



The NOvA Far Detector Data Acquisition System

Jaroslav Zalesak

**The Institute of Physics of the Czech Academy of Sciences / Fermilab
and**

**Peter Shanahan, Andrew Norman, Kurt Biery, Ronald Rechenmacher, Jonathan Paley,
Alec Habig, Susan Kasahara, Denis Perevalov, Mathew Muether**

for the NOVA Collaboration



CHEP 2013 in Amsterdam, The Netherlands

14 Oct 2013



The NOvA Experiment



Goals:

- Measure $\nu_\mu \rightarrow \nu_e$ oscillations:
 - Measure θ_{13}
 - Determine the mass hierarchy
 - Constrain δ_{CP}
 - Determine the θ_{23} octant
 - Measure ν_μ disappearance.
 - Precision measurement of $|\Delta m_{32}^2|$, $\sin^2 2\theta_{23}$
- **Other physics:**
 - Near Detector neutrino cross-sections
 - Sterile neutrinos
 - Supernova search
 - Monopole search
 - Dark matter searches

A long-baseline neutrino oscillation experiment



The NOvA Detectors




Far Detector:

- Surface Detector
- 14 kton
- 60 x 15.6 x 15.6 m³
- 896 alternating X-Y planes
- 334 064 cells
- “Totally Active” Calorimeter
- 3D - tracking
- Liquid Scintillator filled PVC
- Optimized for EM shower reconstruction & muon tracking


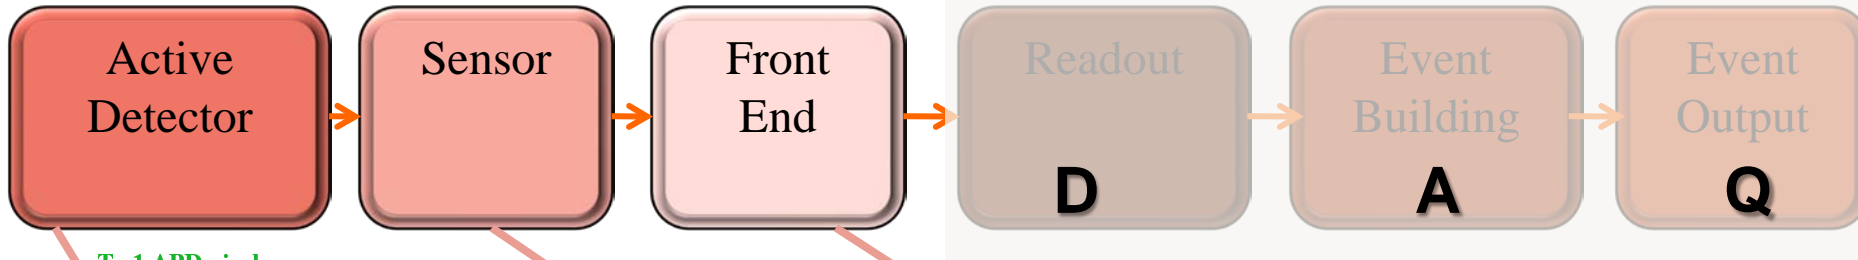
Near Detector:

- Identical to Far detector
- Underground Detector
- 1:4 scale size
- Optimized for NuMI cavern rates - 4x sampling rate electronics
- Construction just began


FarDet: Assembly status 7Oct2013




NOvA Sensor and Front End



ASIC:
Amplification, shaping, multiplexing to ADC
Front End Board (FEB)



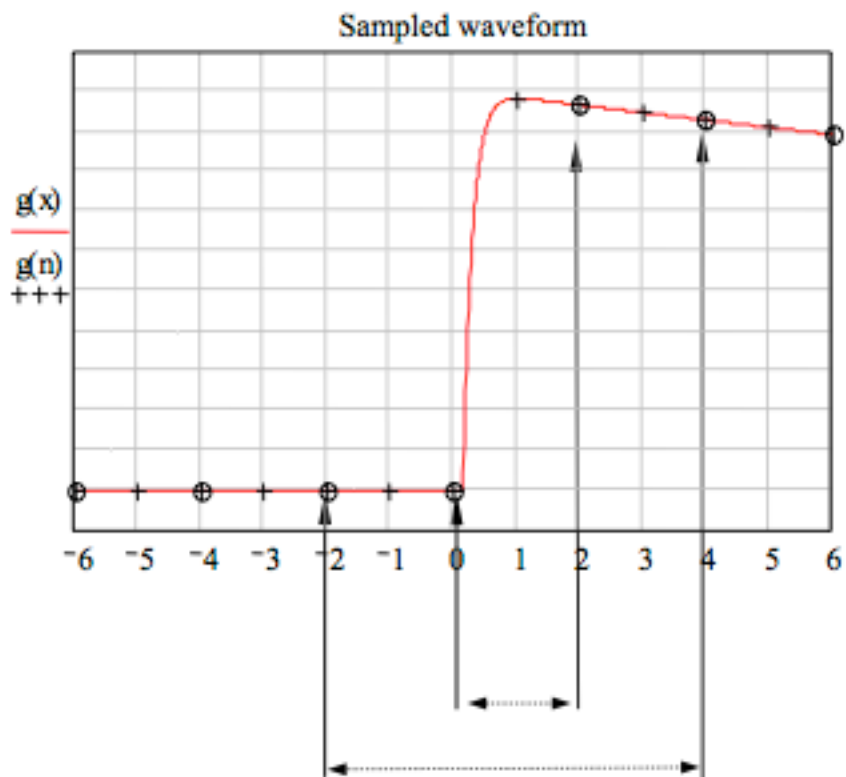
FPGA:
Digital signal processing



Avalanche Photo Diodes (APD):

- 32 channels
- 85% Quantum Efficiency
- Gain ~100
- Cooled to -15C for 2PE dark noise

FEB Signal Processing



DSO Mode

- Digital Sampling Oscilloscope
- Take (1000) contiguous samples
- Use to measure noise level and set appropriate DCS threshold

ASIC Output Pulse Shaping (FarDet)

- 380 ns rise time, 7 μ s fall time
- Sample each channel every 500 ns
- Actual clock runs at 62.5 ns,
- Four 8:1 multiplexors

Near Detector

- 60ns rise time, 500ns fall time
- Sample each channel every 125 ns
- Same base clock, but 2:1 multiplexing

