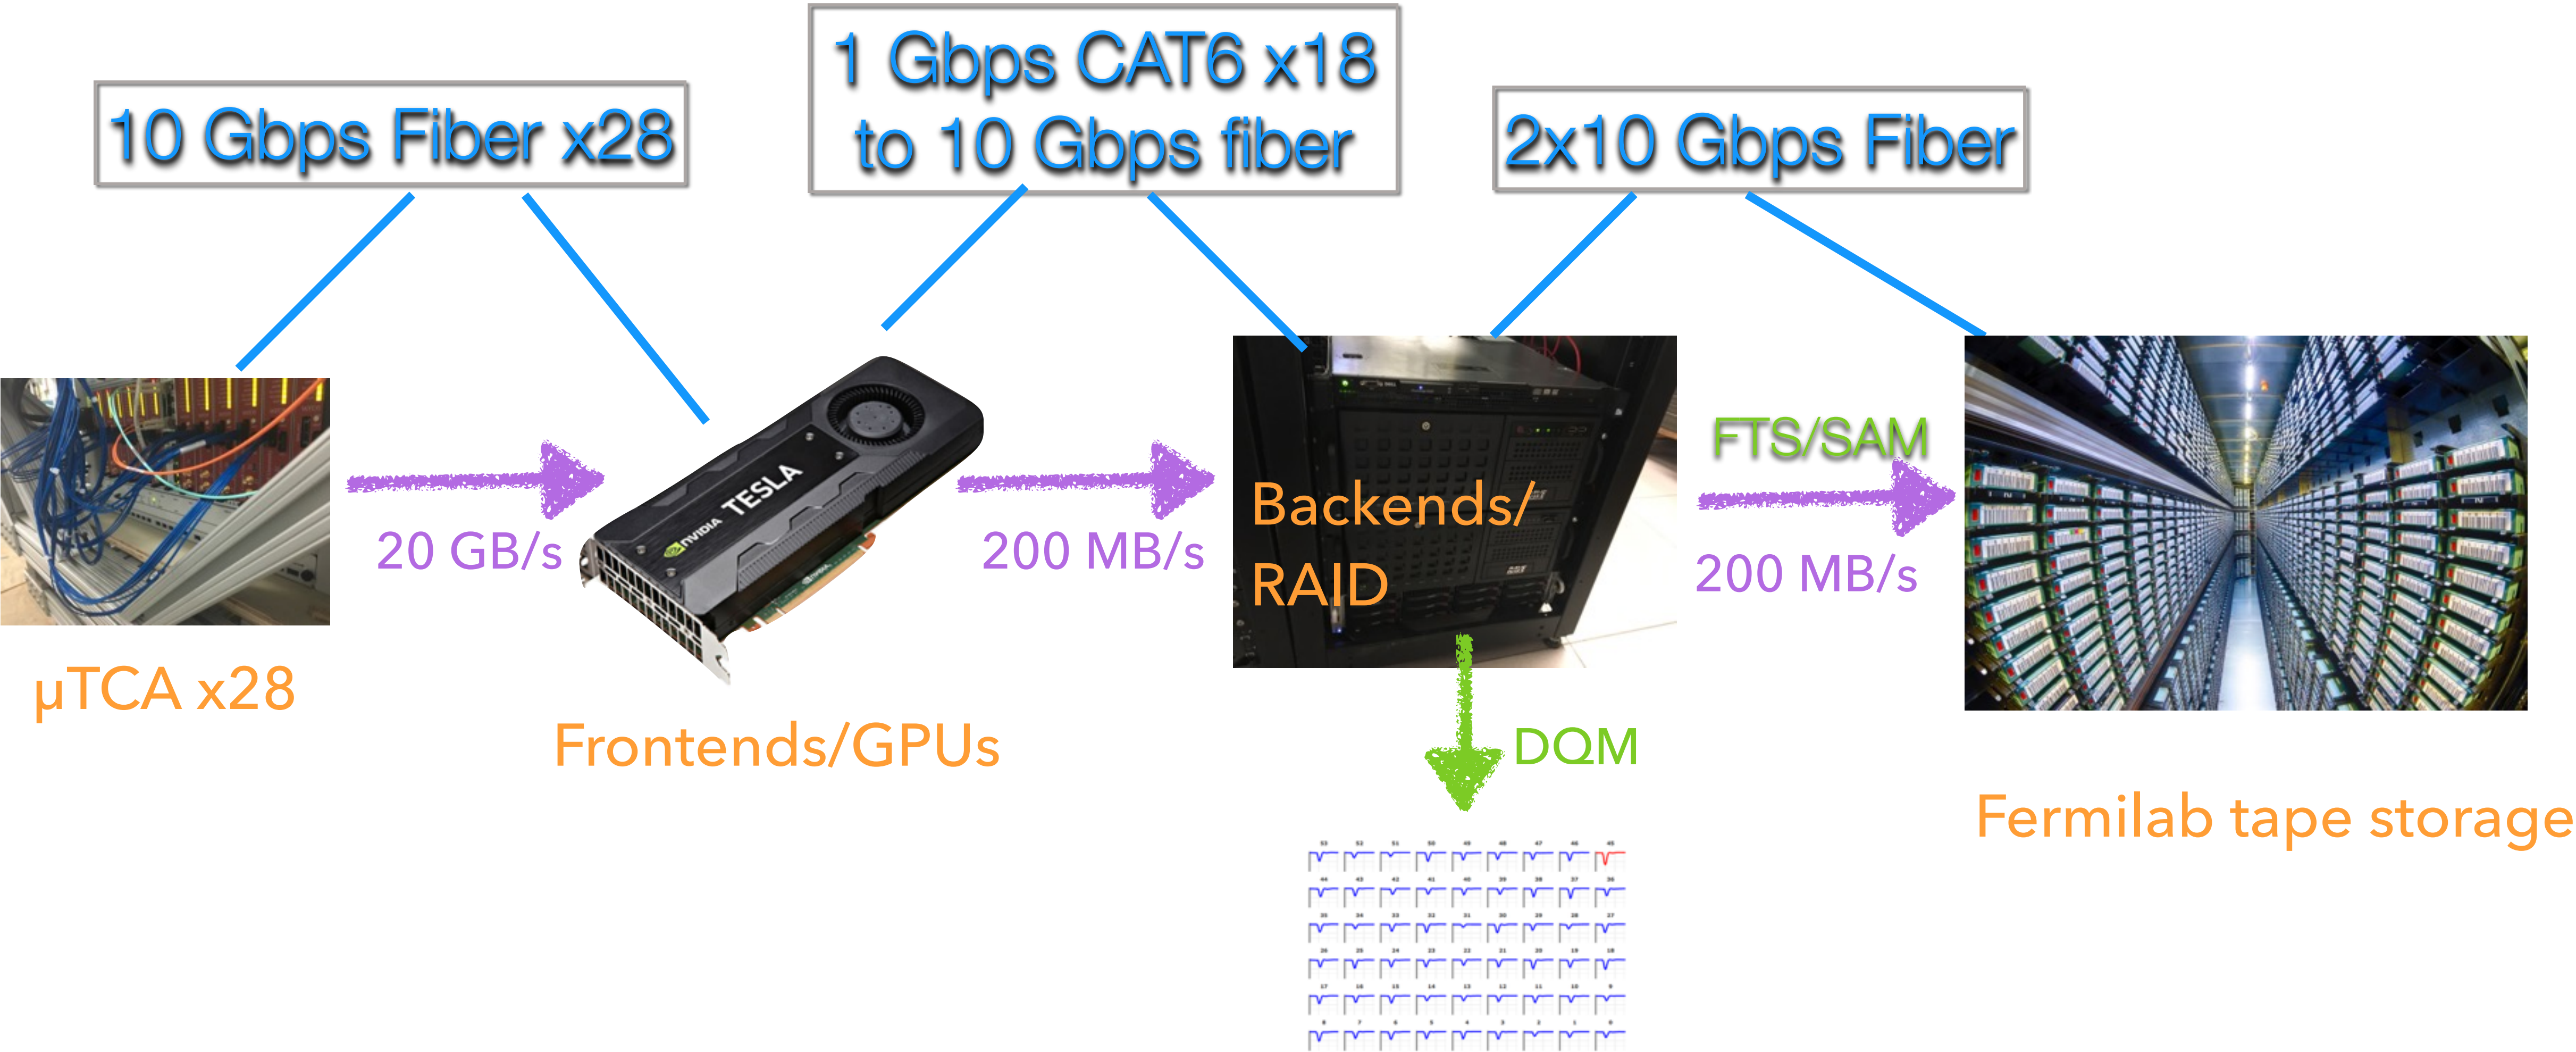
# DAQ Design

- Layered array of commodity, networked processors
- Frontend layer for readout of detectors.
- Backend layer for assembly of event fragments.
- Slow control layer.
- Online analysis layer using *art*+JS.
- Field DAQ operates independently, but with a similar design.





# Data Flow





# DAQ Hardware




- The DAQ hardware includes:
  - 17 frontend machines
  - 5 backend machines
  - 2 dedicated near line analysis machines
  - 3 computers for HV control
  - 3 servers
  - 24 beagle bones running slow controls
- Each frontend contains two Nvidia Tesla K40 GPUs
  - 2880 CUDA cores at 740 MHz
  - 288 GB/s memory bandwidth
  - 12 GB on board memory
  - ECC memory protection
- 70 TB RAID for temporary data storage.

