Online Data Challenge 2013

Primary participants: *Elliott Wolin, Sean Dobbs, David Lawrence*

When: August 26 – 29, 2013

Where: Hall-D Counting House

Objective: Test data flow and monitoring between final stage Event Builder (EB) and tape silo *(i.e. neither the DAQ system nor the offline were included)*

Input Data

- Pythia-generated events simulated, smeared, and passed through L1 event filter*
- Events digitized and written in EVIO format
  - *mc2coda* library used to write in the new event building scheme specification provided by DAQ group
  - Translation table derived from Fernando's spreadsheet detailing the wiring scheme that will be used

**event filter may have used uncalibrated BCAL energy units, but resulted in roughly 36% of events being kept.*

Computers systems

(many of these on loan)

n.b. all L3 machines connected via InfiniBand

node	type	cores	RAM	function
gluon01a gluon02 gluon03 gluon04 gluon05	3.4GHz Intel i5	2+2ht	16GB	console (human interface)
gluon20 gluon21 gluon22 gluon23	1.9GHz AMD Opteron	8	8GB	monitoring
gluon40 gluon41 gluon42 gluon43	1.9GHz AMD Opteron	8	8GB	L3 trigger & monitoring
gluon44 gluon45	2.53GHz Intel Xeon	8+8ht	48GB	L3 trigger, event source, & monitoring server
gluon100 gluon101 gluon102 gluon103	1.9GHz AMD Opteron	8	8GB	L3 trigger
halldraid1	2.0GHz Intel Xeon	4+4ht	12GB	RAID disk

L3 and monitoring architecture

for 2013 Online Data Challenge

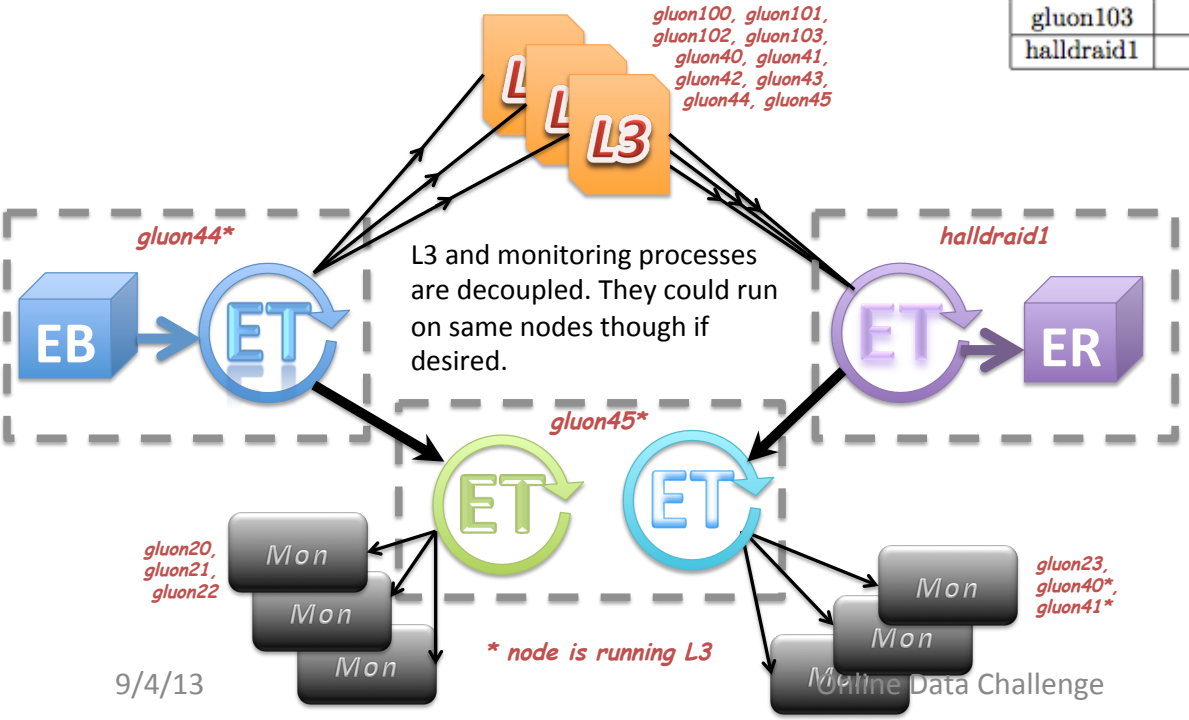

L3 Infrastructure Test

- 10 nodes used to pass events from Event Builder (EB) to Event Recorder (ER)
  - EB on *gluon44*, ER on *halldraid1*
- Two "pass-through" modes used:
  - Simple buffer copy without parsing (40kHz)
  - Buffer copy with parsing and application of translation table (~13kHz)
- DL3TriggerBDT algorithm from MIT
  - Worked in single threaded mode, but crashed with multiple threads
  - Cause of crashes unknown and under investigation
    - MIT and TMVA code itself have been eliminated as causes