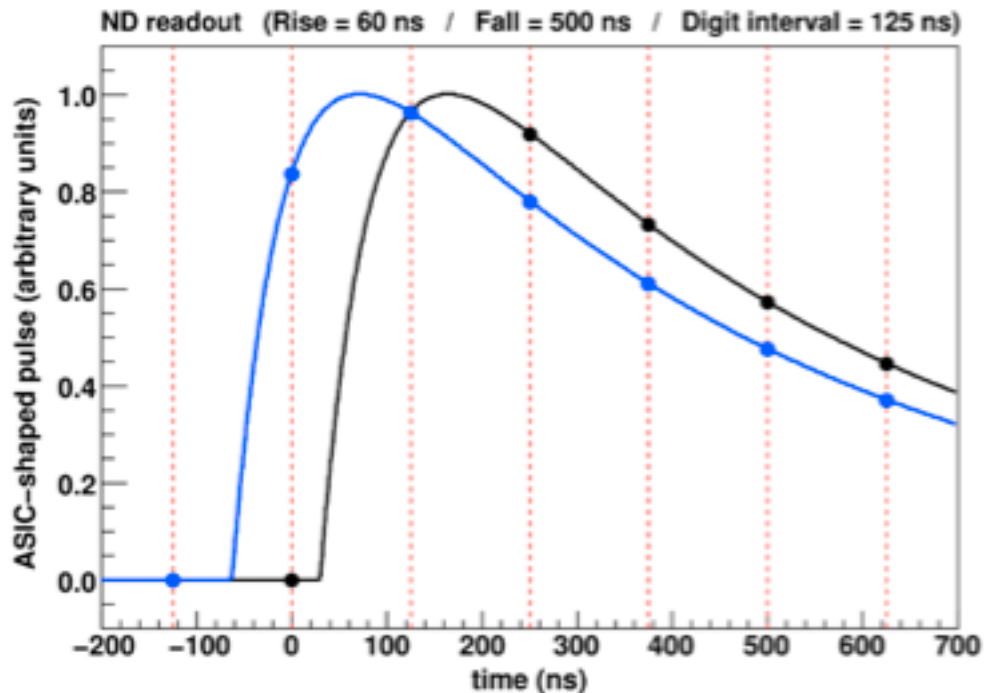
DCS Mode

- Dual Correlated Sampling
- Output data when $V_i - V_{i-2}$ exceeds configured threshold
- Data: pulse height = $V_i - V_{i-2}$, timestamp = T_i

Fast Timing



Two pulses with identical DCS output Pulseheight and Time



Multi-point sampling

- Apply DCS-threshold, but read out multiple contiguous samples
- With compression, can read out up to 12 samples with only 2 additional data words
- Apply matched filtering to improve timing resolution substantially
- Factor of 3 improvement demonstrated in Data with 4 points

Benefits

- Reduce Near Detector Pile-up
- FarDet: track direction from timing

DAQ Requirements



- ❑ Record beam spills in a window of at least $30 \mu\text{s}$, with no dead time
 - Every 1.33 seconds


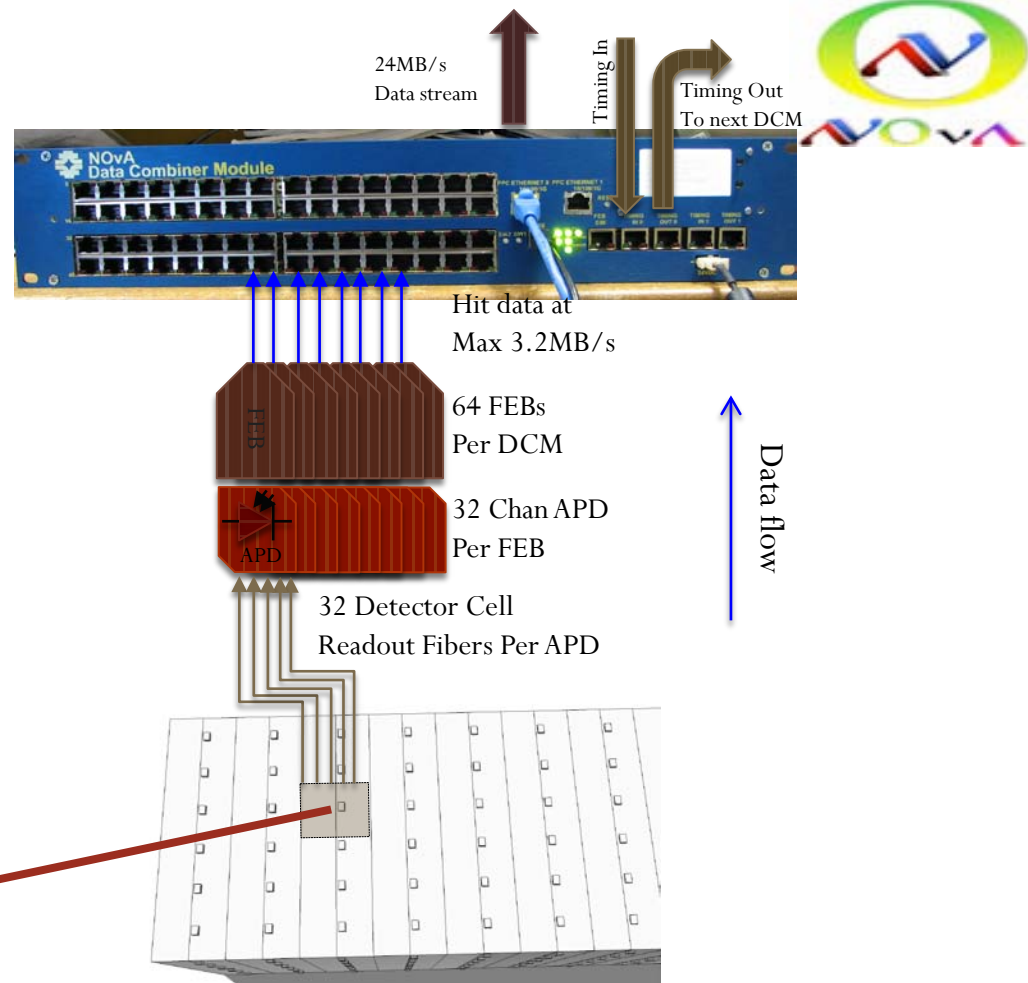
- ❑ Record ~ 100 times as many periodic cosmic spills with no dead time

- ❑ Cosmic rays (200 kHz!) and noise add up to a hit rate of about 52 MHz, or 620 MB/sec
 - Process 620 MB/sec readout with no dead time
 - Write out 1.4 MB/sec saved events.

NOvA DAQ Layout

DCM (Data Concentrator Module)

- Designed to aggregate the continuous data stream 64 front end boards (2048 detector cells) into partially time sorted data block
- Combines a large FPGA with a single board computer to build 5ms
- Run custom embedded linux with support for NOvA firmware
- 168 DCMs for the Far Detektor
- 620 MB/sec readout, ~4 MB/sec/DCM