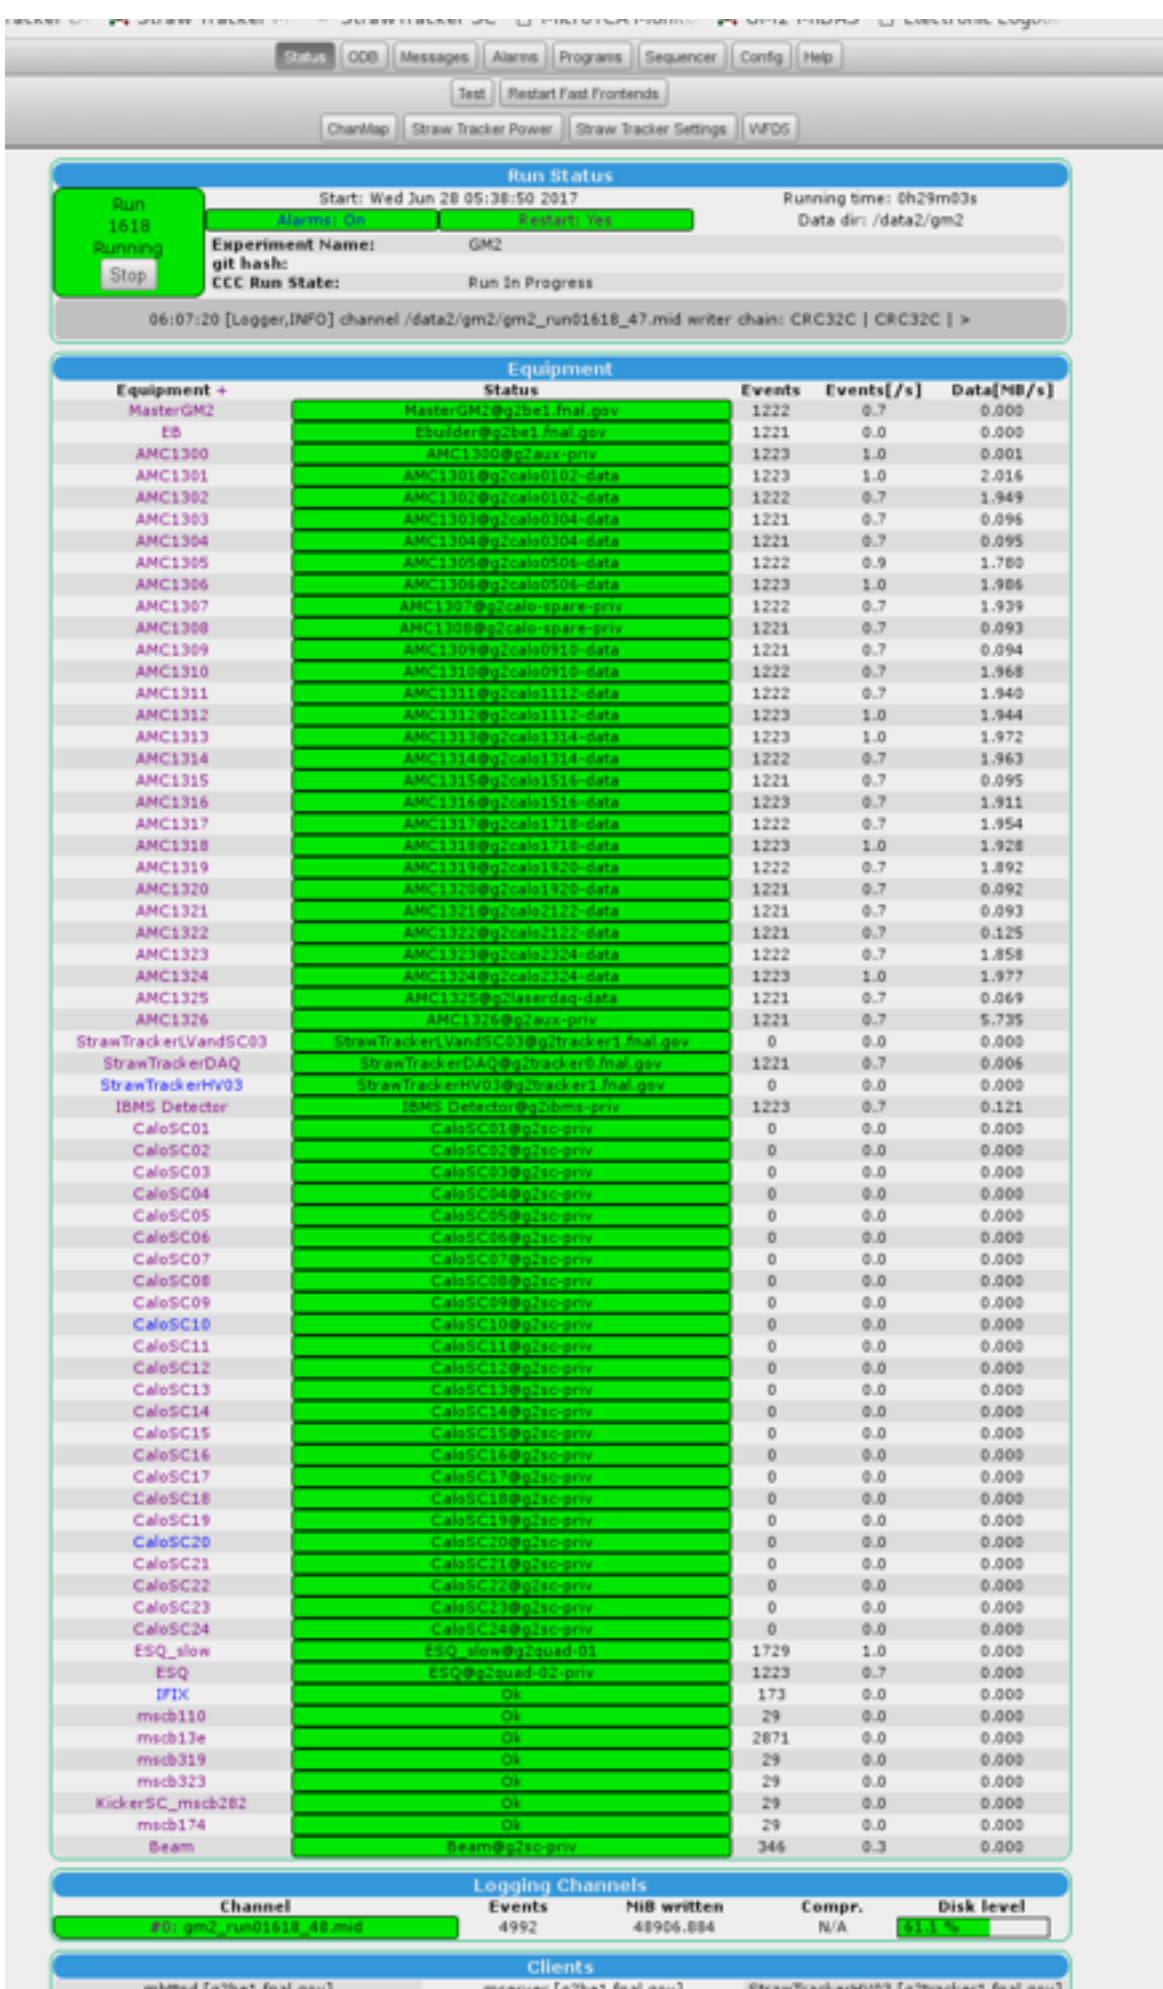




# MIDAS Software



- MIDAS is a DAQ software package developed at PSI and TRIUMF and used by T2K, SCDMS, MEG, and many other experiments.
- Provides a web-interface for control of the experiment, data logging, and event building.
- We write frontend code in C++ and CUDA that processes the data and sends it to the event builder.
- 32 fast frontends reduce data volume for each event by a factor of 100 and store data in Midas banks.
- 35 slow control frontends process slow data.
- Includes alarm system and sequencer.
- Software configuration is dumped to a JSON file and saved to a PostgreSQL database at the end of each run.


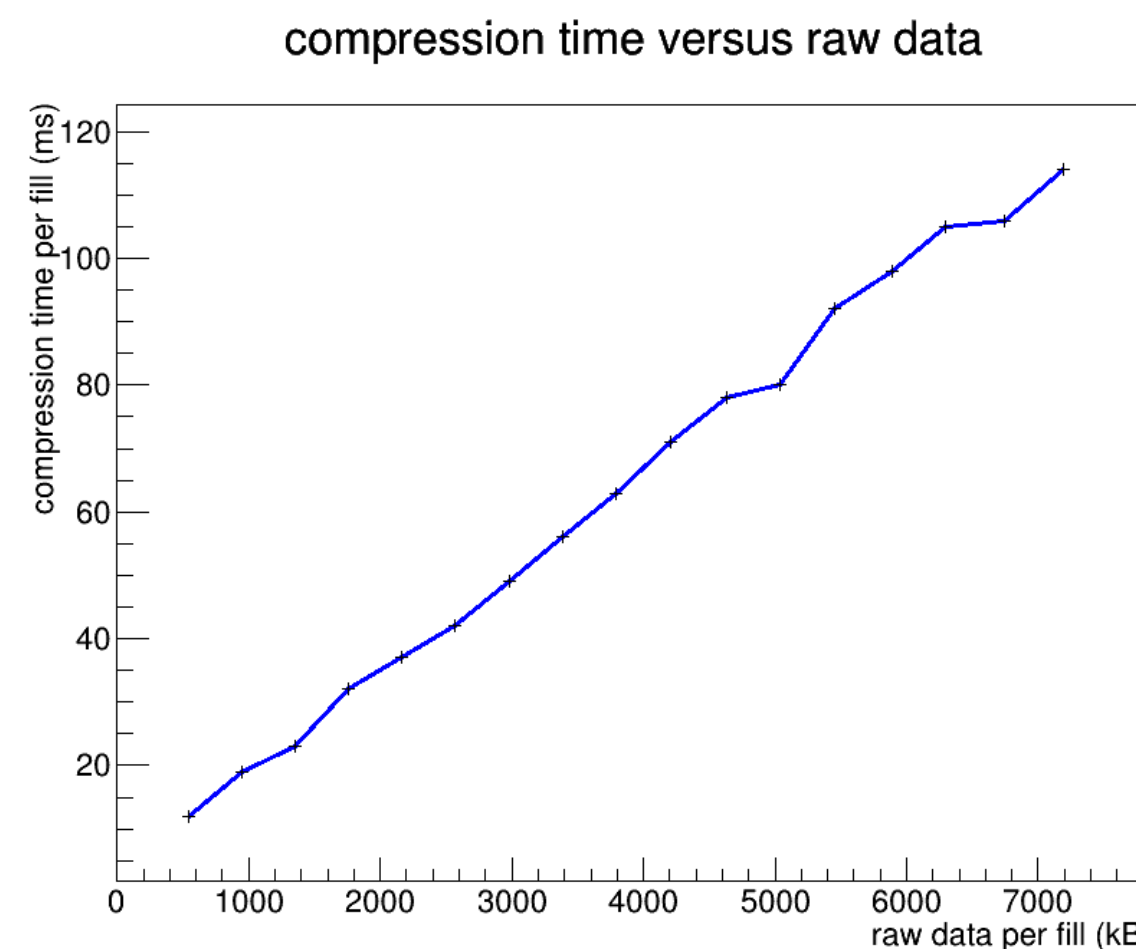
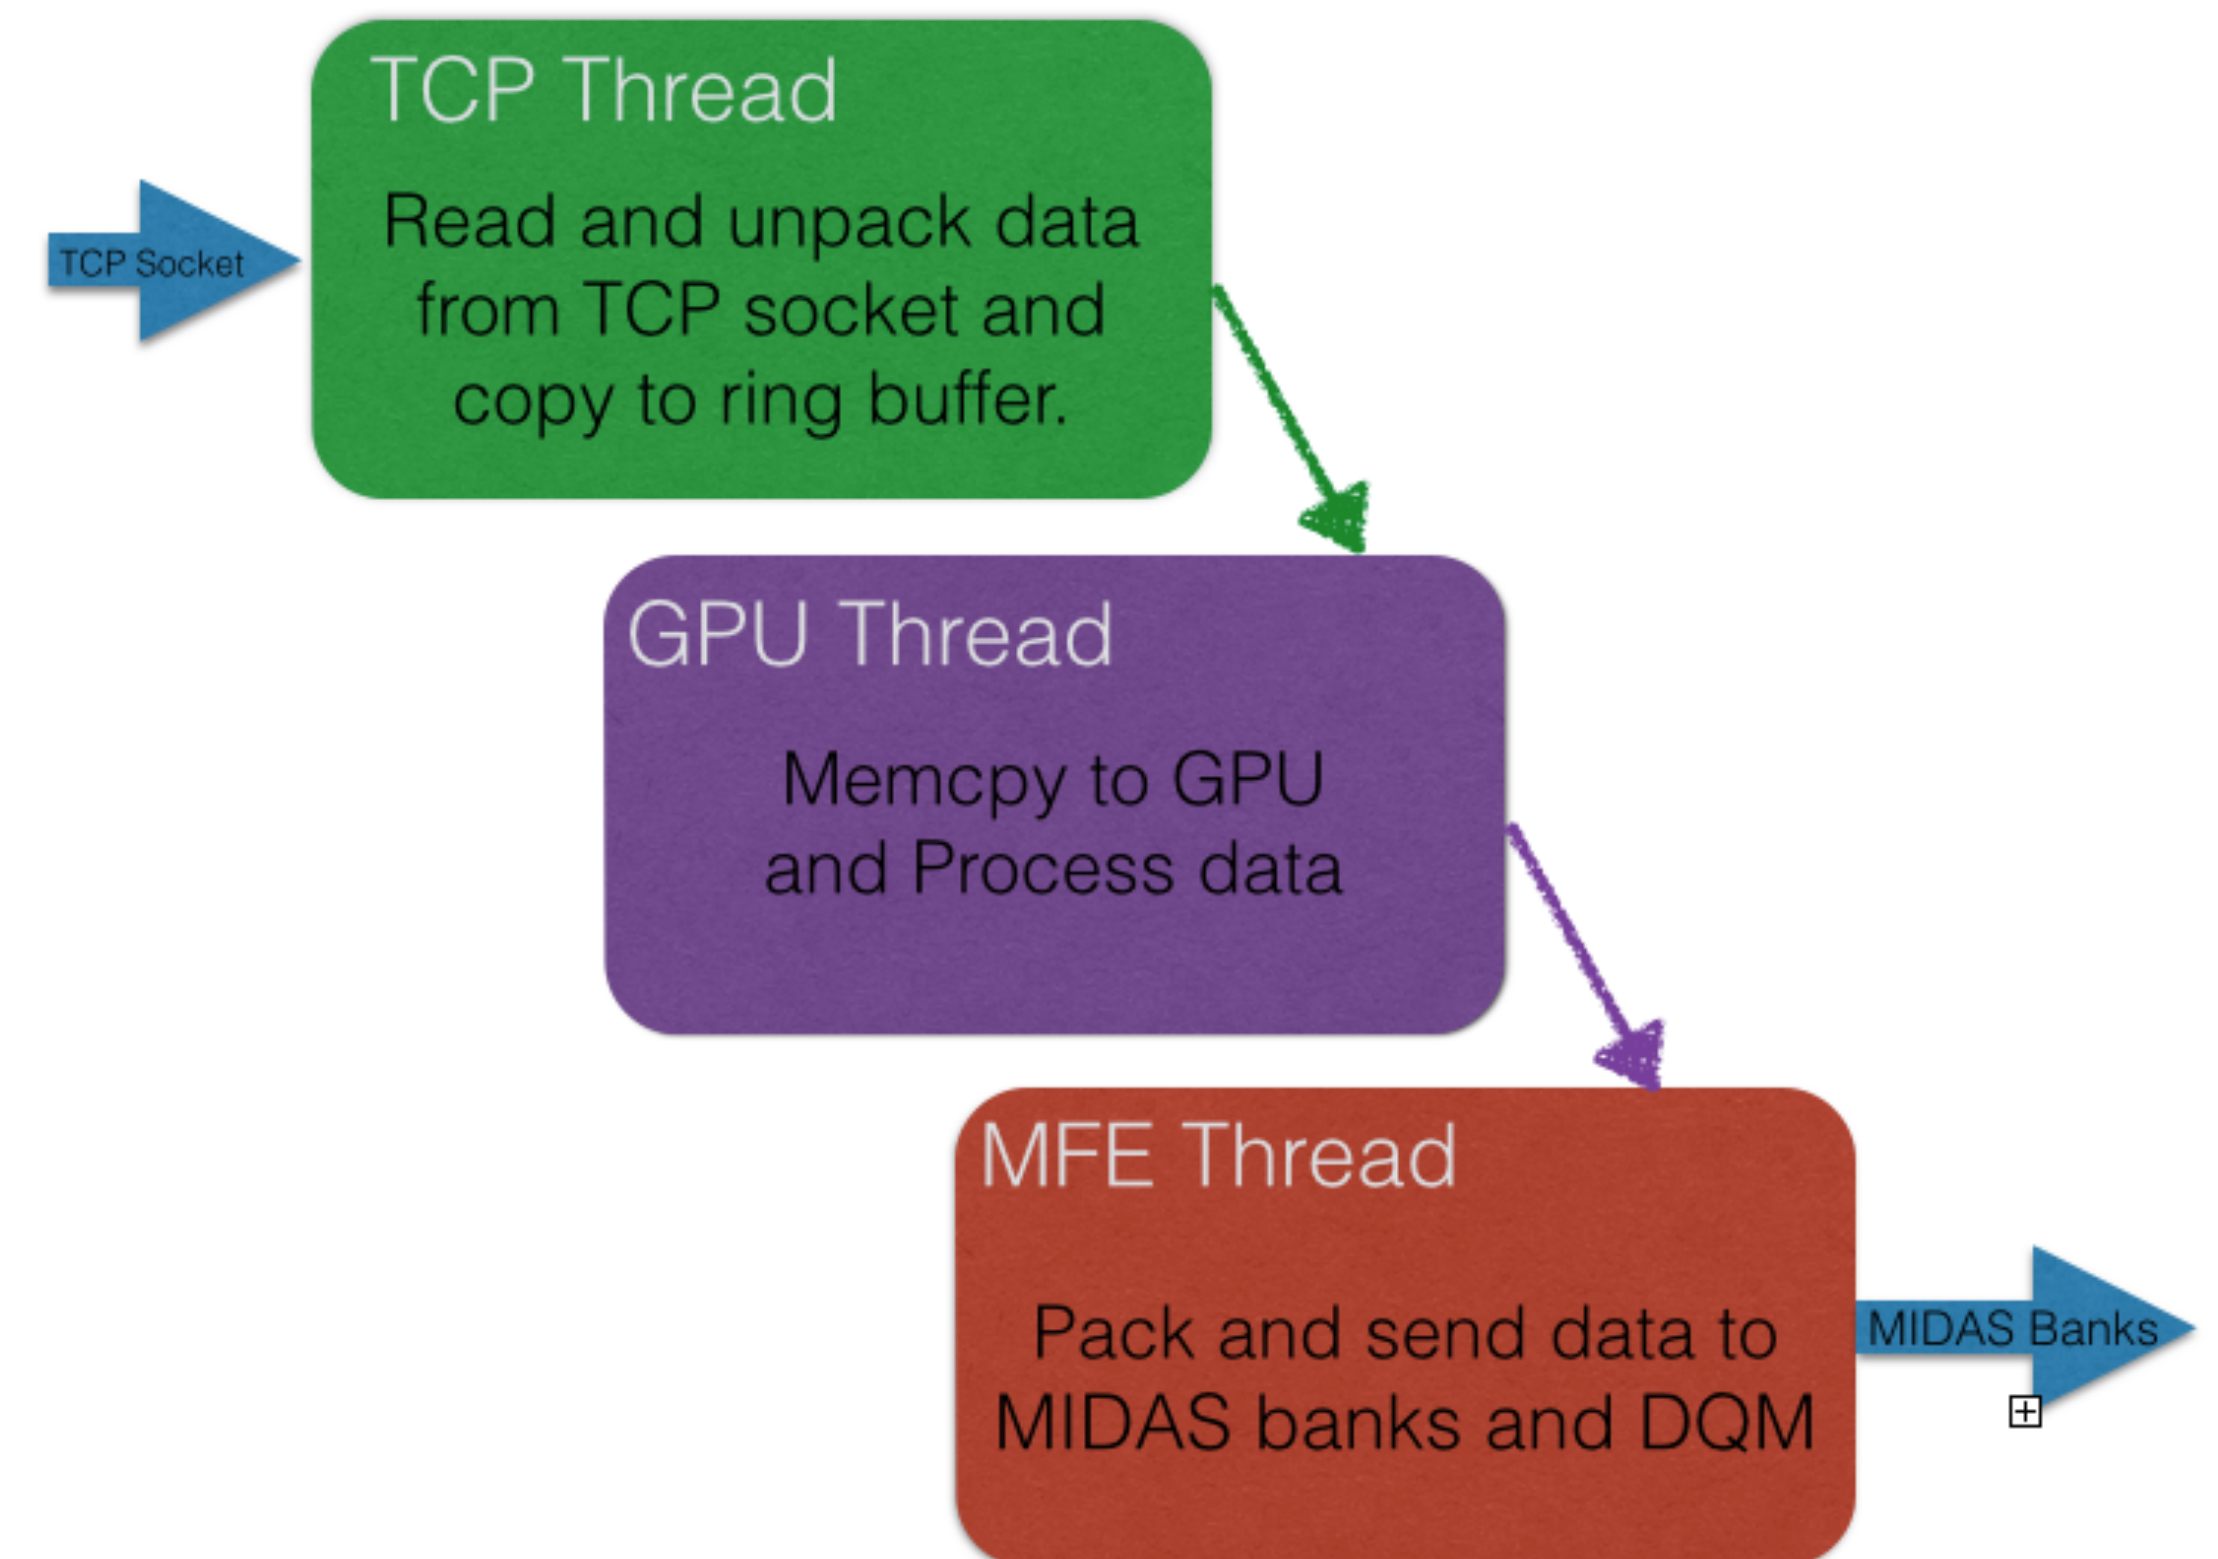