Screen Capture of ET systems monitors

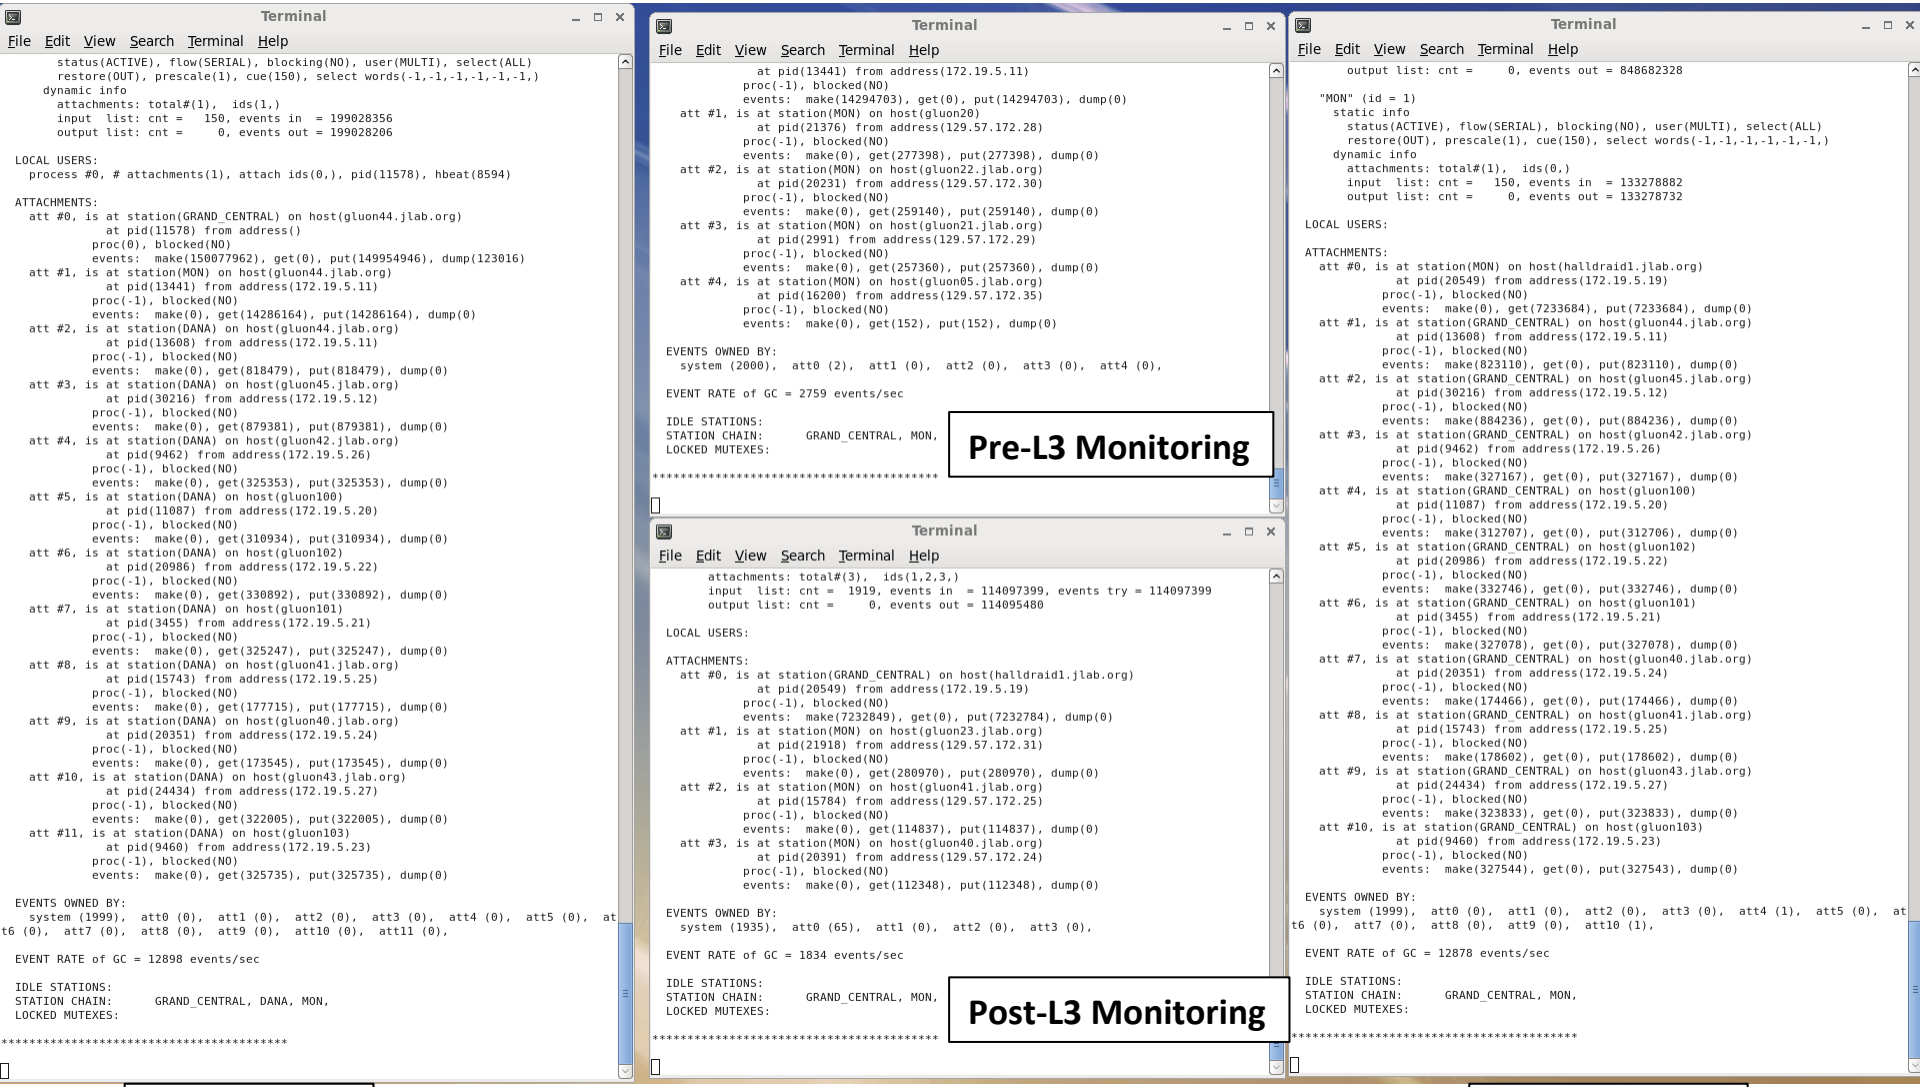

Monitoring System (RootSpy) test

- Histograms produced by several plugins were displayed via the RootSpy GUI
- Overlay with archive histograms
- RootSpy archiver (writing summed histograms to file)
  - Integration with CODAObjects
  - Still need to fully implement *final histograms* mechanism