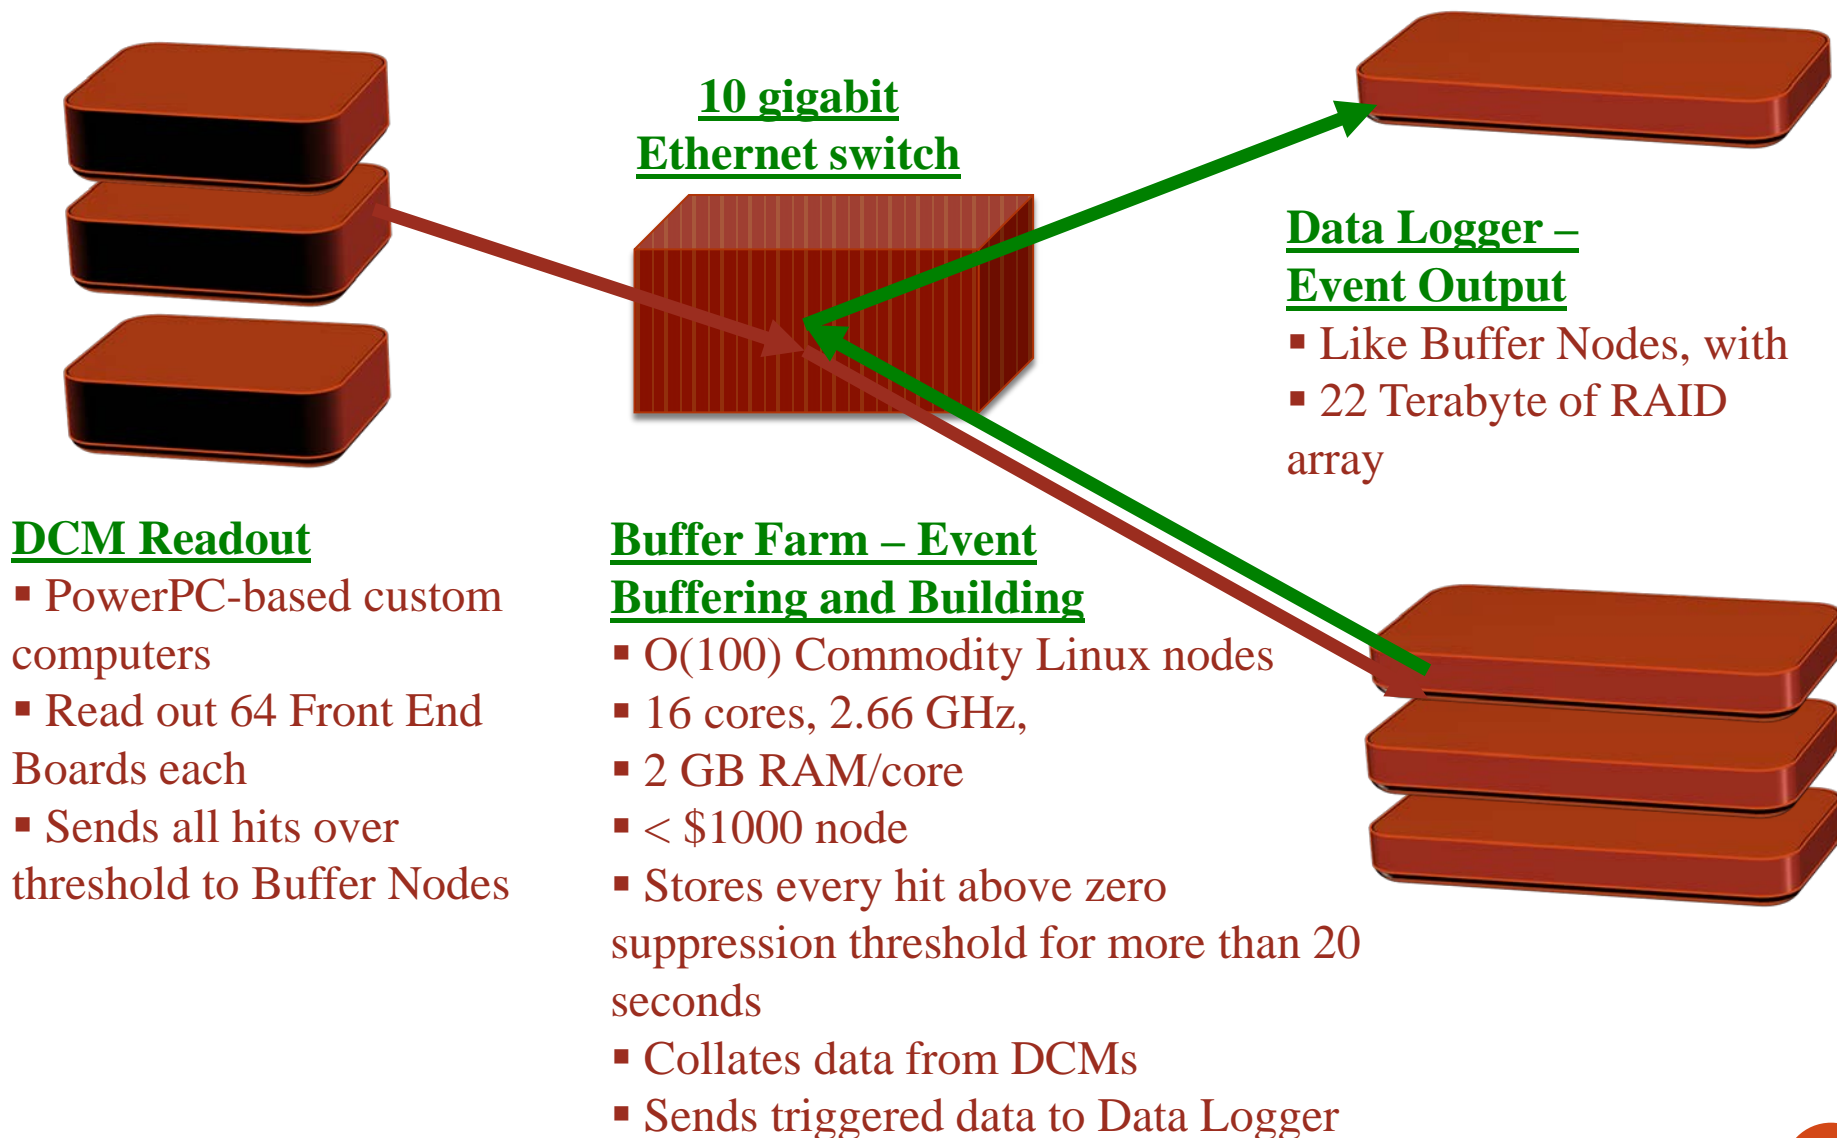
1 Diblock:

- 12 DCMs mounted on detector
- 6 per view
- Each DCM reads out a 2-module-wide slice

DAQ Cluster (Buffer Farm and Control Nodes)

- Ash River – Computing center in Detector building
- Fermilab – Lattice Computing Center

NOvA DAQ Architecture



NOvA Trigger Architecture



Trigger Inputs


- Neutrino beam spills times
- Random/Periodic gates
- Data-driven triggers



Trigger Outputs

- $(T_0, T_0 + \Delta T)$ trigger window
- Sent as software message to Buffer Farm and Data Logger
- No Hardware trigger in NOvA!