# Calorimeter-GPU Frontend




- Each frontend process reads data from one uTCA crate over 10 Gb ethernet with TCPIP.
- Frontend is multithreaded with mutex locks.
- Data is processed in Nvidia Tesla K40 GPUs using CUDA code that is integrated into the frontend.
- Midas banks are losslessly compressed using zlib.
- Full configuration of the uTCA crate is performed via the MIDAS ODB.





# Why GPUs?

- The GPUs dramatically improve performance by parallelizing processing.
- Technology was developed for commercial applications, so it is well supported.
- Easier to code in C++ than to learn specialized language.
- Without GPUs, we could not keep up with our data rates.






# GPU Processing



- The frontend includes CUDA routines for data processing.
- Each GPU processes data from one calorimeter.
- Raw fill is copied to GPU memory, where it is reduced using T-method (island chopping), Q-method (histogramming), pedestal calculation, and template fitting.
- The output of each process is written in one MIDAS bank.

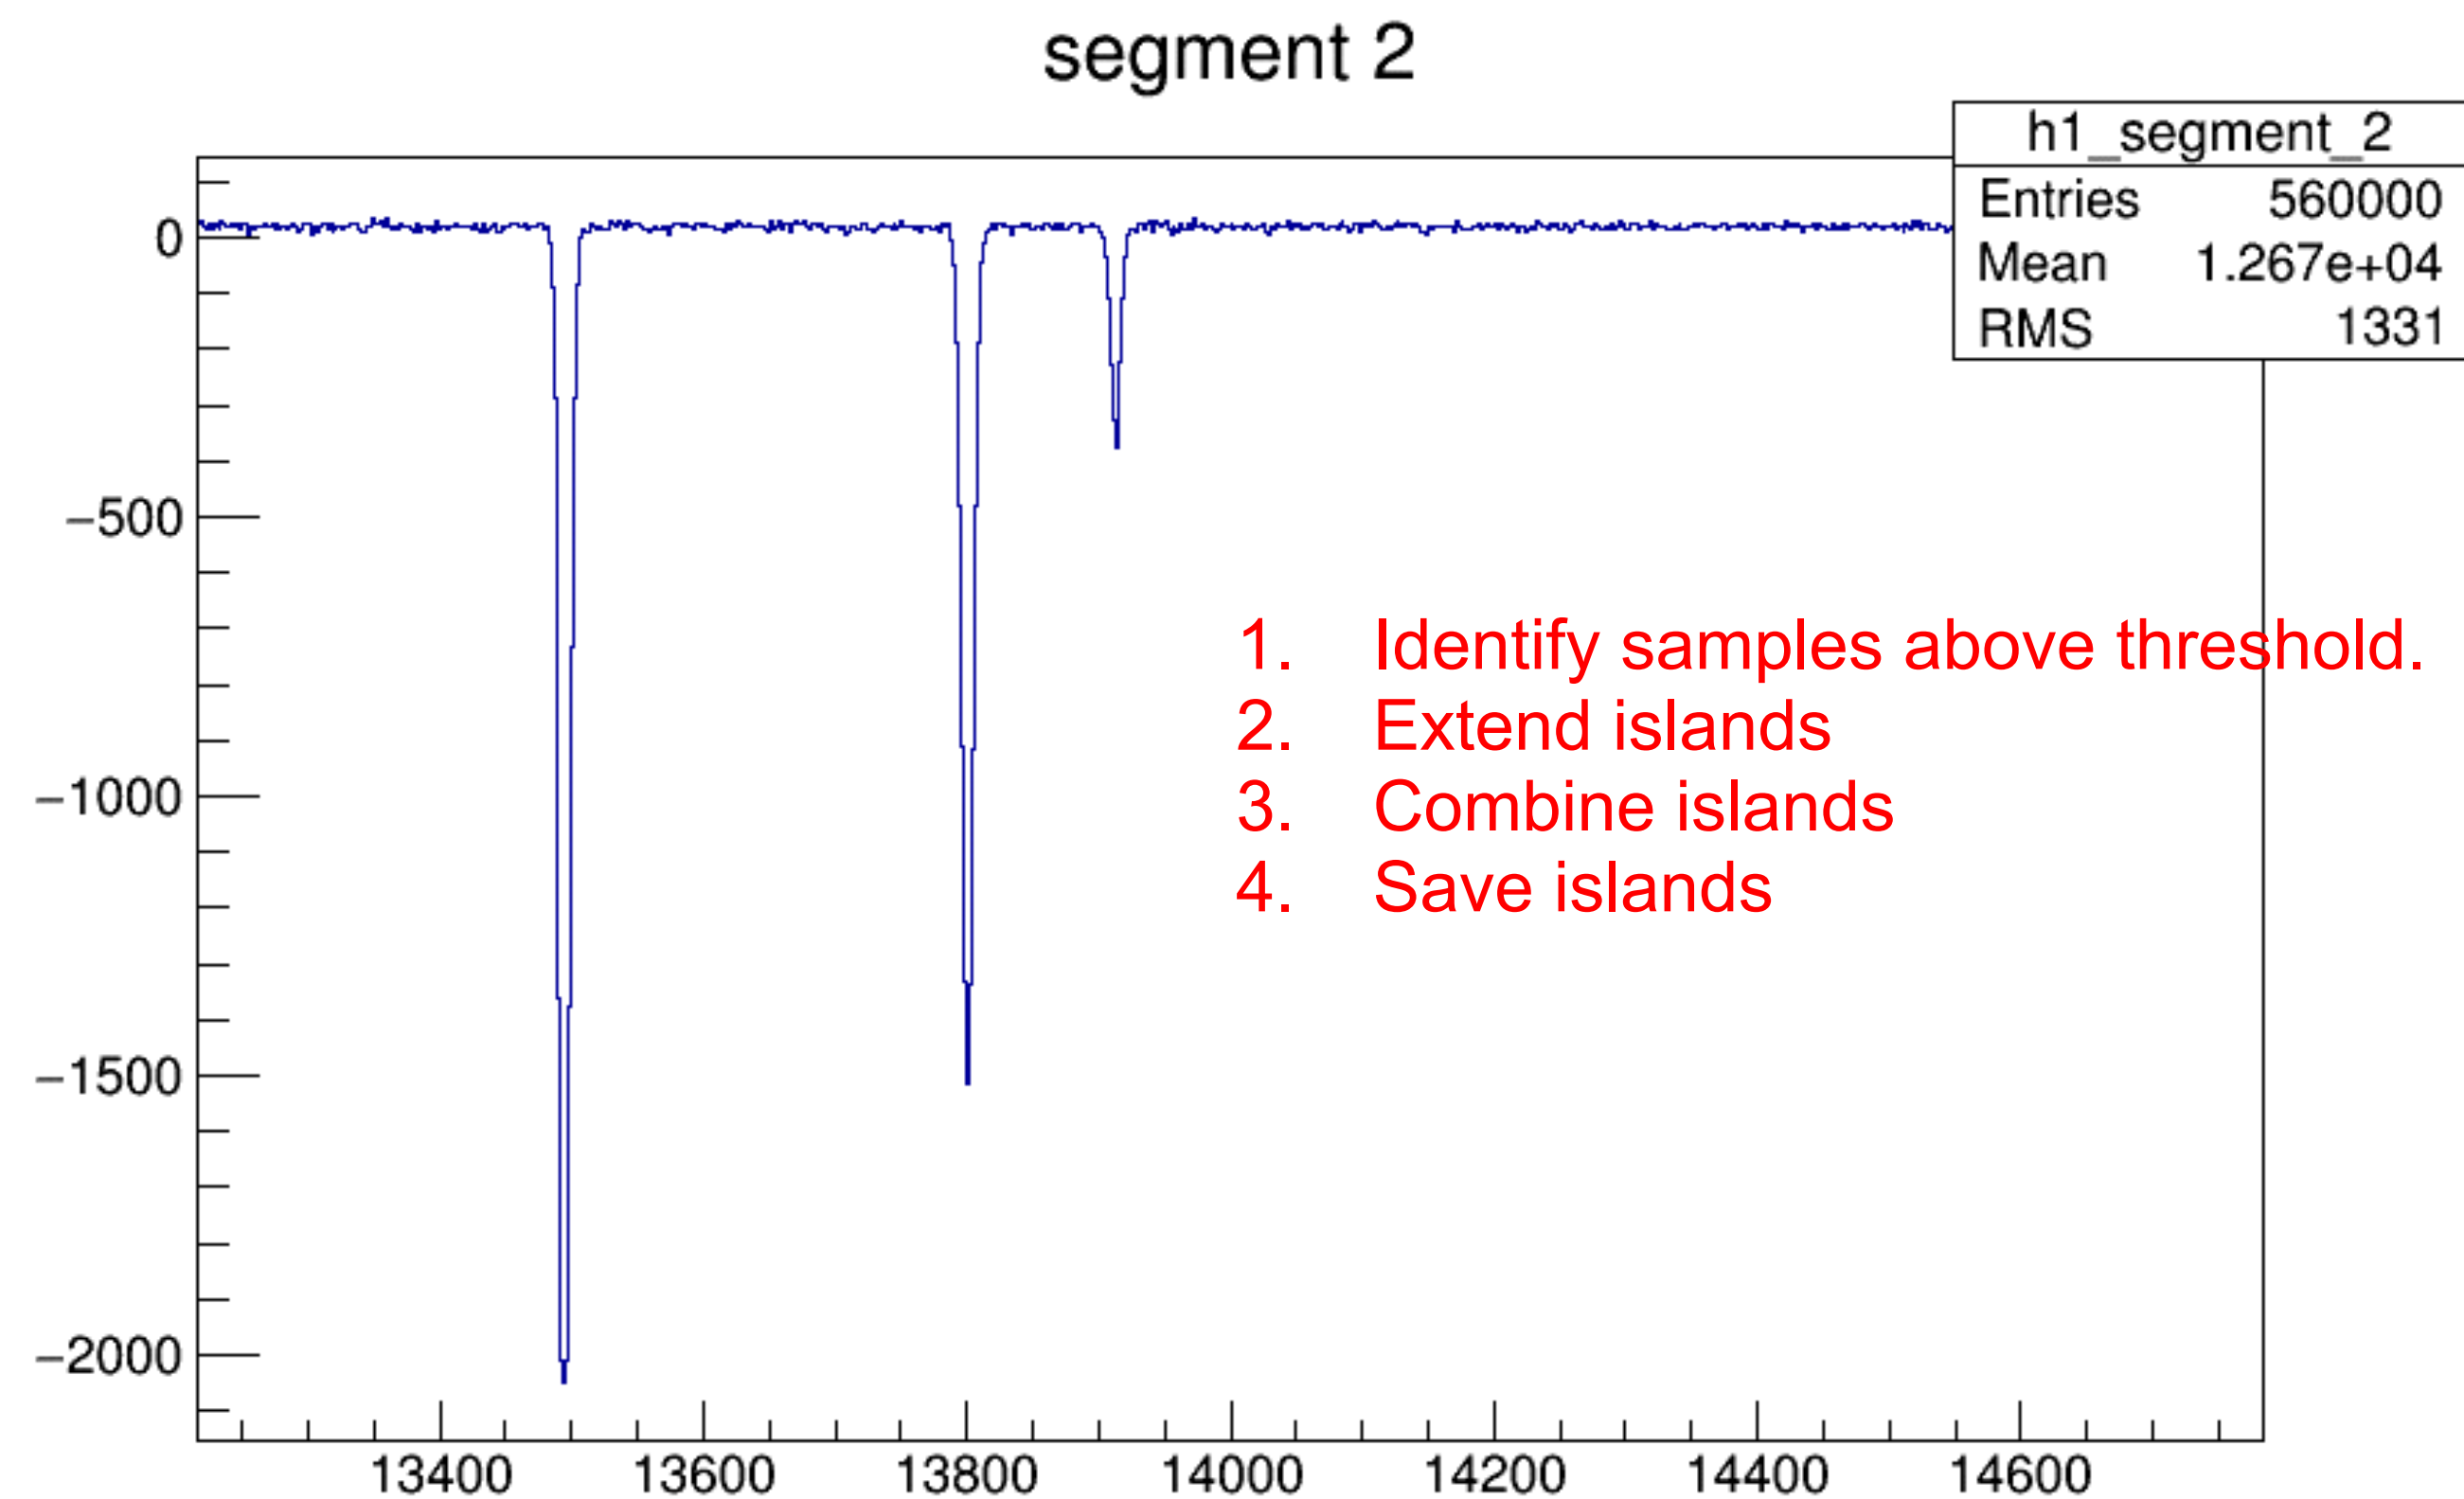




# T-Method



- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.






# T-Method



- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.






# T-Method





- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.





# T-Method



- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.





# T-Method

- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.