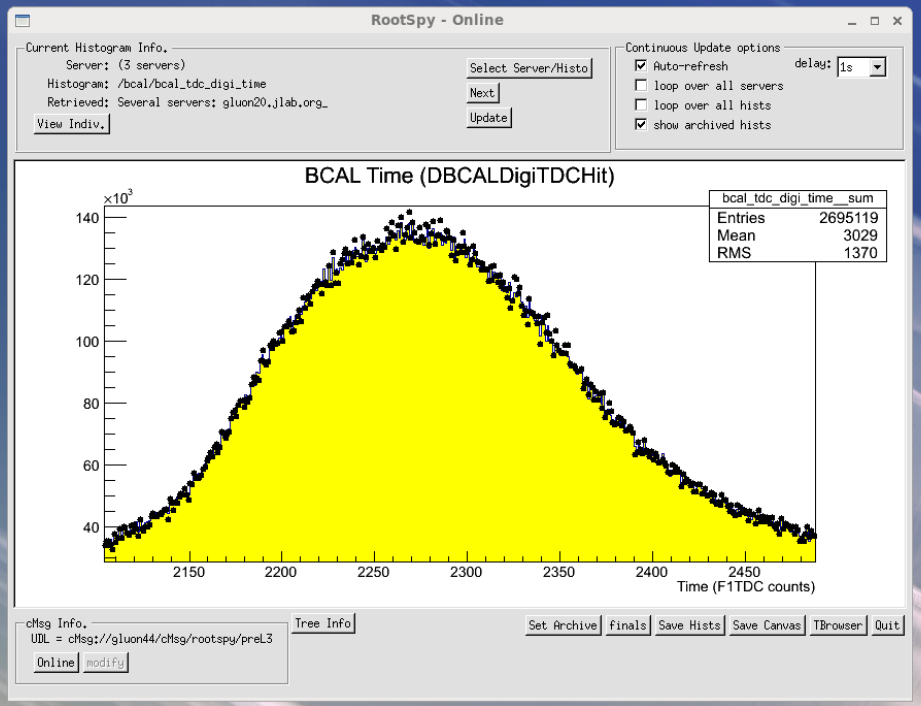

Pre-L3 Monitoring

Post-L3 Monitoring

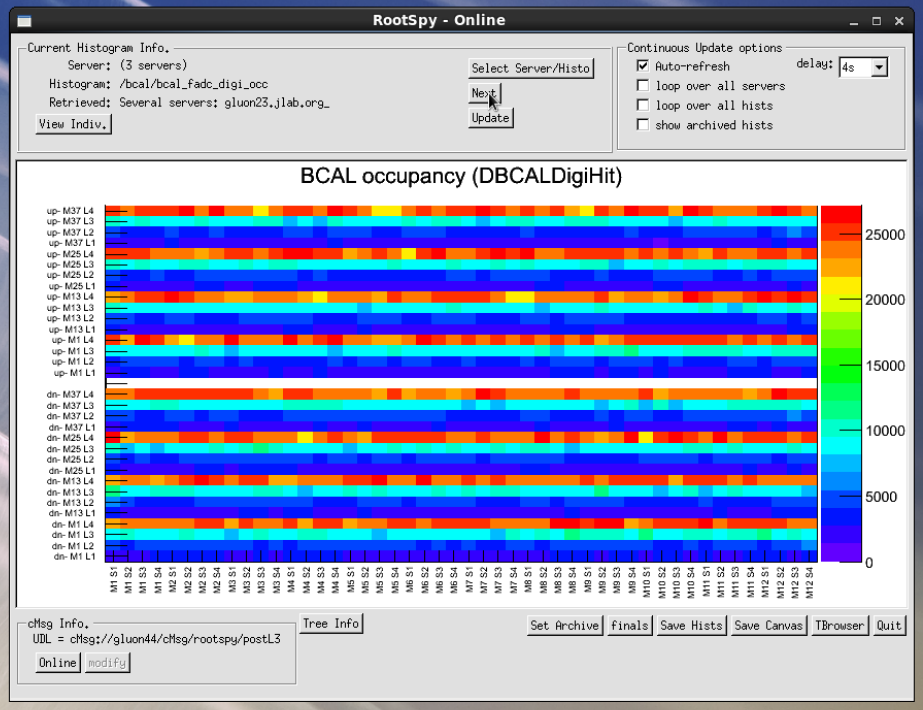

Other monitoring

- Ganglia installed and working for monitoring general health of all computer nodes

```
HDOPS gluon04:~> janactl -t 0.1 thinfo
Sent command: get threads to janactl

Threads by process:
-------------------------------------------------
gluon100.jlab.org_347:
   thread: 0 0x7f681e1fc700  4590 events  76.0909Hz (45.5811Hz avg.)
   thread: 1 0x7f681ebfd700  4667 events  71.113Hz (46.3516Hz avg.)
   thread: 2 0x7f681f5fe700  4295 events  64.6926Hz (42.6587Hz avg.)
   thread: 3 0x7f681ffff700  4441 events  73.4932Hz (44.1011Hz avg.)
   thread: 4 0x7f6834dfa700  4797 events  80.417Hz (47.6384Hz avg.)
   thread: 5 0x7f68357fb700  4735 events  46.3996Hz (47.3674Hz avg.)
   thread: 6 0x7f68361fc700  4635 events  74.3584Hz (46.026Hz avg.)
   thread: 7 0x7f6836bfd700  4596 events  75.2129Hz (45.6399Hz avg.)
gluon101.jlab.org_26978:
   thread: 0 0x7f40ea1fc700  4722 events  25.347Hz (46.889Hz avg.)
   thread: 1 0x7f40eabfd700  4762 events  78.416Hz (47.282Hz avg.)
   thread: 2 0x7f40eb5fe700  4602 events  33.9704Hz (45.6912Hz avg.)
   thread: 3 0x7f40ebfff700  4843 events  60.7648Hz (48.0869Hz avg.)
```

RAID to Silo test

- Transfer from RAID disk to silo tested
  - At least 50MB/s achieved, but possibly higher
  - Certificate and jput set up, but we were informed later that a different mechanism should be used for experimental data from the halls
  - Will arrange for IT division experts to come run tests and educate us on the proper way we should be transferring to the silo

Summary

- EB to ER Data flow piece tested
  - L3 infrastructure tested and works in pass-through mode at 40kHz (mysterious issues with L3 plugin still being tracked down)
- Monitoring system tested
  - Identical pre-L3 and post-L3 monitoring systems
  - RootSpy GUI used with multiple producers
  - RootSpy archiver
- RAID to Tape silo tested
  - Successfully transferred > 1TB from counting house to silo at >= 50MB/s
  - Rate seemed slower than anticipated by factor of 2, but measurement mechanism not accurate due to staging
  - Alternate transfer method has been advised and will be pursued