10 gigabit Ethernet switch



Buffer Farm

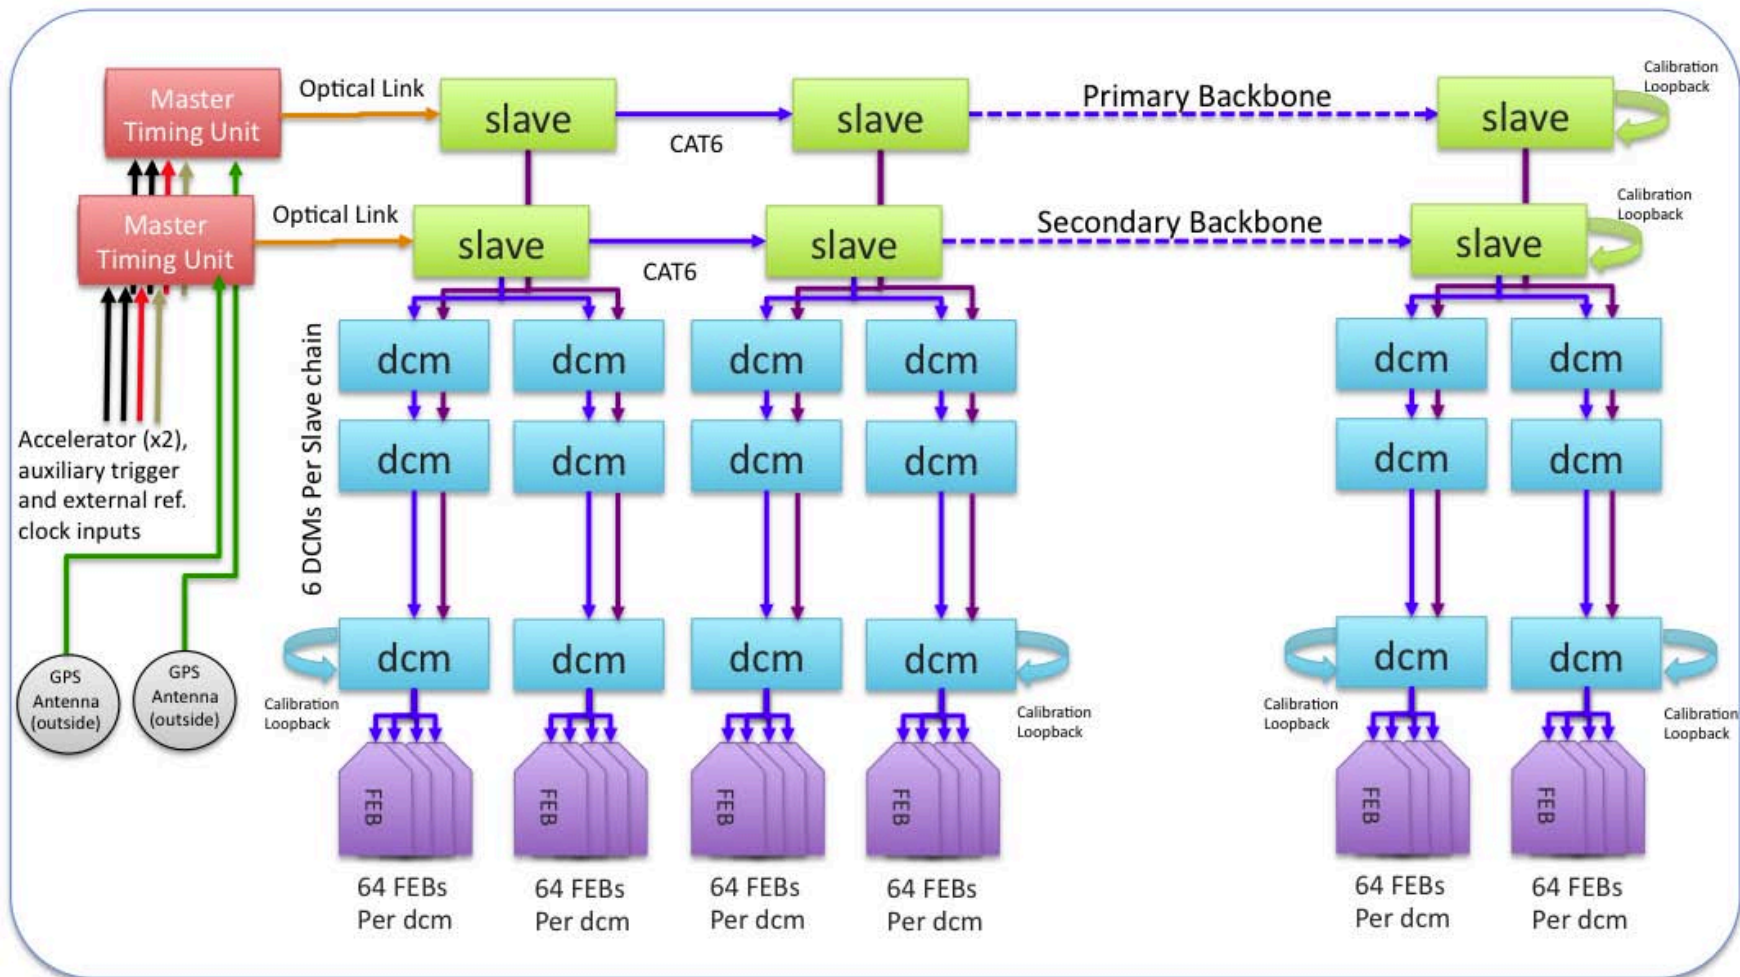
- All Data touching trigger window is sent to Data Logger
- Data-driven trigger:
 - Desirable cosmic events
 - Neutrinos
 - Supernovae
 - Exotic signatures

Data Logger

- Receives trigger as cross-check on data from Buffer Farm



NOvA Timing System




NOvA Clock:

- 64 MHz clock represented in 56 bits (24+32), covers 36+ year span with 15.625 ns ticks
- All FEB/DCM/TDUs are synchronized via GPS clock

DDT Design Overview



11520 FEBs
(368,4600 det. channels)



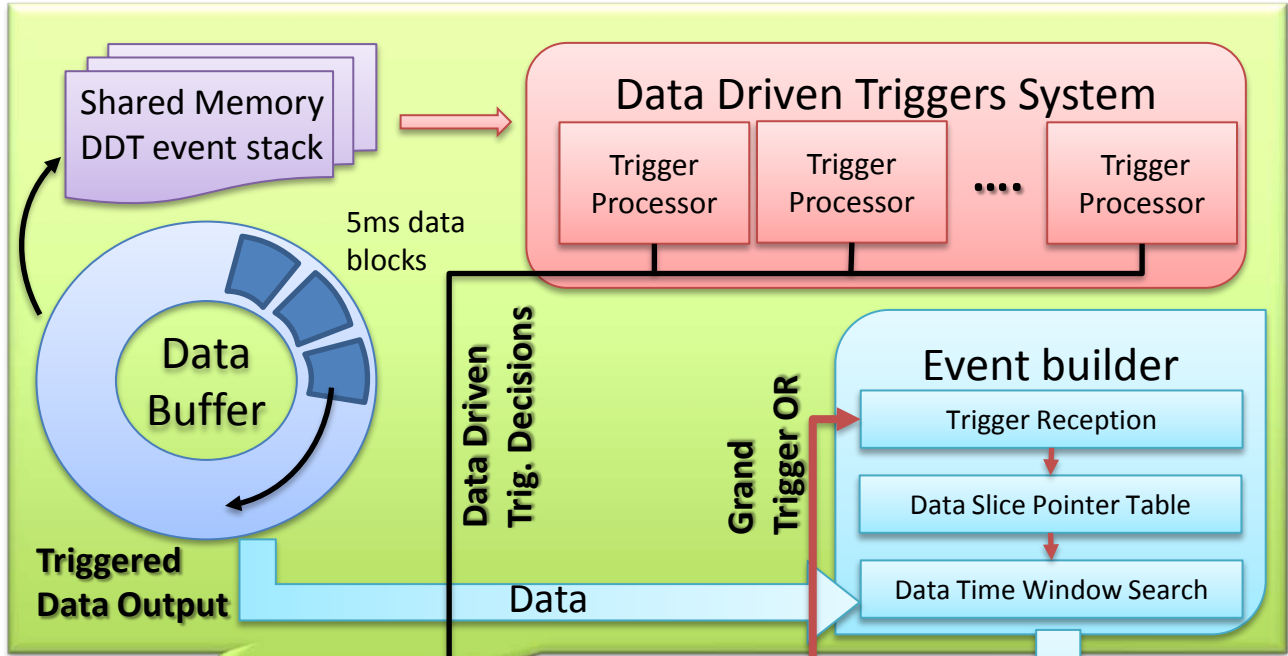
Zero Suppressed
at $\frac{1}{2} - \frac{2}{3}$ MIP
(6-8MeV/cell)