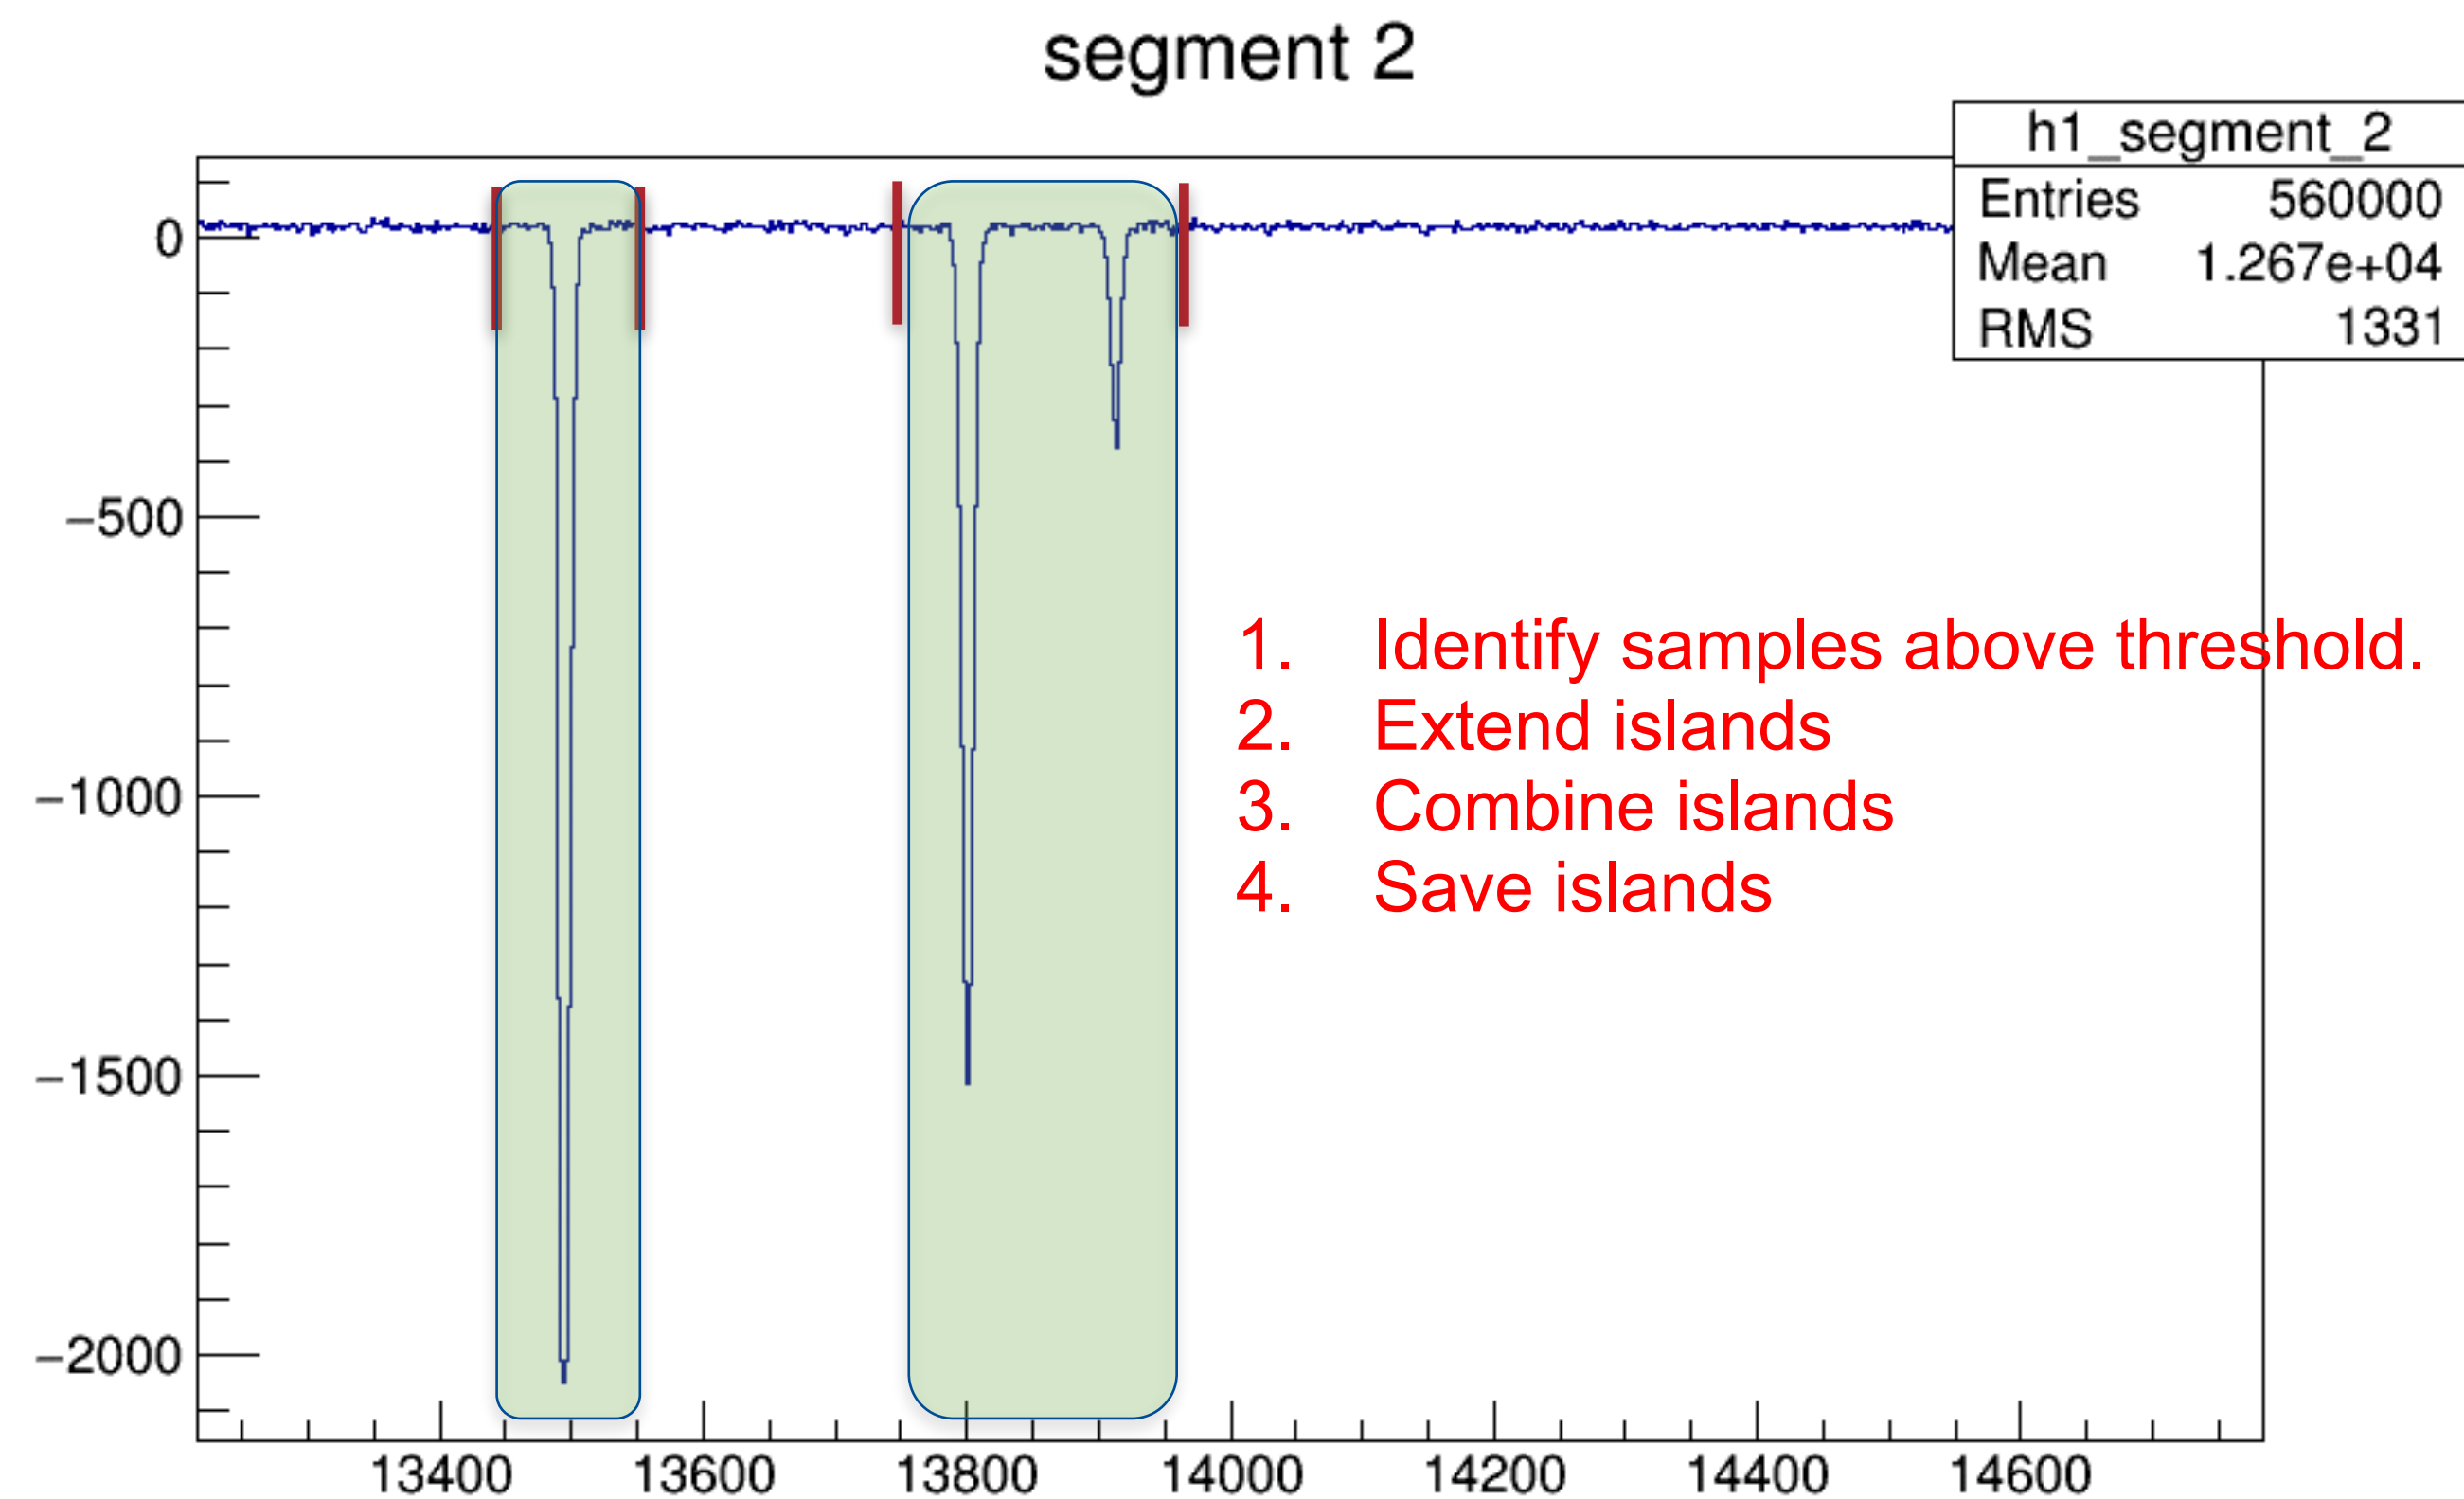




# T-Method



- Identify and save regions of the waveform containing positron hits.
- A typical waveform will have  $\sim 180$  islands.







# Q-method



- Full waveforms are decimated in time and summed over many fills to create a histogram that is saved in the data file.
  - i.e. If we decimate in time by 10 and flush every 100 fills, we reduce the data rate by a factor of 1000, so from 20 GB/s to 20 MB/s.
- Use smaller bins at lower times and wider bins at later times to insure that we can extract the pedestal.



$\times 10^6$






# GPU Template Fits



- Fit templates are loaded into the GPU memory and used to fit each peak above a certain threshold.
- A bank containing the fit results will be saved for each fill in addition to the chopped islands.
- Reduces the processing necessary in the online DQM.








# Processing time



- Must process each event in 83 ms to keep up with average beam rate of 12 Hz.
- Most time is spent reading data from TCP socket and copying it to the GPU.
- Processing time in the GPU is very small.





# Data Quality Monitor



- Online DQM streams data from MIDAS mserver.
- Uses existing art modules to process the data.
- Data is extracted from the art job using ZeroMQ and published to a web GUI using node.js and plotly.







# DAQ Performance During First Run




- The MIDAS DAQ performed well during our first run.
- Day shifts were mostly dedicated to beam line commissioning, so production data was taken at night.



# Summary



- The DAQ for Muon  $g-2$  at Fermilab has been fully installed and performed well during commissioning.