Minimum Bias
0.75GB/S Stream
DCMs
CoTS Ethernet 1Gb/s

200 Buffer Nodes
(3200+ Compute Cores)

Buffer Nodes

Beam Spill Indicator
(Async from FNAL @ .5-.9Hz)
Calib. Pulsar (50-91Hz)



Triggered
Data Output

5ms data
blocks


Data Driven
Trig. Decisions

Data

Grand
Trigger OR

Trigger
Broadcast

Global Trigger
Processor




Data Concentrator Module



“Kernel module”

Serves FPGA data to CPU as device driver

64 FEBs / DCM



56-bit 64 MHz time marker
Roll over every 36+ years

Microslice Header
Time Marker Low
Time Marker High
Nanoslice 0 FEB 0
Nanoslice 1 FEB 0
Nanoslice 0 FEB 2
Nanoslice 0 FEB 2
Nanoslice 0 FEB 3

Hit data: Nanoslice (500 ns)

Serial FEB-DCM link limit:
~ 15 hits/FEB/50 μ s

DCM Application

- DAQ system interface for control of DCM and FEB settings
- Monitors data quality
- Produces “Microslices” – data between consecutive time markers
- Builds Millislice for transfer to buffer node: collection of all Microslices within 5 ms

31	29	28	26	25	21	20	16	15	9	8	3	2	1	0
DFT	SEQ	Version	Pixel Adr	FEB Status	FEB ID	D	D	L	M	P	S			
Nanoslice Header														
Timestamp counter														
Pulse Height					Reserved									

Buffer Nodes / Data Logger



Round-robin input

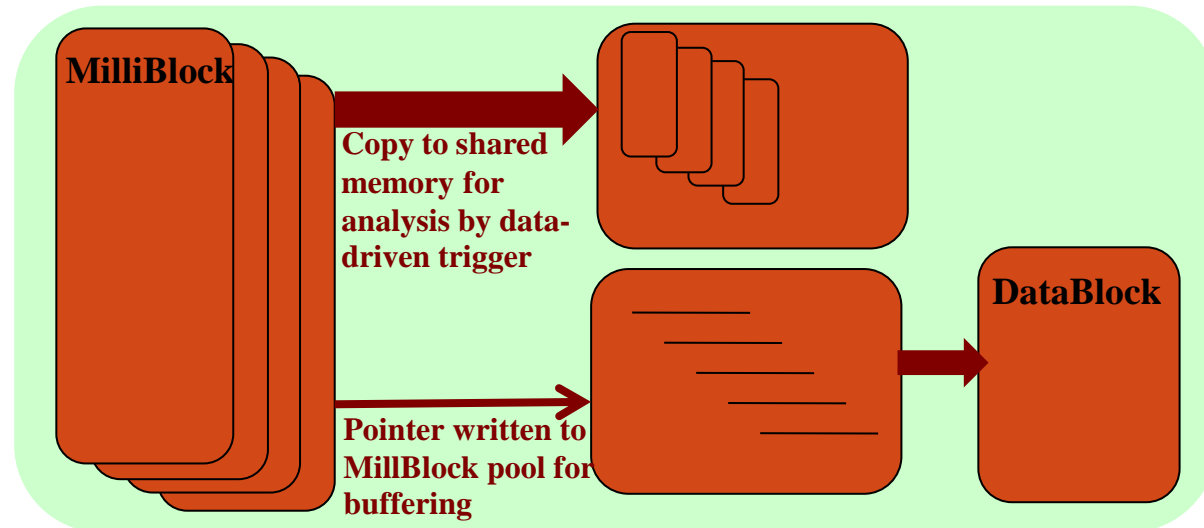
- All DCMs send to the same Buffer Node for a give 5ms period (sequence number)
- With 100 buffer nodes, this means a buffer node is used twice each second

MilliBlock

- 1st Millislice seen in a sequence causes creation of new Milliblock
- Complete when last DCM with sequence number is seen

Trigger data search:

- Create a DataBlock for output to DataLogger
- Check all MilliBlocks in buffer
- Does MilliBlock time overlap with trigger time? Check MicroBlocks:
 - Copy all MicroBlocks touching Trigger Window into DataBlock



Event (“Data Atom” until complete)

- Receipt of Trigger or DataBlock with event no
- Header, Trigger Info, DataBlocks, Tail
- Complete when data received from Trigger and all Buffer Nodes
- Note: The trigger, not the hit data, defines the event. The same data can appear in multiple events.