- The DAQ hardware includes 17 frontend machines, 5 backend machines, 2 dedicated line analysis machines, 3 computers for slow control, 3 servers, and 24 beagle bones running 67 MIDAS frontends.
- Takes an input data rate of 20 GB/s, and reduces that to  $< 200$  MB/s in the event builder via GPU processing and lossless compression.
- Commissioning run in June 2017 yielded first “wobble plot” created from  $\sim 700$ k positrons.








# Backup

# Overview



- The DAQ must assemble data from an “Internet of things” to form a complete picture of each muon fill.
- Expect an average data rate of 12 Hz for muon fills, plus some laser and pedestal fills.
- This provides data at a rate of 20 GB/s, which must be reduced via processing of data in GPUs.



# Input sources

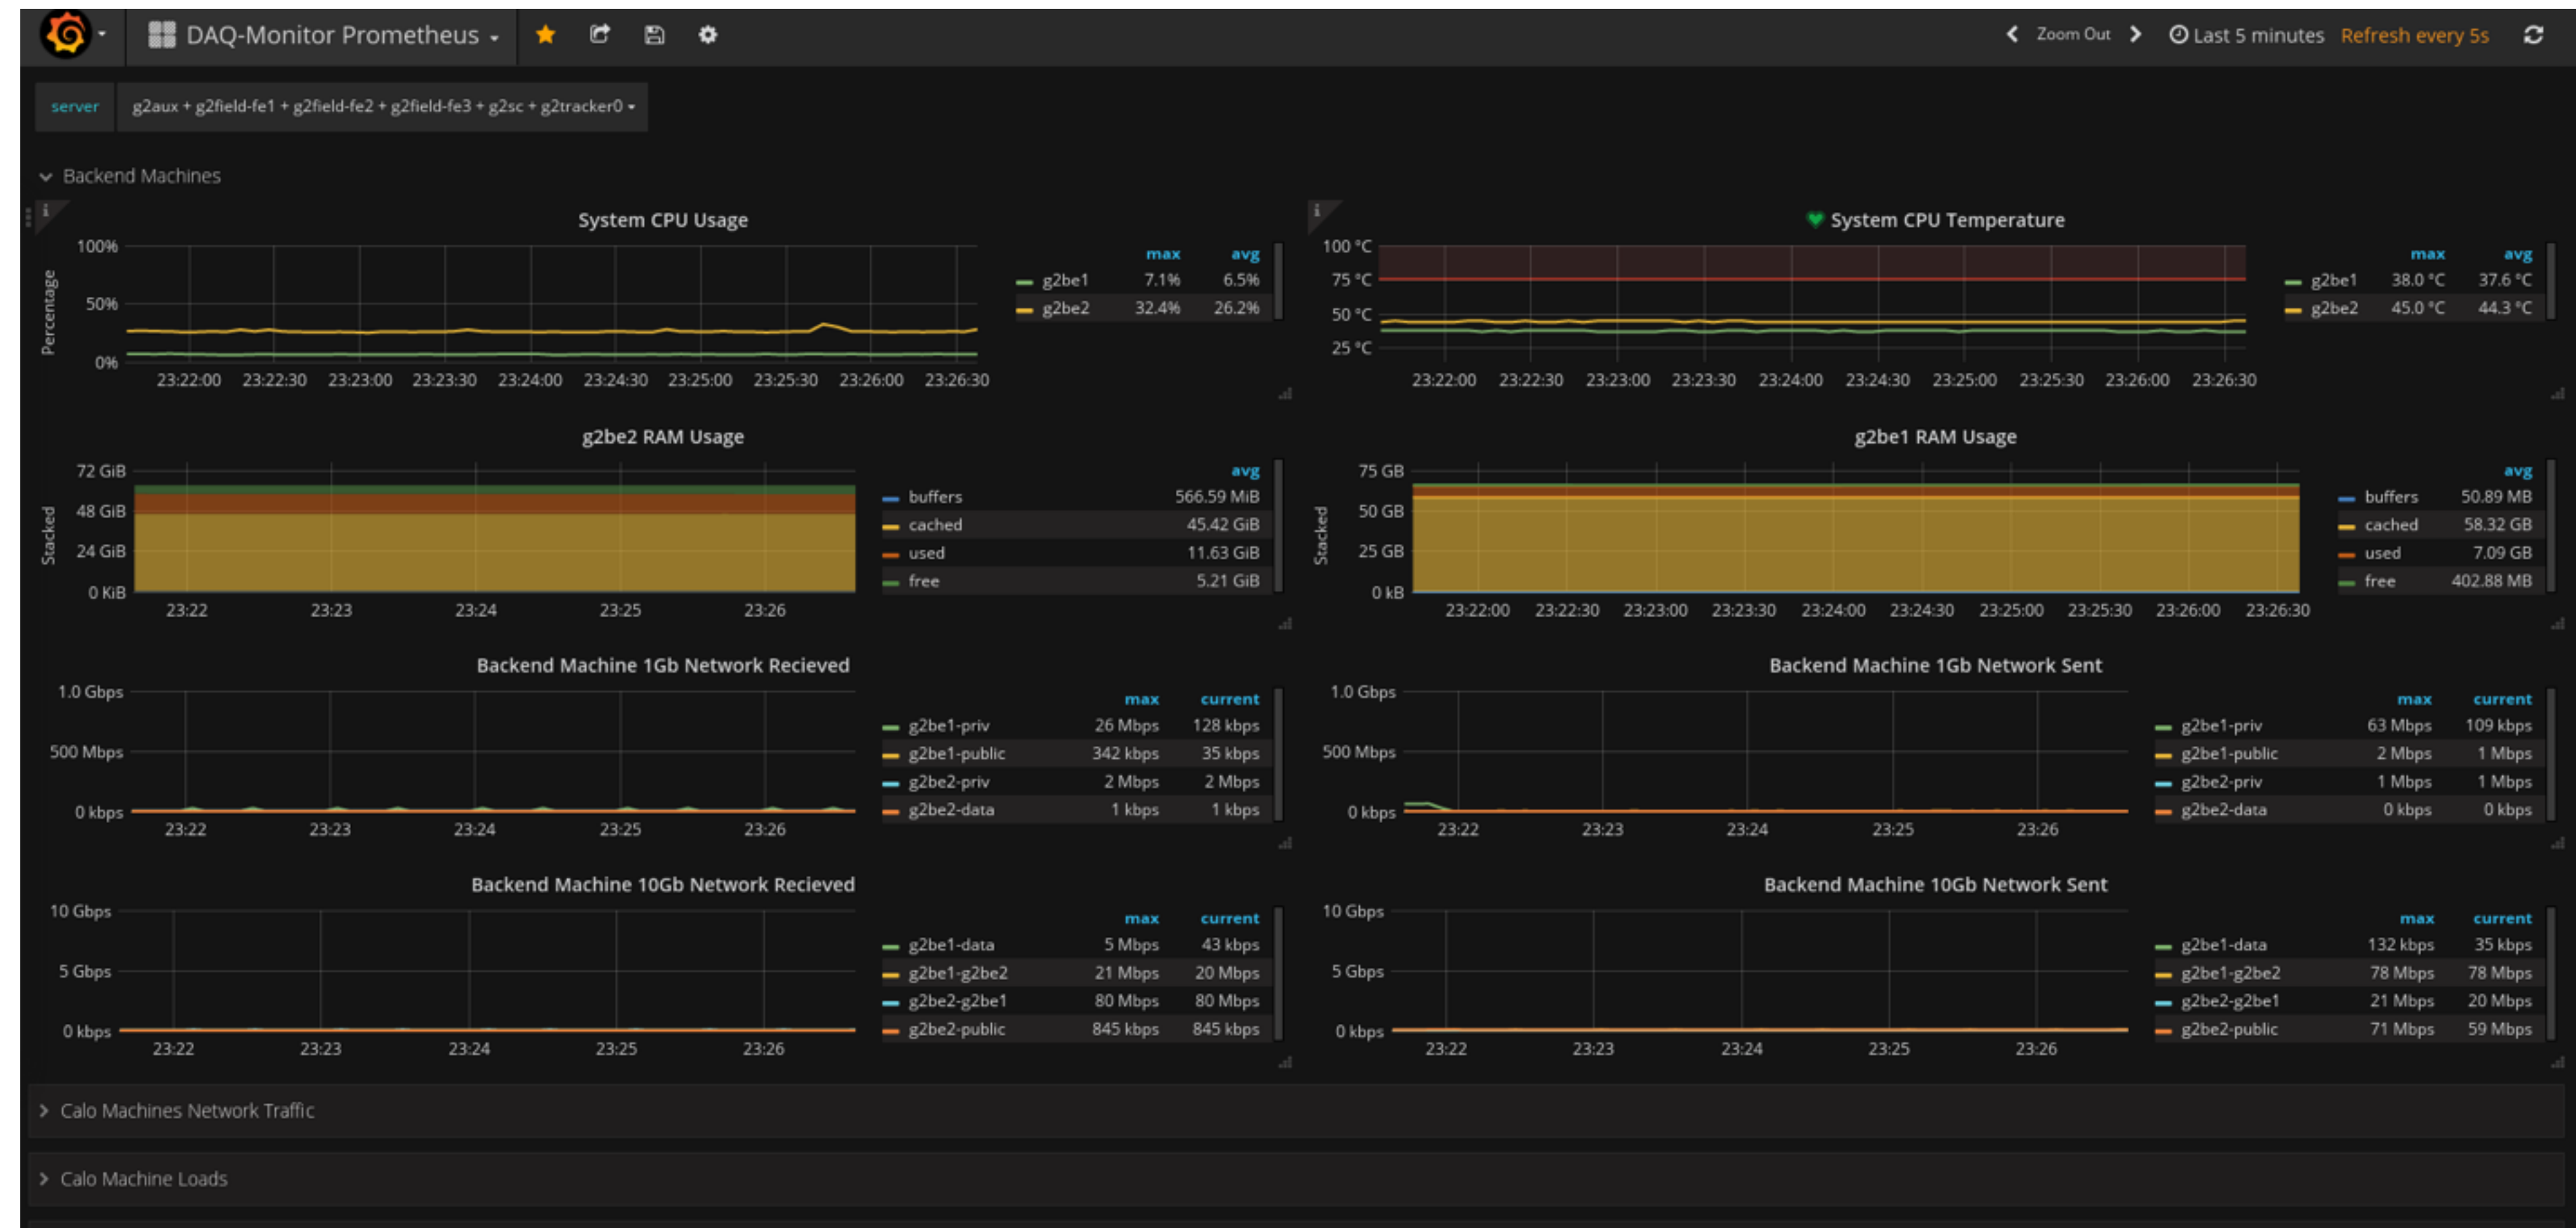


- Digitization is performed in custom uTCA based waveform digitizers.
- Each digitizer runs at 800 MSPS, so each time bin is 1.25 ns, and a 700 us fill is 560,000 clock ticks.
- Each uTCA crate contains 12 WFD5s or 60 channels of digitization.
  - Crate 0 reads data from the clock and control center (CCC)
  - Crates 1-24 each read data from one calorimeter (+ spare channels)
  - Crate 25 reads data from the laser system
  - Crate 26 reads data from the Auxiliary detectors (Harps, Quads, and Kickers)
  - Crate 27 reads data from the three tracker detectors.
- Data from each crate is sent to a DAQ computer via a dedicated 10 Gb fiber. The total data rate is 20 GB/s.
- The data is then processed in Nvidia K40 GPUs.

# DAQ Health Monitor



- Monitored health of DAQ systems using netdata for system monitoring, prometheus for short term data storage, and grafana to display data.






# MIDAS Sequencer



- The sequencer was used extensively for calibration runs such as filter wheel scans and bias voltage scans.
- A typical sequence would be:
  - Execute script to move wheel.
  - Update ODB values
  - Take data for 10 minutes
  - Repeat



**Progress**  
Sequence is finished

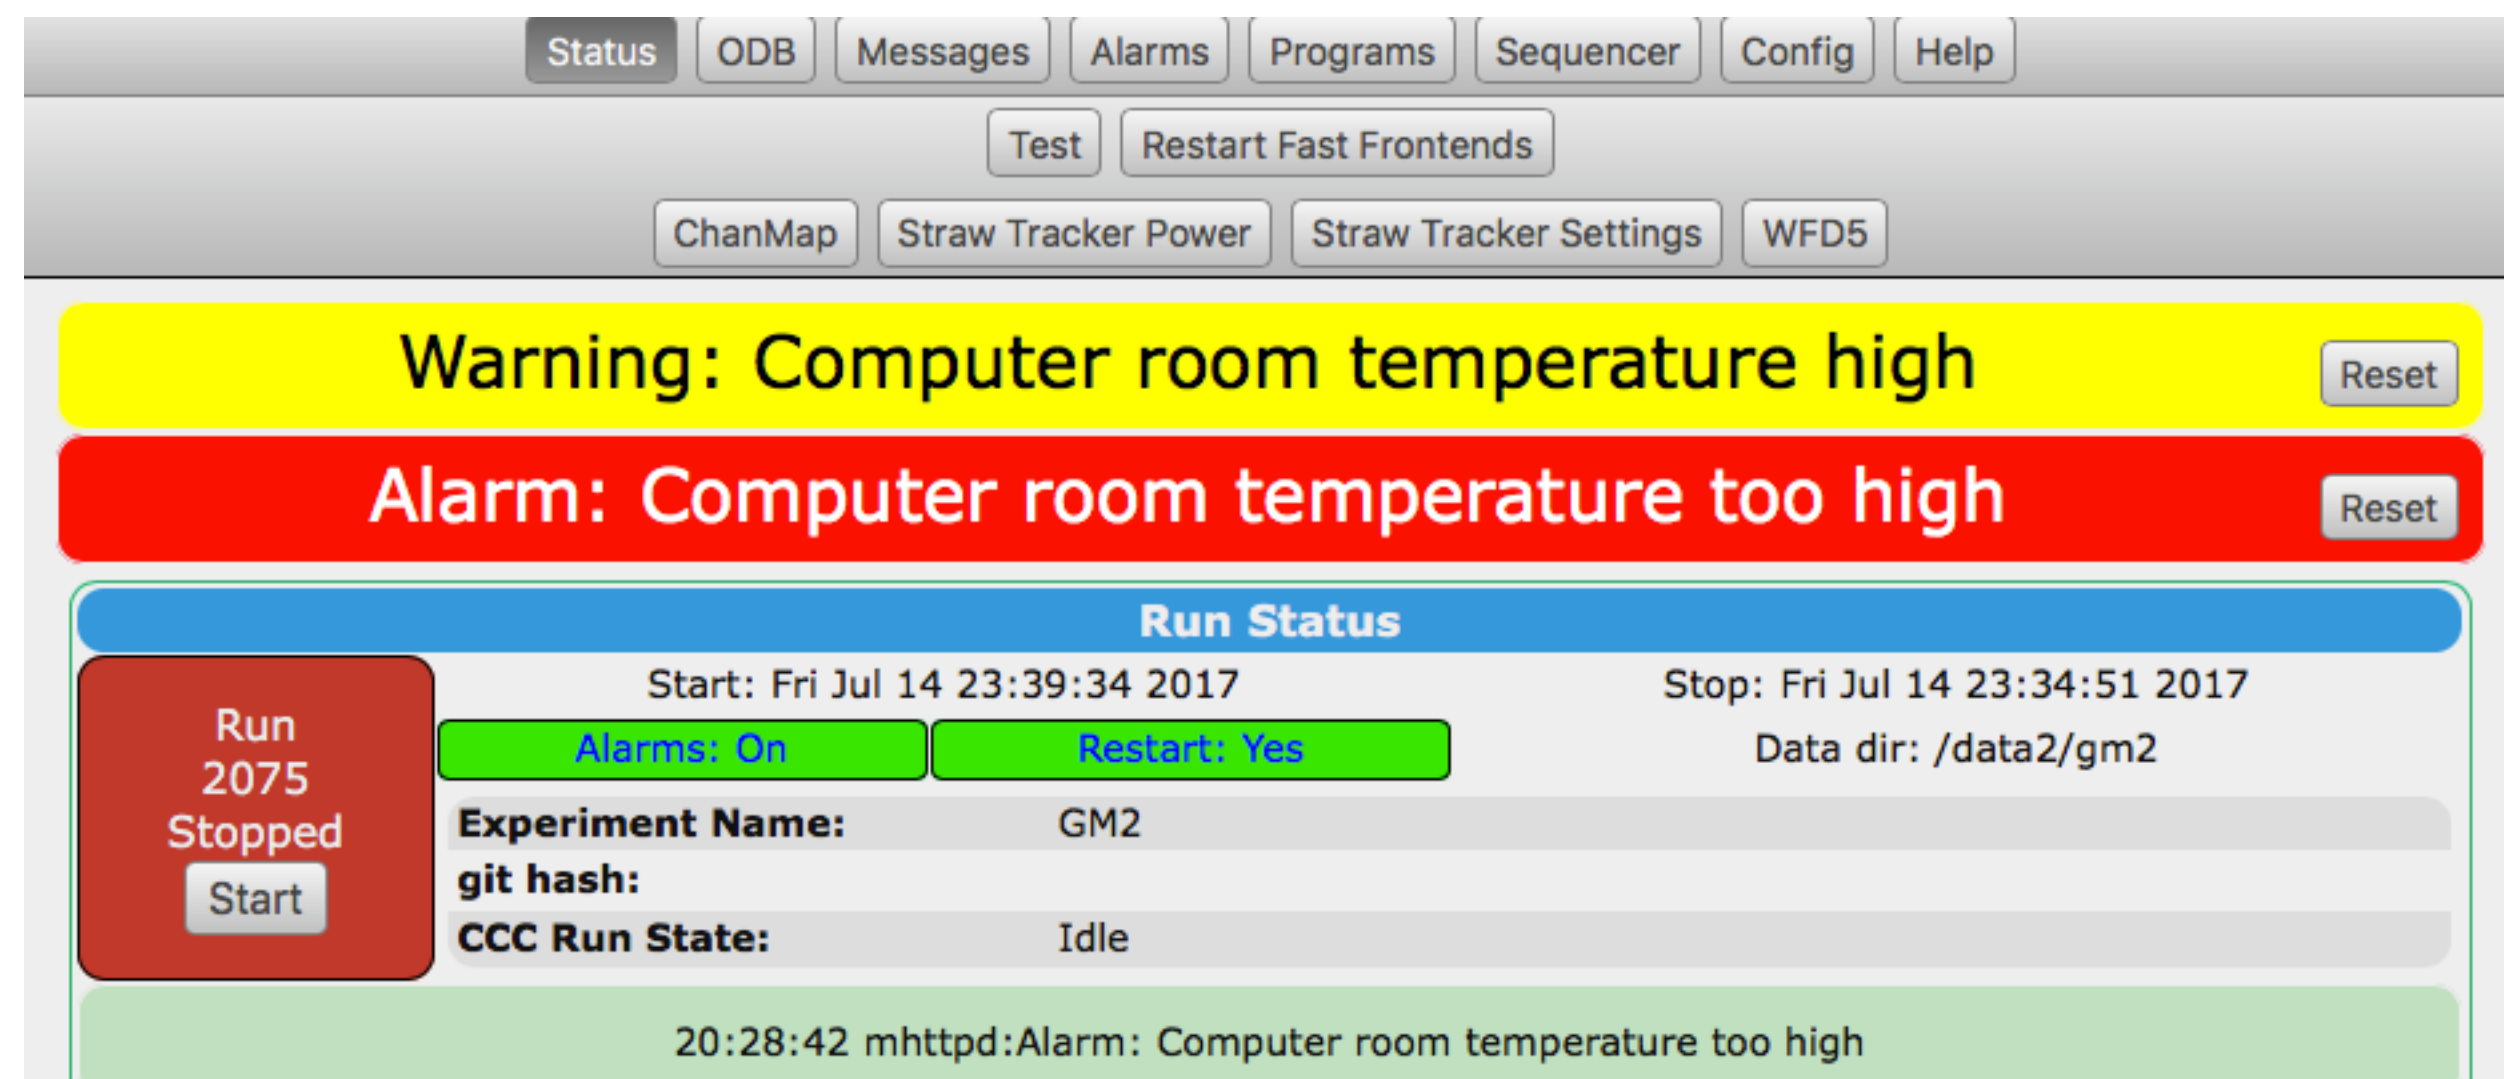
**Sequencer File**  
Filename: wheel\_setting.msl [Show XML](#)

```
1 COMMENT "Run a sequence of laser runs"
2 RUNDESCRIPTION "Calibration run"
3
4 LOOP angle, 1, 2, 3, 4, 5, 6, 7, 1
5   ODBSET "/Experiment/Edit on start/Comment (256 max)", "Automated run"
6   ODBSET "/Experiment/Edit on start/Quality YNCT", "C"
7   SCRIPT /home/daq/test_wheel.sh, $angle
8   WAIT Seconds 2
9   TRANSITION START
10  WAIT Seconds 200
11  TRANSITION STOP
12 ENDLOOP
13 ODBSET "/Experiment/Edit on start/Comment (256 max)", "Enter comment"
14 ODBSET "/Experiment/Edit on start/Quality YNCT", "T"
```

# MIDAS Alarms



- The MIDAS alarm system was used as the primary alarm system.
- Alarms were set on temperatures and voltages from MSCB devices.
- Other slow frontends set alarms automatically when encountering an error.
- Periodic alarm reminded shifters to perform shift checks.
- Had problems at first with alarm audio, which was traced to a recent lack of mp3 support in scientific linux — this was later rectified with a recent update.





# Custom Controls Page



- With this number of frontends, configuring settings via the standard ODB tree is very cumbersome.
- A set of custom Javascript pages were written to manipulate ODB values en masse.

The screenshot shows a web-based configuration interface for the Muon g-2 experiment. At the top, there is a navigation bar with buttons for Status, ODB, Messages, Chat, Alarms, Programs, History, MSCB, Sequencer, Config, and Help. Below this, there are dropdown menus for 'AMC1300' and 'TQ01', and buttons for 'Threshold', 'Det x-segmt', and 'Det y-segmt'.

The main content area is divided into two sections, each with a blue header bar:

**Section 1: /Equipment/AMC1300/Settings/TQ01/RiderXX/ChannelXX/threshold value**

Channel #	Rider 01	Rider 02	Rider 03	Rider 04	Rider 05	Rider 06	Rider 07	Rider 08	Rider 09	Rider 10	Rider 11	Rider 12
00	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200
01	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200
02	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200
03	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200
04	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200	-200

Below this table are three buttons: 'Enabled', 'Chan used', and 'Positive crossing'.