Output

- Send data to configurable set of streams (files on disk) based on trigger ID

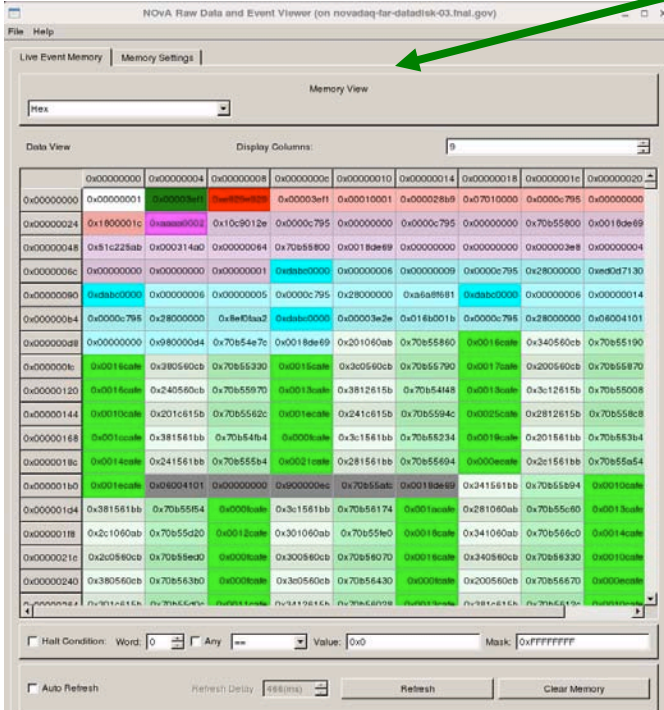
Data Monitoring / Applications



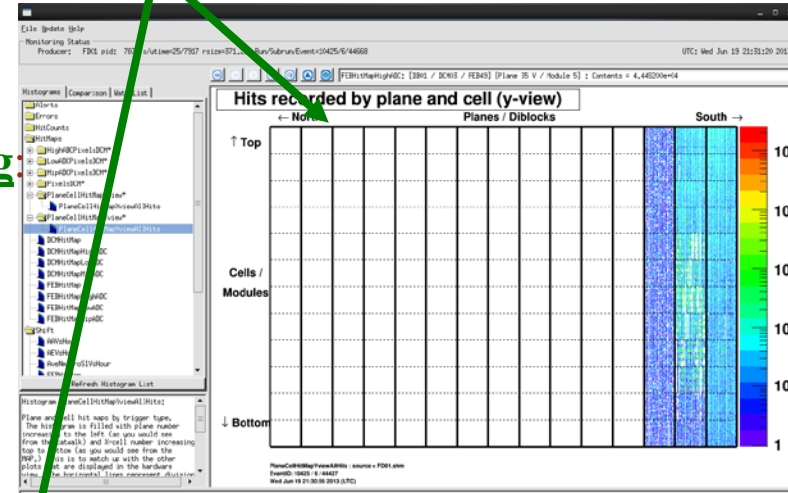
Data Logger Node:



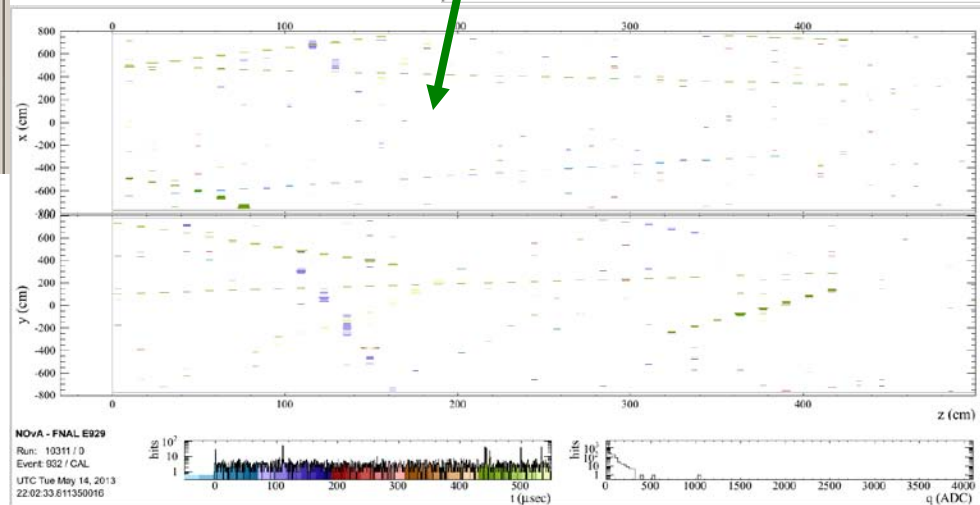
Event Dispatcher:
servers data to
monitoring applications



Online monitoring:
histogram
producer/viewer



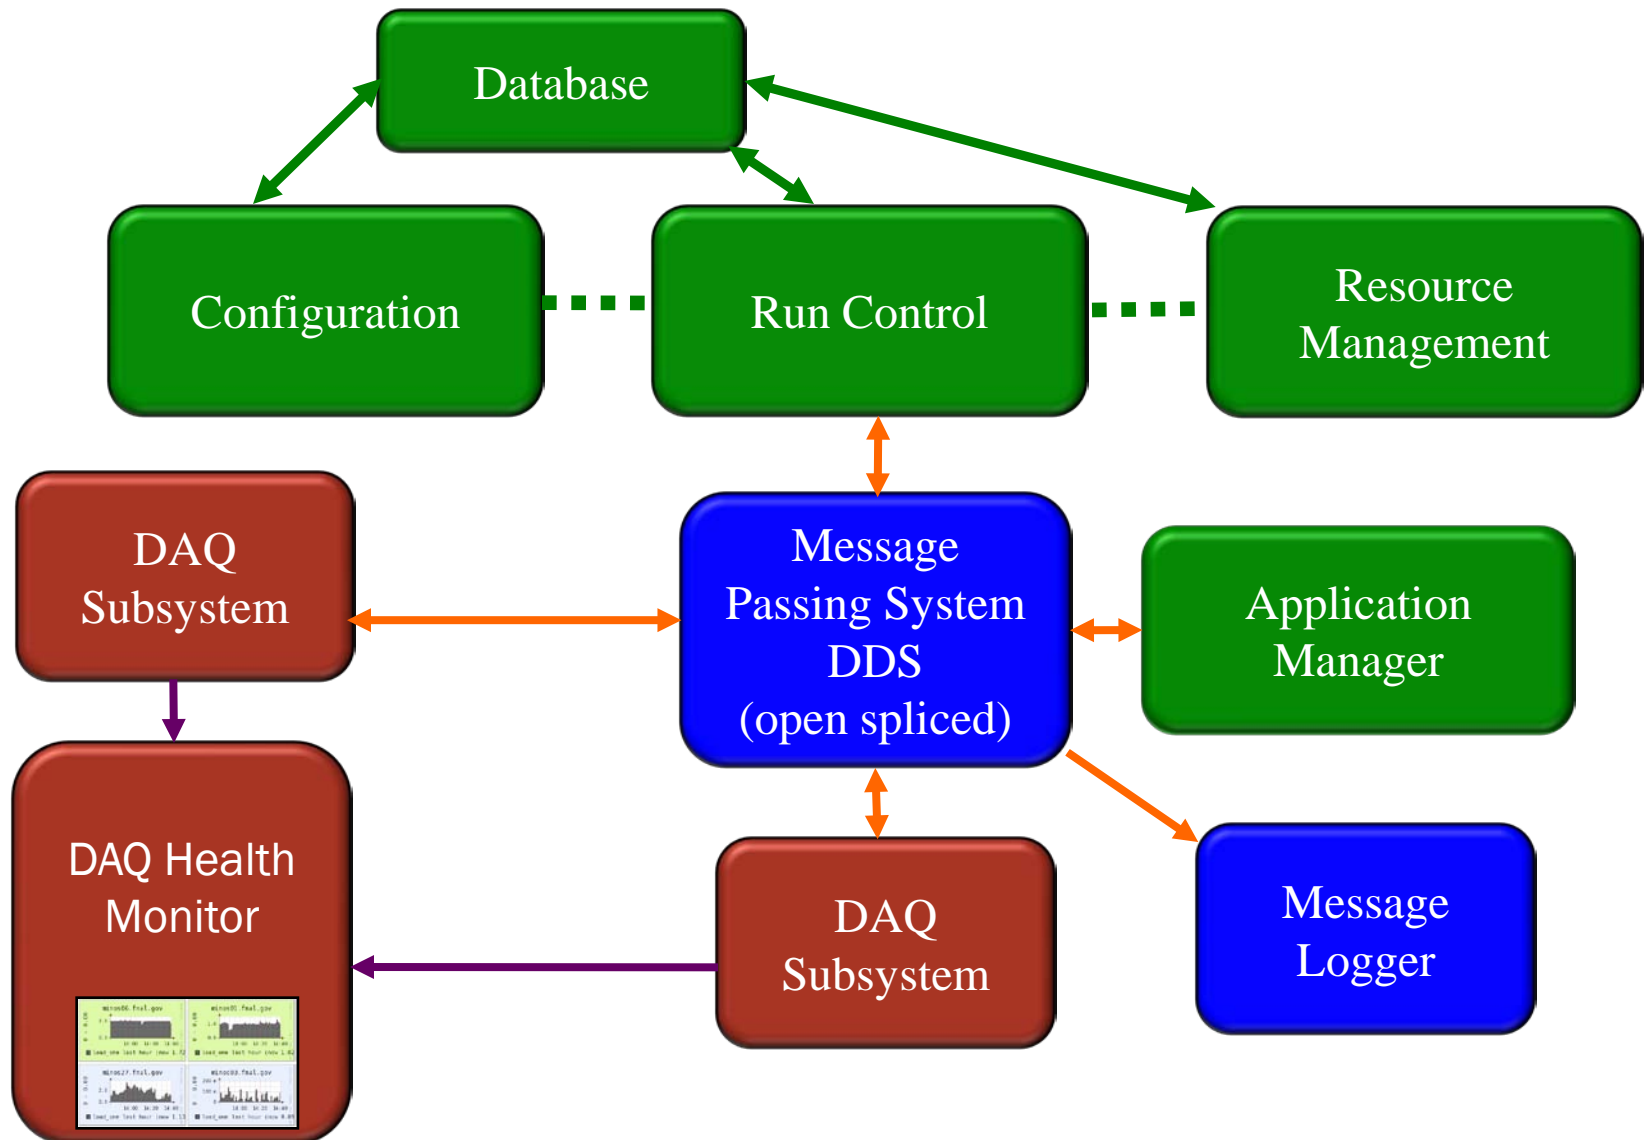
Event Display:



Memory Viewer:

- Displays raw data with highlighting to indicate data block type

DAQ Control & Monitor Systems



Summary



- ❑ **Far Detector commissioning is in progress and Near detector construction has begun.**
- ❑ **NOvA is achieving its DAQ design goals with a commodity online buffer farm and commodity network.**
- ❑ **We have built the foundation for a unique free-running DAQ system**
- ❑ **We understand how to perform continuous readout and buffering**
- ❑ **We have solved the problem of correlation and synchronization → see talk by Evan Niner**
- ❑ **We are exploring the power and opportunities of Data Driven Triggering → see poster by Zukai Wang**

BackUp

Resource Management



Resource Manager

- Keeps track of which resources are in use
- Resource ~ a DCM or other node

Partition

- A totally independent DAQ system
- One partition can be used for taking physics data on an established part of the detector, while another is used for checkout and commissioning

FarDet Detector Resource Manager (on novadaq-far-master)

File Edit Connection Configuration

Null Partition Partition 1

Id	Host	App State	RC State	Port
+ Managers				
- bNEVB Lists				
- bng03				
bnevb021	novadaq-far-farm-21	Enabled	LoadingRunC...	7555
bnevb023	novadaq-far-farm-23	Enabled	LoadingRunC...	7555
bnevb022	novadaq-far-farm-22	Enabled	LoadingRunC...	7555
bnevb024	novadaq-far-farm-24	Enabled	LoadingRunC...	7555
bnevb025	novadaq-far-farm-25	Disabled	CfgSelected	7555
bnevb026	novadaq-far-farm-26	Enabled	LoadingRunC...	7555
bnevb027	novadaq-far-farm-27	Enabled	LoadingRunC...	7555
bnevb028	novadaq-far-farm-28	Enabled	PreparingNew...	7555
bnevb029	novadaq-far-farm-29	Enabled	LoadingRunC...	7555
bnevb030	novadaq-far-farm-30	Enabled	LoadingRunC...	7555
+ bng02				
+ bng01				
- Timing Chains				
+ db01s				
- db01t				
dcm-2-01-01	dcm-2-01-01	Enabled	PreparingNew...	
dcm-2-01-02	dcm-2-01-02	Enabled	PausingRun	
dcm-2-01-03	dcm-2-01-03	Enabled	PausingRun	
dcm-2-01-04	dcm-2-01-04	Enabled	PausingRun	
dcm-2-01-05	dcm-2-01-05	Enabled	PausingRun	
dcm-2-01-06	dcm-2-01-06	Enabled	PausingRun	
+ db02s				
+ db02t				
+ db03s				
+ db03t				

Run Control

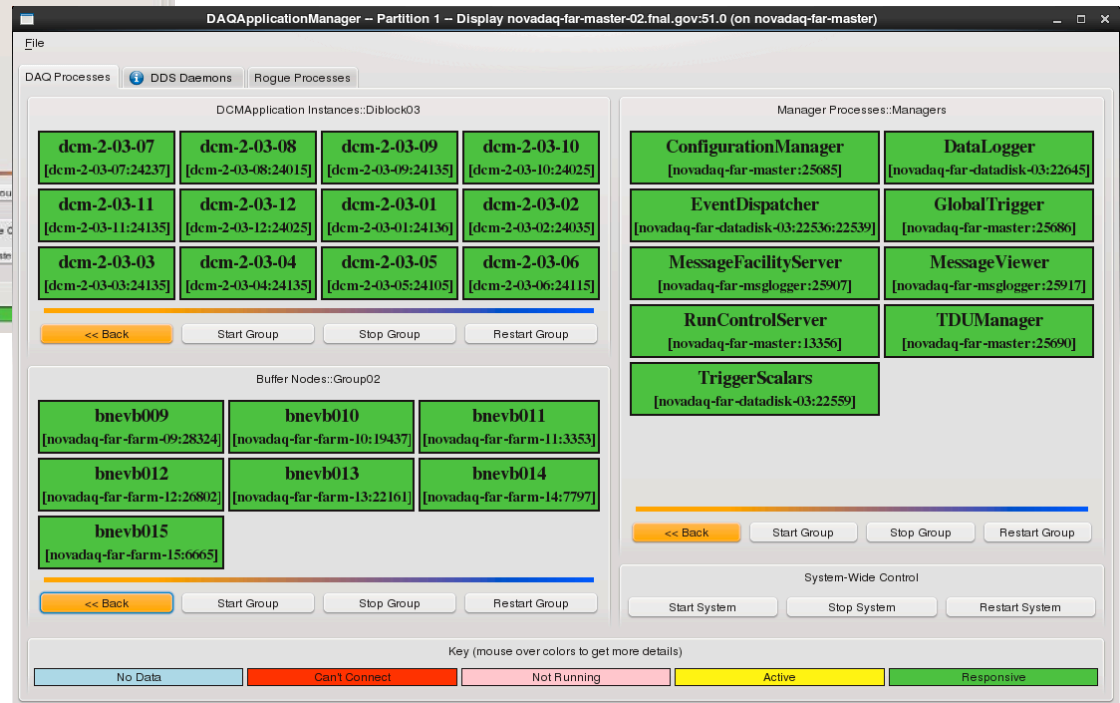
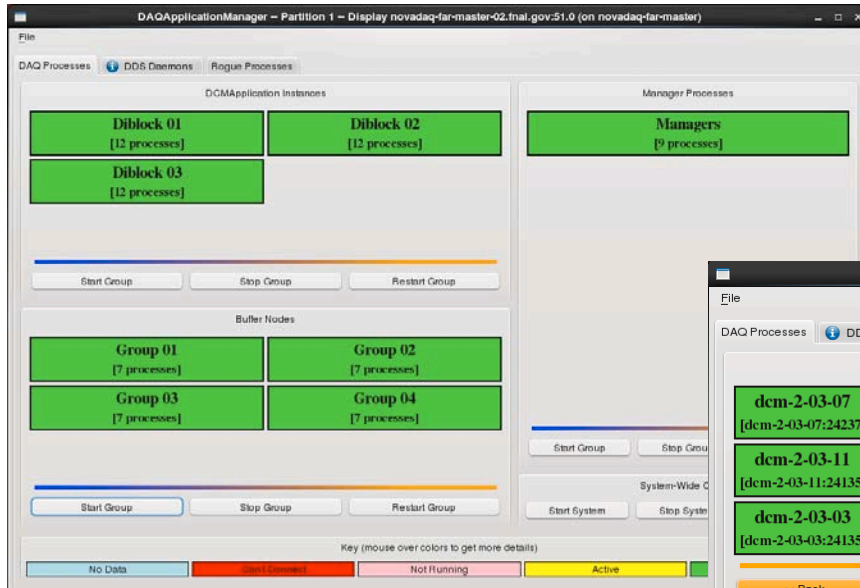