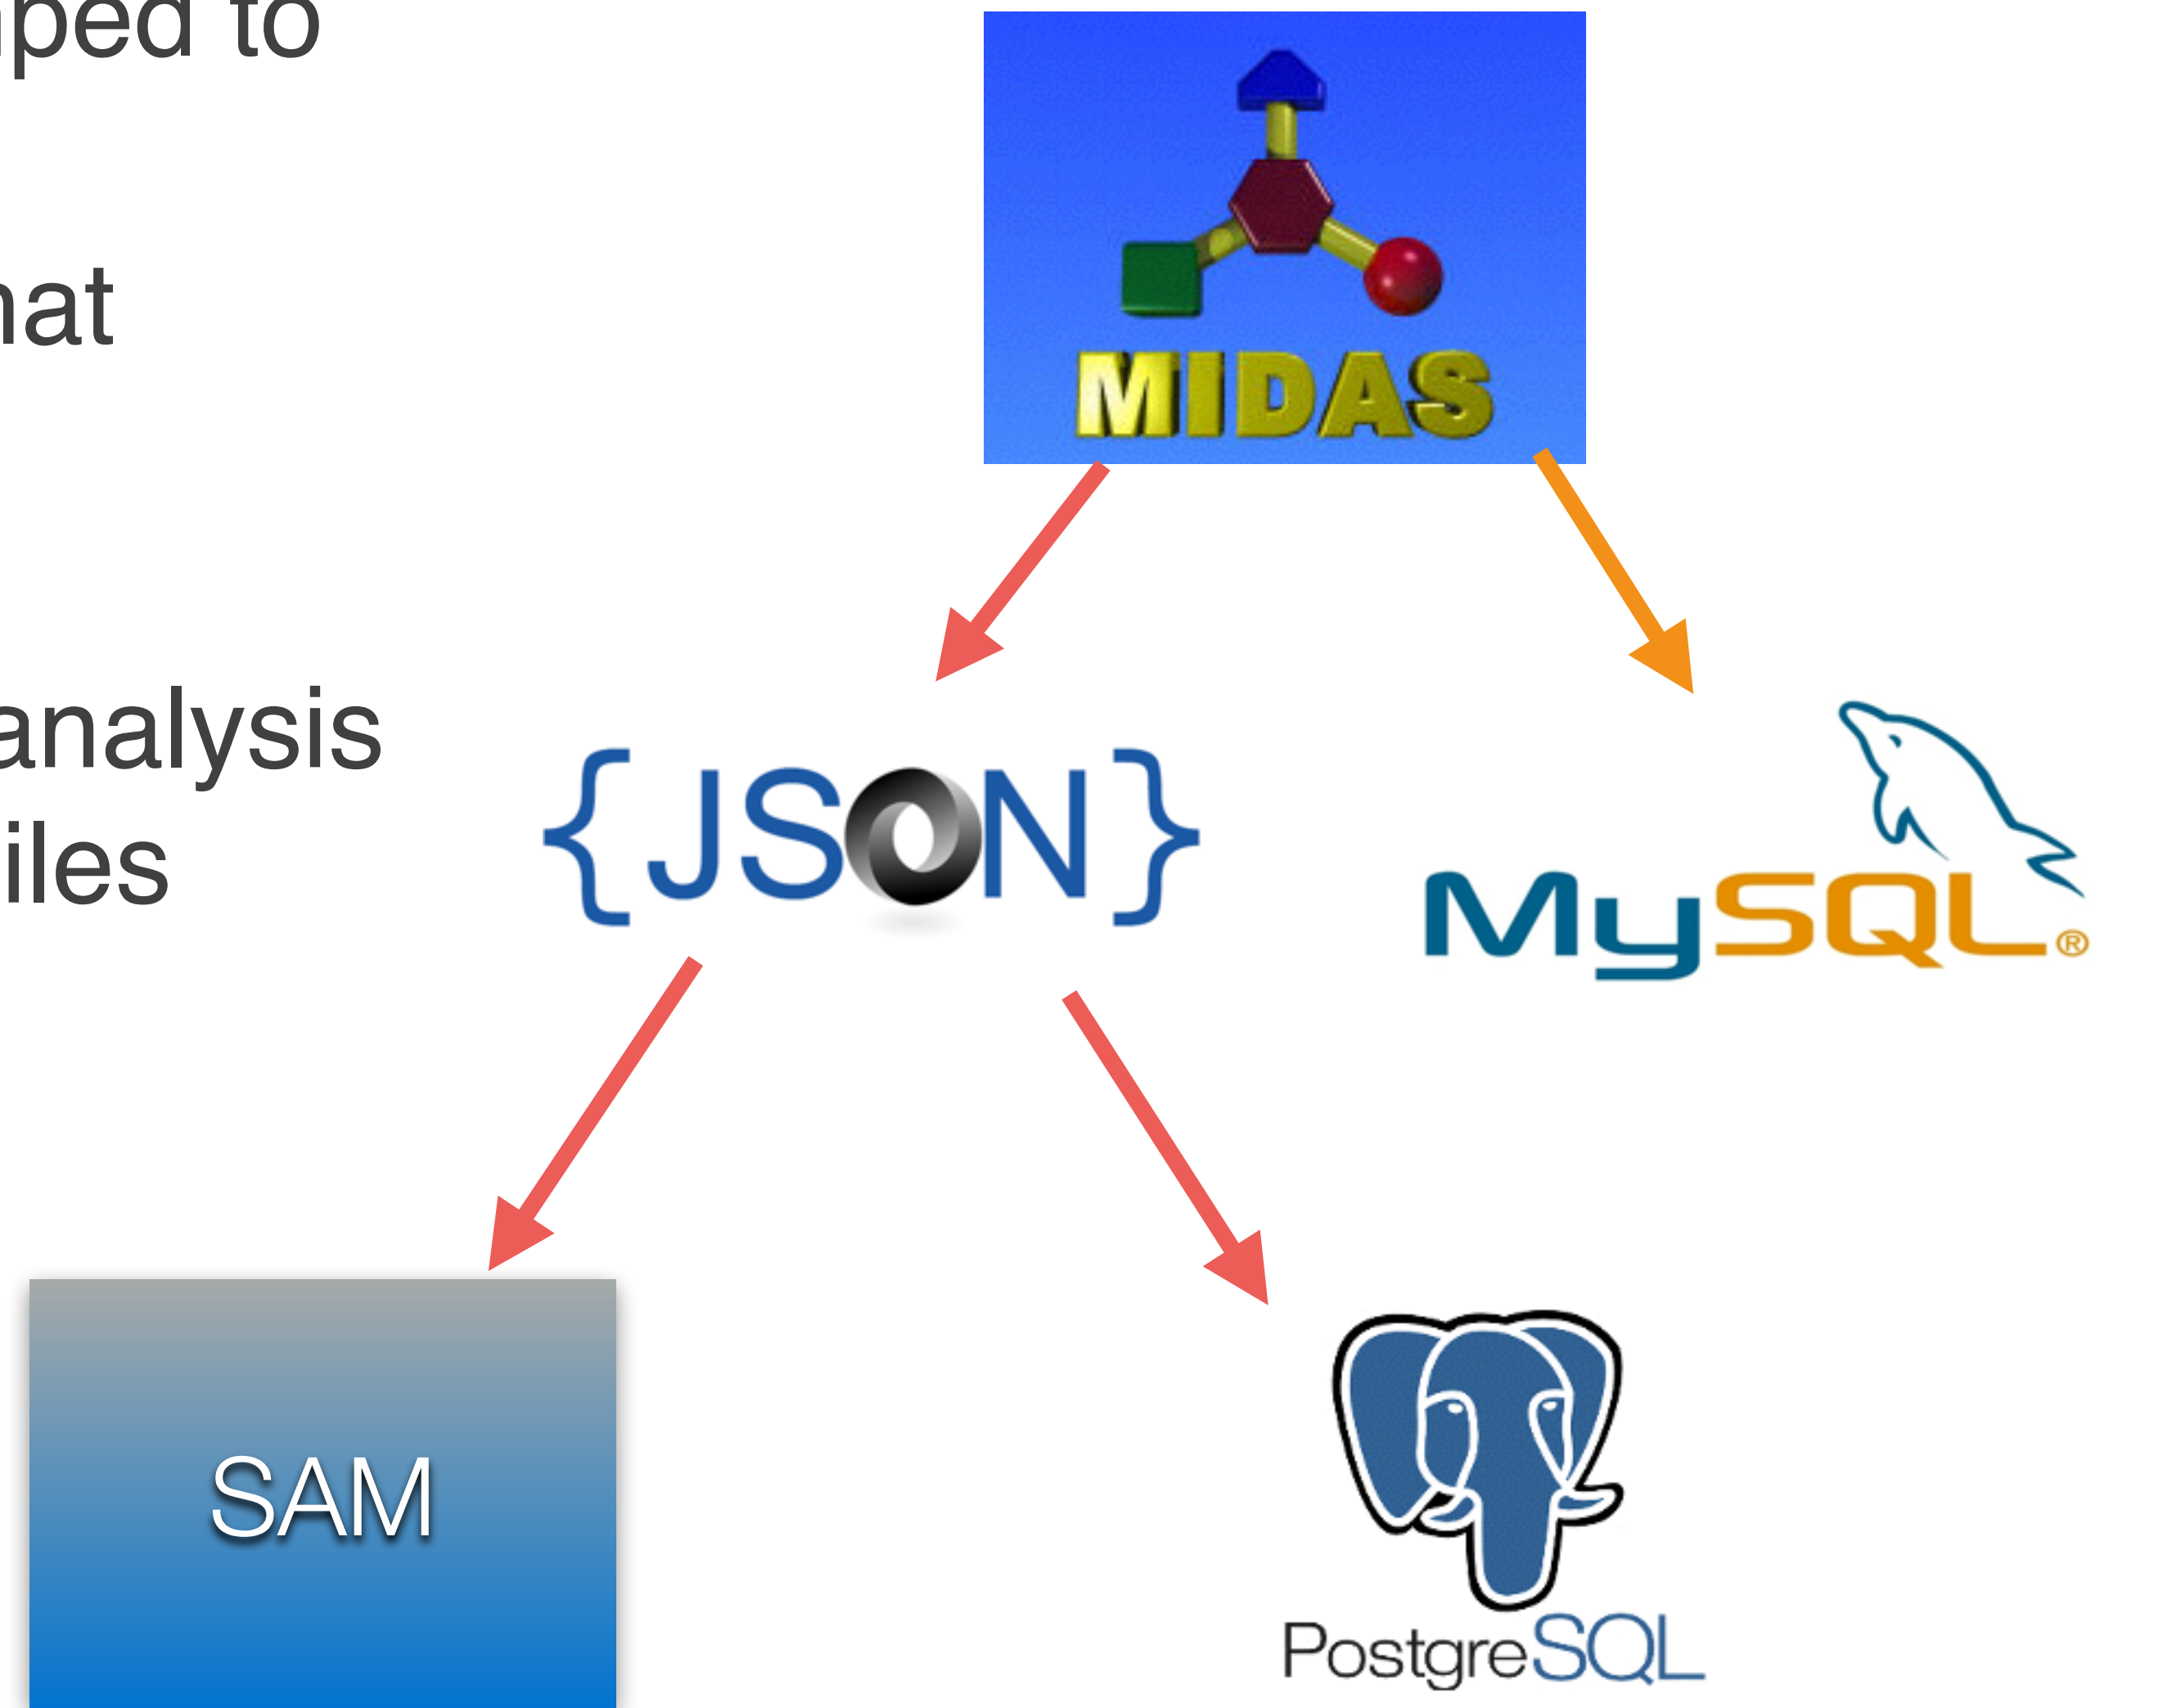
**Section 2: /Equipment/AMC1300/Settings/RiderXX/ChannelXX/enabled**

Channel #	Rider 01	Rider 02	Rider 03	Rider 04	Rider 05	Rider 06	Rider 07	Rider 08	Rider 09	Rider 10	Rider 11	Rider 12
00	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
01	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
02	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
03	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
04	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

At the bottom of the interface is a red button labeled 'WRITE to ODB'.

# MIDAS ODB Archive

- At each end of run, the ODB is dumped to a JSON file.
- A python routine is then executed that imports the entire JSON file into a PostgreSQL database.
- Metadata that is used in the offline analysis is also extracted from these JSON files using python plugins.





# Field DAQ



- Field DAQ runs in independent MIDAS experiment.
- Contains seven asynchronous frontends reading data from fixed magnetic field probes and from a trolley that periodically transverses the ring to perform precision measurements of magnetic field.
- Data is correlated with the fast DAQ offline using GPS timestamps.

StatusODBMessagesAlarmsProgramsHistoryMSCBHelp

Restart Front-EndsRestart LoggerRestart ServerRestart DQMs

Trolley ControlPlunging Probe Control

Run Status

Run 395RunningStop

Start: Thu Jun 1 08:52:27 2017Running time: 0h26m35sData dir: /home/newg2/gm2Data/

Alarms: OnRestart: No

Experiment Name: g2-fieldTrolley Status: n

09:19:00 [Fixed Probes,DEBUG] issued trigger

Equipment

Equipment	Status	Events	Events[/s]	Data[MB/s]
Fixed Probes	Fixed Probes@g2field-fe2-priv	398	0.3	0.995
TrolleyInterface	Frontend stopped	0	0.0	0.000
GalilFermi	Frontend stopped	0	0.0	0.000
Surface Coils	Frontend stopped	0	0.0	0.000
Monitor	Frontend stopped	0	0.0	0.000
Fluxgate	Frontend stopped	0	0.0	0.000
PS Feedback	Frontend stopped	0	0.0	0.000

Logging Channels

Channel	Events	MiB written	Compr.	Disk level
#0: run00395_00.mid.gz	399	949.078	N/A	17.1 %

Muon g-2 DQMRun395 Event5Subsystems

Ring Yoke Probe

Update Options

RefreshAutomatic OnEvery 5 secondsDisplay OptionsAll Probes

Average Frequency: ---

Healthy Probes : 324

Front-end Online

Run in Progress

Trolley Run

# Slow Controls



- DAQ includes six SCS3000 mscb devices.
- 24 beaglebones reading slow control data from calorimeters.
- HV and LV frontends for tracker system.
- Slow frontend reading magnet properties from IFIX via an OPC client.
- Beamline frontend periodically reading output of beam components from database.
- Slow control data is stored in a Postgres database and displayed using a custom Django web display.