The image shows two windows from the FarDet Detector Run Control interface. The left window, titled "FarDet Detector Run Control, Partition 0 (on novadaq-far-master)", contains a menu bar (File, Configuration, Connections, View, Help) and a grid of buttons for resource management: Recover Resources, Select Resources, Reserve Resources, Release Resources, Select Configuration, Prepare Configuration, Load Connections, Make Connections, Load Hardware Config., Configure Hardware, Load Run Config., Configure Run, Begin Run, Pause Run, and End Run. Below the buttons is an "Execute command:" field. The status section shows "RC Server Status: Connected" and "RM Server Status: Connected". Progress indicators for "Run: 0" and "Subrun: -1" are at 0%. The "Num. Events" is 0, and the "Run Type" is set to "Commissioning". The "Shifter" is "Tian Xin". A log window at the bottom shows messages: "You do not have control over RC server.", "Control taken by Tian Xin at 127.0.0.1:40703", "Executing RC transition WaitingToDiscovering", and "Executing RC transition DiscoveringToDiscovered".

The right window, titled "Resource Selection Window (on novadaq-far-master)", displays a table of resources with columns for "Id", "Disabled", "Acknowledged", and "Group Acknowledged". The "Id" column lists various components, with "Timing Chains" expanded to show sub-items like "db04s", "db04t", "db05s", "dcm-2-05-07", "dcm-2-05-08", "dcm-2-05-09", "dcm-2-05-10", "dcm-2-05-11", "dcm-2-05-12", "db05t", and "db06s". The "db05s" and "dcm-2-05-10" rows are checked in the "Id" column. The "Acknowledged" column shows "New" for several items. The "Group Acknowledged" column has "All changes Ok" for some groups. At the bottom, there are buttons for "Select All", "Deselect All", "Acknowledge All Deviations", "Do Pedestal Run", "OK", and "Cancel".

- User selects resources for partition via Run Control
- Once applications are started, Run Control is used to execute configuration steps, and start/stop run

Application Management



Application Manager:

- starts/stops/monitors applications and Message Service Daemons

Message Facility



Message Facility – Log Messages:

- Used in offline ART framework
- For Online, has Server destination using DDS backbone
- Message rate throttling is run-time-configurable

Message Analyzer:

- Configured to recognize patterns of errors, e.g., many DCMs complain about the same buffer node
- Integration to Automatic Error Recovery actions

MessageFacility MsgViewer (on novadaq-far-msglogger.fnal.gov)

NOVA Message Viewer Total received messages: 43753 @ NULL_PARTITION

Message Filter: rcServer / rcServer / MF-online
dcm-2-01-12 skipped 11 heartbeats!

Categories: BNEVMilliBlock, BufferNodeEVB_c, ConfigurationMans, DLRCC, DLRCC_cppLog, DataBlockReader, DcmControl

Hosts: dcm-2-01-01, dcm-2-01-02, dcm-2-01-03, dcm-2-01-04, dcm-2-01-05, dcm-2-01-06, dcm-2-01-07, dcm-2-01-08

Applications: ConfigurationMans, DataLogger, NGTMaster, NovaGlobalTrigge, TDUControl, bnevb001, bnevb002, bnevb003

Log messages include: WARNING / RunControl, INFO / DcmMonitorTimingHistory, and DCM Write on Partition 0.

NOVA Message Analyzer Total parsed messages: 368333

System Status

Node Status Rule Engine Indicator

Rules	Details	Description	Expression
1 sync	Check DCM sync	ALARM	
2	DCM data corrupt (empty millislices)		
3	negative time jump in dcm		
4	farm node connection failure		

Base conditions	Details	Description	RegExpr
1 c1	dcm-.*	DCM data corrupt (empty millislices)	0
2 c2	dcm-.*	negative time jump in dcm	0
3 c3	dcm-.*	dcm lost connection with farm nodes	0
4 dcm_ti...	dcm-.*	Capture DCM timing ping messages	25013

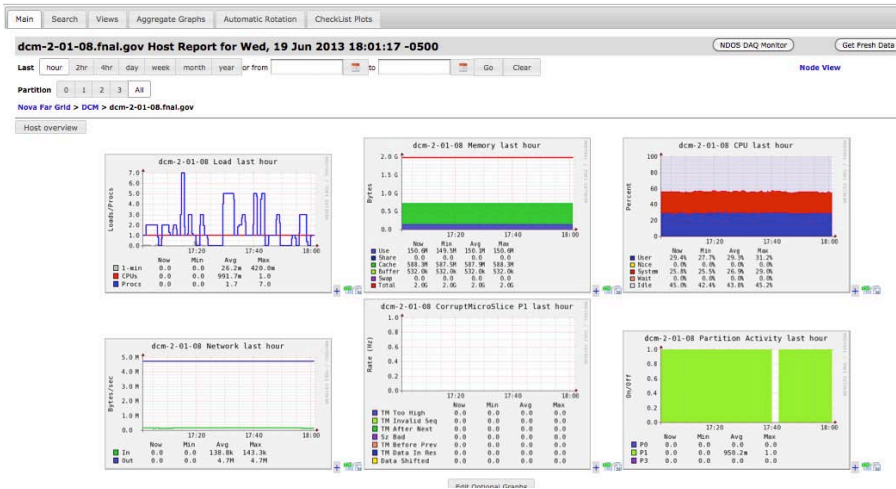
Message Log Analyzer

Click open to select an archived log file

0% Open Start

Reset All OK Cancel

NOvA DAQ Monitoring



Ganglia:

- Widely-used open-source cluster monitoring package
- Daemon on each host sends to central server
- Web-based display of plots of metrics vs. time for “Grid”, “Cluster”, and nodes
- Default metrics: CPU use, free memory, etc.
- Extensible to custom metrics

Nova DAQ Monitor

- System for generating warnings and alarms when metrics vary outside configured limits