# IBMS Frontend

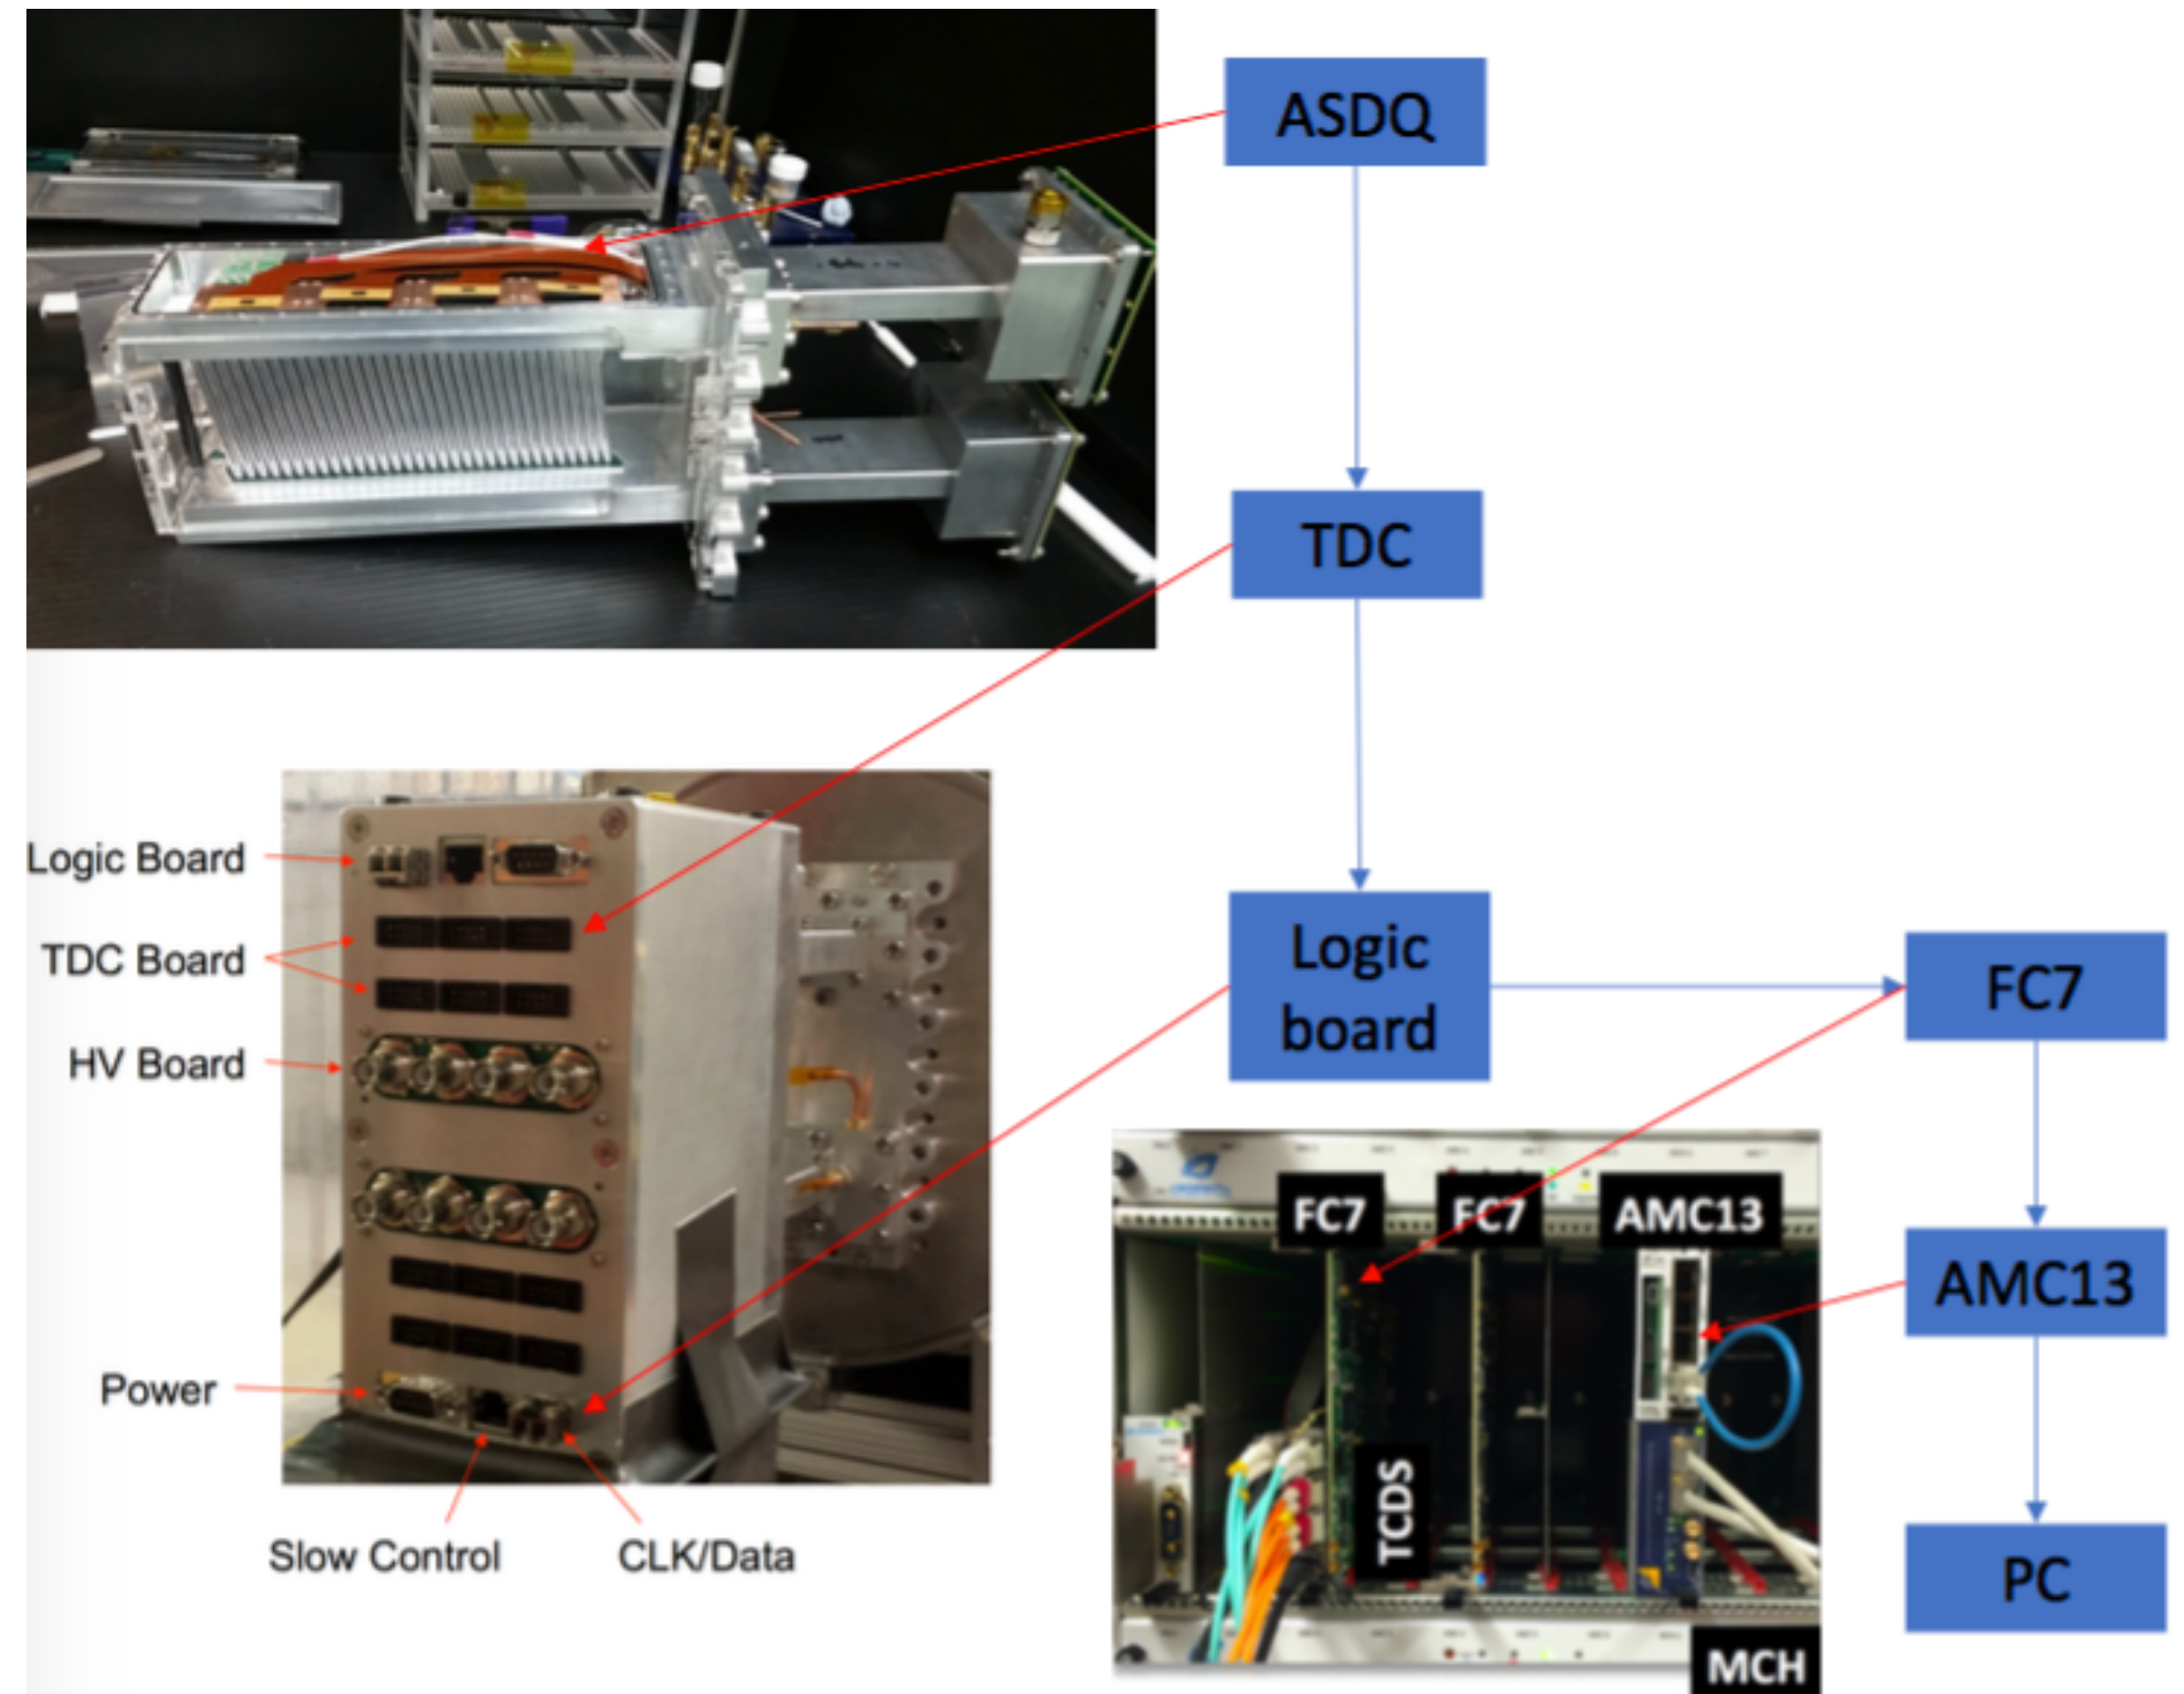


- Data from the inflector beam monitoring system (IBMS) is read out via a CAEN digitizer.
- A custom MIDAS frontend was written to integrate this detector into the DAQ.



# Tracker Frontend

- Three tracker stations will be read via one uTCA crate.
- Reads data from AMC13.
- Instead of digitizers, data comes from multihit TDCs that are read via FC7 cards.




(Thanks R. Chislett for the diagram)



# MasterGM2 frontend

- Communicates with other frontends using RPC calls.
- Provides begin of run and end of run RPCs to all frontends.
- Provides end of fill trigger to synchronous frontends.
- Configures clock and control system.
- Reads trigger times from Meinberg GPS unit and writes them to a MIDAS bank.



Bank:GPS0 Length: 20(I\*1)/5(I\*4)/5(Type) Type:Unsigned Integer\*4  
1-> 0x00018c2 0x580838bc 0x5dcabfb0 0x580838bc 0x5dd48308

# Event builder



- Modified event builder to change how it is enabled — now done entirely in ODB Equipment of each frontend.



Online Database Browser

Find Create Delete Create Elog from this page

/ Equipment / AMC1308 / Settings / Globals /

Key	Value	+
sync	n	
use AMC13 simulator	n	
GPU Device ID	2 (0x)	
Send to Event Builder	y	
Shelf configuration	rider	
FE lossless compression	n	
raw data store	y	
raw data prescale	1000 (0x3E8)	
raw data prescale offset	8 (0x8)	

Enable here!



- Total EB rate maxed out at  $> 1.2$  GB/s (limited by network bandwidth)
- The event builder combines up to 270 banks for each event.